VISUAPPS Viewer

User Manual

Date 12.01.2024
Build 3256

▰	VISUAPPS GmbH Suchenwirtplatz 3/22 1100 Vienna
SN	Build number 3256
⩘	12.01.2024
⧖	12.01.2029
⚠	Please read sections with this warning carefully.
📖	Please read the manual.
👤#	This product may display a unique patient identification.

CONTENTS

1

USAGE CAUTIONS

The usage of VISUAPPS Viewer for diagnostic decision support is only allowed for trained personnel. VISUAPPS GmbH offers training courses for appropriate usage.

Diagnostic readings with the VISUAPPS Viewer are only allowed on medical display devices certified for diagnostic usage. The operator is responsible to maintain monitor calibrations according to the display specification.

If the application displays the warning: Image viewing only, not for diagnostic purpose, diagnostic decisions must not be performed based on this software.

Interactions with the image via mouse or the context menu are supposed to be interactive. If the application is not responding to user interactions, this may be an error condition and no diagnostic reading must be performed.

The correctness and visibility of the *Patient Demographics* has to be checked every time an image is being displayed for reading. Missing or wrong patient demographics are an indication of a software error condition. No readings must be performed in this situation.

If the lines in the DVH Rendering are all white, no linkage to a valid *RTStruct* is available. No diagnostic reading must be performed in this case.

When switching to *True Size Zoom* mode or performing geometric monitor calibration the user has to make sure that the browser zoom is set to 100%.

When calibrating the image's pixel spacing, the real-world calibration reference has to be placed parallel to the modality's projection plane. Furthermore it has to be in plane with the pathology of interest. If these requirements are not fulfilled, no calibration is possible. The calibration tool has to be removed manually in this case, as the pixel spacing is not valid.

2

TERMS OF USE

The purpose of the VISUAPPS Viewer is to display medical data from instrumental diagnostic modalities. The functionality comprises planar and volumetric displays including simple measurement methods. The intended use of this product is to support the doctor in the readings and decision finding process. This product does not perform any diagnostic decision finding.

> **The usage of VISUAPPS Viewer for diagnostic decision support is only allowed for trained personnel. VISUAPPS GmbH offers training courses for appropriate usage.** ⚠

Primary readings with the VISUAPPS Viewer are exclusively permitted on a calibrated monitor, certified for diagnostic purposes and adequate ambient light conditions. The operator is liable to fulfill the local legal regulations for diagnostic reading conditions.

> **Diagnostic readings with the VISUAPPS Viewer are only allowed on medical display devices certified for diagnostic usage. The operator is responsible to maintain monitor calibrations according to the display specification.** ⚠

> **If the application displays the warning: Image viewing only, not for diagnostic purpose, diagnostic decisions must not be performed based on this software.** ⚠

3

USER INTERFACE

The VISUAPPS Viewer presents a visual environment for investigating imaging data. The user interface consists of two main parts: the *Commander* and the *Image Area*. The *Commander* is a side-bar on the left side of the application. It provides information on the data currently loaded and access to various tools (see Section 3.1). The central area hosts the *Image Area*. In this area the images are displayed according to the *Hanging protocol*[1] (see Section 3.2).
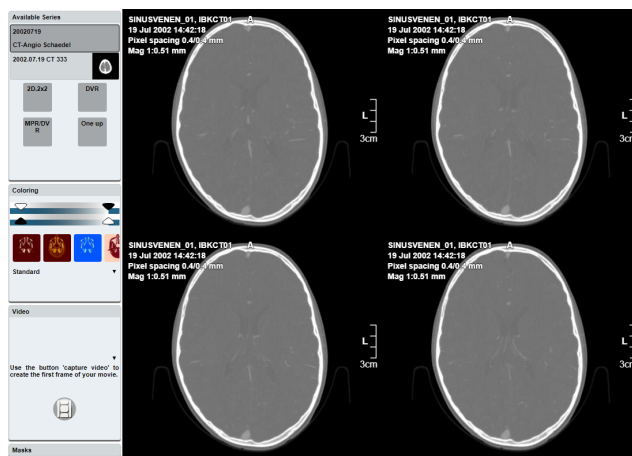

Figure 3.1: The Application User Interface

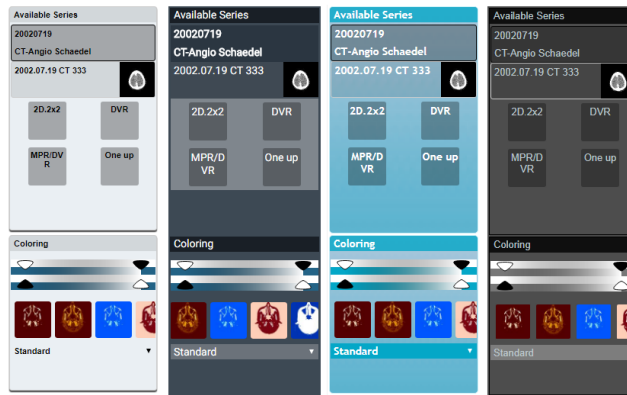[1] ftp://medical.nema.org/medical/dicom/final/sup60_ft.pdf

Figure 3.2: The Commander in various visual styles

3.1   Commander

The *Commander* hosts modules, which provide functionality depending on the the currently displayed data. More precisely the *Hanging protocol* defines the set of available modules. Limiting the available functionality to an applicable subset avoids the user interface to become overloaded and confusing.

The visual appearance of the *Commander* may vary according to the used style (see Figure 3.2). The style of the VISUAPPS Viewer is usually adopted to visually match the encapsulating system. Even though the visual appearance differs, the functionality and user interaction remains the same for all styles. Hence, this document will consistently use one style without loss of generality.

From top to bottom, the commander contains the modules in panels, which can be minimized interactively. A single click on the title opens and closes the module. The *Hanging protocol* may specify modules to be minimized by default if they would be applicable only in rare cases.

In the following the individual modules are described in more detail.

3.1.1   Available Series Module

The *Available Series Module* (also referred to as *Available Series Tray*) displays the available data structured in a DICOM *Study/Series* hierarchy (see Figure 3.3). The *Study Level* displays the study date and the study description. On the *Series Level* the first line presents the series data, the modality and the number of instances in this series. If present, the second line shows the series description. Next to the textual information a sample image from the series (*Preview Token*) is displayed. Each series contains a set of available *Hanging protocols*. The availability of *Hanging protocols* depends on the data and is defined in the *Hanging protocol*. The *Available Series Module* is used to select data to be displayed in the *Image Area*.

A single click on the *Study Level* label will expand or collapse the contained series. When clicking on the *Series Level* label, the associated *Hanging protocols* will
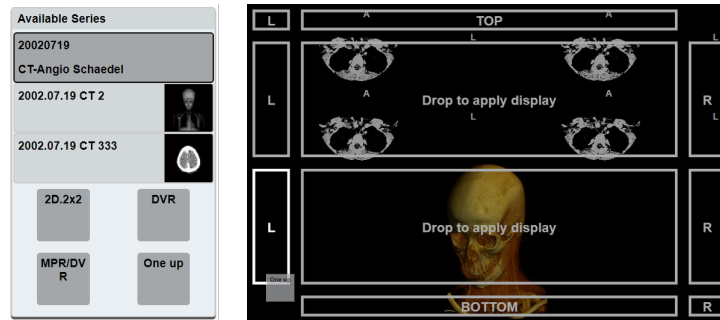
Figure 3.3: Available Series and Drag'n'Drop Interaction

be shown or hidden. A click on the *Hanging protocol* icon will replace the current content of the *Image Area* with the hanging, specified in the selected *Hanging protocol*. Dragging the *Hanging protocol* icon into the *Image Area* provides the possibility to append the selected hanging to the currently displayed content (see Figure 3.3). It is also possible to drag the *Series Level* label directly into the *Image Area*. In this case the *Hanging protocol* "One up" will be appended to the display content.

When dragging an item over the *Image Area*, a panel dims the displayed content and supplies placement information (see Figure 3.3). The drop panel provides options where to append or replace the content. When dropping the item on existing displays the current content will be replaced (option *Drop to apply display*). At the edges of the *Image Area*, options are available to append the new content. The option *TOP* will resize and move the existing displays towards the bottom and append the new content on top. The other options *L* (left), *BOTTOM* (bottom), and *R* (right) work accordingly. The corners provide options to use the entire row/column for the new content instead of just inserting it into the respective slot. In the example in Figure 3.3 this option is possible for *L* and *R*.

3.1.2 Coloring Presets Module

The *Coloring Presets Module* allows the user the application of coloring presets and fine-tuning of the currently applied preset (see Figure 3.4). There are three different sources for presets: *predefined presets*, *modality presets*, and *computed presets*.

Predefined presets are presets provided by the manufacturer. They may have been created for a certain purpose (e.g., Lung preset, Bone preset) or to allow the user access to certain functionality (e.g., MIP preset). Only presets suitable for the currently displayed data and rendering technique are listed.

Modality presets are provided by the modality and come along with the imaging data. If a preset name is available in the data this information is taken, otherwise the modality presets are named Standard(c/w) where *c* denotes
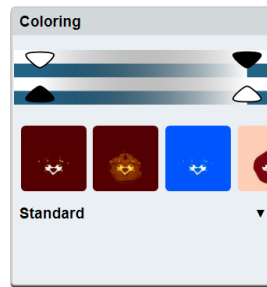
Figure 3.4: Coloring Presets

the center and w the width. If the modality presets were computed from a *Lookup Table* the term LUT is appended to the preset name.

Computed presets are available for the modalities CR (Computed Radiology) and MG (Mammography). Based on the given information a *Sigmoid*–interpolated preset is computed. This preset is named Sigmoid.

A visual representation of the preset is displayed in the upper part of the module. The depicted color transition represents the color lookup table, which is currently applied. The top–row handles define the opacity. The handles at the bottom define the color.

A new preset can be selected in two ways: The desired preset can be *(a)* selected from the drop-down list in the lower part of the module or *(b)* chosen by clicking on one of the rendered presets from the presets gallery. To modify the selected preset, the opacity and color handles can be re-positioned using the left mouse button.

3.1.3 Video Creation Module

The *Video Creation Module* provides the possibility to create a video sequence. This functionality is especially useful for the volumetric display, however, it may also be applied to the image display as well. The video sequence is defined by a series of *key frames* (see Figure 3.5). Each key frame represents a certain scene, which can be defined by modifying the display parameters (e.g., zoom, pan, window/level, etc...). The *Video Creation Module* will create a smooth transition from one key frame to the next and create a video out of it.

Multiple videos may be edited simultaneously. For each video the *Quality* and *Speed* parameter can be set. The *Quality* parameter provides the following options:

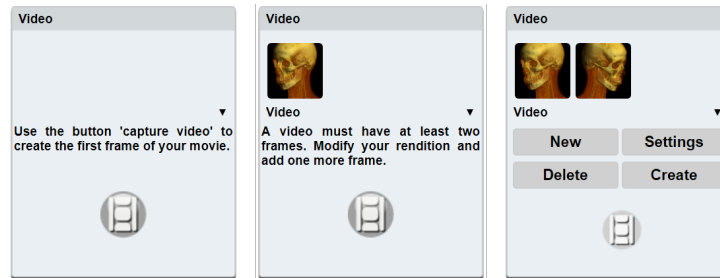| Resolution	Width	Height	Fps
YouTube	320	240	25
DVD	720	480	25
HD/Blue-Ray	1920	1080	30

Figure 3.5: Video Creation: Creating key frames

The *Speed* parameter modifies the replay speed of the video. The default video sequence duration from one key frame to the next is 3s (*Regular*). Setting the value to *Slow* stretches the duration to 5s and *Fast* reduces the duration to 1s.

Initially a wizard dialog is presented, as at least two key frames have to be defined for a video sequence. The user either clicks on the button in the wizard, or on the *Add Key Frame* button (see Section 5.18) in the *Context Menu* (see Section 3.2.3). Reordering the key frames in the key frame tray can be done by dragging the key frame into its new position using the left mouse button. Dragging the key frame onto the waste bin icon removes the key frame from the video sequence. The waste bin icon appears only while dragging a key frame. Clicking on the key frame will apply the key frame's display parameters to the current display in the *Image Area*.

The button *New* initializes a new, empty video sequence. The previous video sequence is still accessible via the drop-down box above. Video parameters (see above) can be modified for an existing video sequence via the *Settings* button. *Delete* removes the current video sequence. Finally *Create* creates a new video sequence and starts the download as soon as the video is created. **This may take up to several minutes!**

3.1.4 Mask Editor

The *Mask Editor* enables the modification of the mask's visual appearance. The default appearance of the masks is defined in the DICOM data. All available masks are displayed in the mask tray and are textually listed in the drop-down list. The mask's color can be modified with the color chooser (see Figure 3.6). Positioning the color handle to the leftmost position will disable the color contribution of the mask. The mask's opacity can be modified according to Figure 3.7: Positioning the slider to the leftmost position will hide the data values covered by the mask in volumetric displays and hide the mask overlay in image displays. Moving the slider towards the right side will gradually fade in the hidden structure. In the interval of the gray-scale ramp the mask's opacity increases in the image display. In the volumetric display the opacity is given by active preset's opacity. Continuing towards the rightmost end of the scale, the mask gradually becomes fully opaque
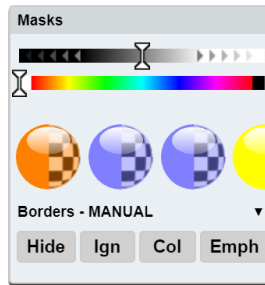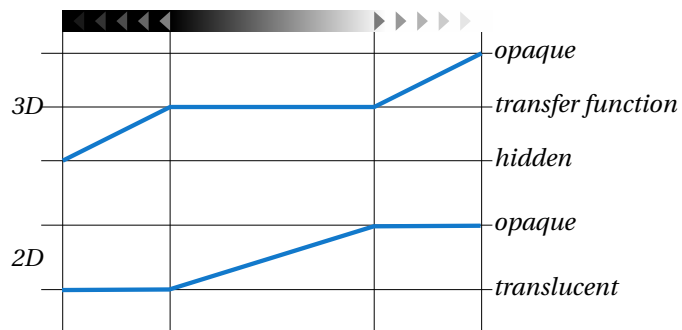
Figure 3.6: Mask Editor

Figure 3.7: The Opacity Parameter

in volumetric displays – independent of the data.

There are four presets for a quick access to commonly used applications:

Hide makes the data marked by the mask completely translucent (e.g., background suppression).

Ign will ignore the mask completely, just as if the mask would not exist at all.

Col colors the mask to mark certain structures.

Emph enhances the mask so that it becomes visible even though the data would be translucent.

To modify a mask, it has to be selected first. This can be done either with the mask tray or the drop-down list. The user can adjust the visibility setting – color or opacity – by dragging the corresponding handle with the left mouse button.

3.1.5   Sequence Splitting Module

The *Sequence Splitting Module* presents a list of image sets, as defined in the *Hanging protocol*. In contrast to the *Available Series Module* (see Section 3.1.1) the structure may differ from the DICOM *Study/Series* hierarchy. This is useful,
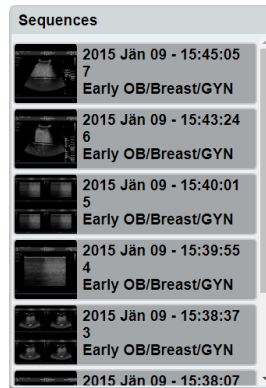
Figure 3.8: Sequence Splitting

especially if different content is put into the same series by the modality. A common use-case are ultrasound studies. This module is only available for some *Hanging protocols.*

The interaction in the *Sequence Splitting Module* is comparable to the interaction in the *Available Series Module.* A single click on one of the image sets will replace the displayed content in the *Image Area* with a "One up" hanging of the image set's data. When dragging a display set over the *Image Area* the drop panel appears and allows the user interactions like appending and replacing, as described in Section 3.1.1.

3.2 Image Area

The *Image Area* displays the selected images, allows the user manipulations of the image, and provides access to the context menu. The displays are organized in a tiled grid arrangement. Superimposed on each image, a set of annotations is displayed (see Section 3.2.2). All image display interactions operate with interactive response times.

> **Interactions with the image via mouse or the context menu are supposed to be interactive. If the application is not responding to user interactions, this may be an error condition and no diagnostic reading must be performed.**

The *Image Area* displays a single *Presentation Group* from the *Hanging protocol* at once. *Presentation Groups* are different hangings of the data, which usually correspond to the workflow of a certain medical reading procedure. If multiple *Presentation Groups* are available, the tool *Presentation Group Switch* is available (see Section 5.13).

3.2.1   Mouse Interaction

Within the scope of a display, a subset of frequently used tools are applicable directly using the mouse buttons. The actions are applied by dragging the corresponding button in the image. These actions can also be invoked while the *Context Menu* (see Section 3.2.3) is opened. The modifier key (e.g., the control key – ctrl) has to be pushed before pressing the mouse button.

Display	Modifier	Button	Action
IMG		Left	Pan
		Middle	Window/Level
		Right	Zoom
		Scroll	Navigation
	ctrl	Right	Navigation
		Scroll	Zoom
WAV		Left	Pan
		Middle	Window/Level
		Right	Zoom
		Scroll	Zoom
MPR		Left	Pan
		Middle	Window/Level
		Right	Zoom
		Scroll	MPR plane movement
	ctrl	Right	Navigation
		Scroll	Zoom
DVR, MIP SVR, HVR		Left	Rotate (3D)
		Middle	Pan
		Right	Zoom
		Scroll	Zoom

3.2.2   Image Overlays

Important image information is displayed along with the image. Patient and technical information is presented in the *Patient Demographics* areas. The orientation of the image – if present in the data – is indicated in the *Orientation* annotation. The current image scale is denoted in the *Ruler* markup.

Patient Demographics   The *Patient Demographics* and *Technical Attributes* are located in the corners of the images. Their visibility can be toggled using the *Demographics Visibility* tool (see Section 5.14). The content of this *Image Overlay* depends on the image's modality.

> ⚠ **The correctness and visibility of the *Patient Demographics* has to be checked every time an image is being displayed for reading. Missing or wrong patient demographics are an indication of a software error condition. No readings must be performed in this situation.**

Patient Name   depicts the patient's name.

Series Date   is the date of the displayed series.

Series Description   is displayed if existing.

Pixel Spacing   is the pixel spacing of the original data.

Magnification   denotes the magnification of the image. There are two options available:
Calibrated: If both, the image data and the display monitor are calibrated, the magnification is given in *mm* units. This means *1* mm in the image is displayed as *n* mm on the monitor, where *n* depends on the actual zoom. (`Mag 1:n mm`)
Uncalibrated: The magnification unit *px* is used. In this case *1* pixel in the image is displayed as *n* pixels on the screen (`Mag 1:n px`).

Orientation   The *Orientation* annotation indicates the patient's spacial orientation relative to the displayed data. The spacial orientation information has to be present in the data for this annotation to be displayed. The locations of the annotation are by default the center of the top edge and the center of the right edge in the display. The location may vary according to the applied *Hanging protocol*. The abbreviation denotes the patient's physical orientation closest to the corresponding edge (see Figure 3.9).

Abbr.	Direction	Abbr.	Direction
H	Head	F	Feet
L	Left	R	Right
A	Anterior	P	Posterior
U	Unspecified		

Ruler   The *Ruler* depicts the current image scale. By default it is located at the right side of the display, however, the location may vary according to the *Hanging protocol*. The image scale changes if the image magnification is modified. The displayed length corresponds to the actual length of the ruler in real world. Ticks are one unit distance apart from each other (see Figure 3.9).
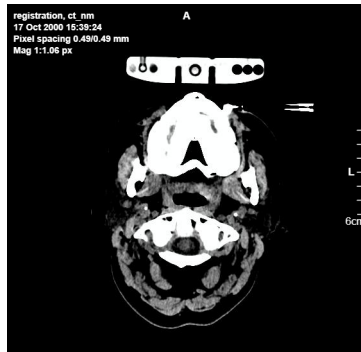
Figure 3.9: Orientation and Ruler

Figure 3.10: Rotation of the context menu

3.2.3   Context Menu

The *Context Menu* can be opened by a single click with the right mouse button into
the *Image Area* (see Section 3.2). The focus of the base section can be changed
by rotating the disk. Disk rotation is accomplished by dragging the gray area of
the disk with the mouse (see Figure 3.10). The selected disk will snap into place.
Alternatively, rotating with the mouse wheel is possible. The mouse pointer has to
be within the *Context Menu*.

There are two types of tools, which can be invoked from the context menu:

Single click tools are invoked by clicking with the mouse button on the tool icon.

Mouse drag tools are started by pressing the mouse button on the tool icon. The
Context Menu will be hidden and the tool is active. Keeping the mouse
pressed, the tool is continuously applied until the mouse is released again.
At this point in time the *Context Menu* will be shown again.

Closing the *Context Menu* itself can be done by clicking on the *Close Context Menu*
button located in the center of the disk. Closing the *Context Menu* can be done with
the right and the left mouse button.

The *Context Menu* hosts multiple tabs with specific functionalities. They can be
selected using the buttons in the top-right corner (see Figure 3.11). The following
tabs are available:
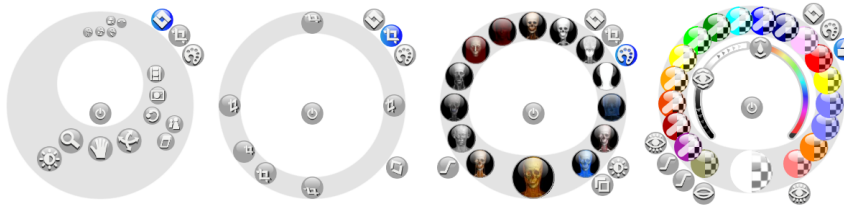
Figure 3.11: Context menu tabs: Main, Clipping, Presets, Masks

Main The *Main* tab provides access to the tools. This tab is always available. Rotating the tab will bring other tools into the focus.

Clipping The *Clipping* tab is only available for the volumetric display. Its purpose is to crop the displayed volume (see Section 5.23). This tab can not be rotated.

Presets The *Presets* tab provides access to window/level presets, if suitable presets for the displayed data exist. Different presets can be selected by rotating this tab (see Section 3.1.2).

Masks The *Masks* tab allows the user to modify the visual appearance of masks. It is only available if masks exist. Rotating the tab will make a certain mask active for modification. This tab provides the same functionality as the *Mask Editor* (see Section 3.1.4).

3.2.4 Displayed Units

Displayed values from user markups are assigned a unit depending on the pixel spacing information from the data.

pix is shown if no additional spacing information exists.

mm will be displayed if the image contains a proper pixel spacing information.

mm[USER] is appended if a manual image calibration was performed by the user (see Section 5.8).

mm[ERMF] is displayed if the image contains an *Estimated Radiation Magnification Factor*. The measured value is corrected by the value from the image information.

mm[<description>] is an indication that the pixel spacing information was calibrated during the acquisition process. The description (`<description>`) of the calibration method is provided by the modality.

mm[GEOMETRY] gives an indication that the pixel spacing information was computed from the geometric properties of the modality. This information is contained in the image data.

Figure 3.12: Monitor calibration

mm[FIDUCIAL] gives an indication that the pixel spacing information has been calibrated by a human operator or an image processing software. The accuracy of the calibration is best for areas close to the central beam and at the same depth as the fiducial marker.

mm[CALIBRATED] if non of the above values are present, however, an indication can be deduced from the image data that the pixel spacing has been calibrated, this information is presented.

3.3   Monitor Calibration

The purpose of *Monitor Calibration* is to specify the monitor geometry in the application in order to enable direct measurements on the monitor. This calibration **does not** calibrate the monitor's luminance.

The monitor calibration is triggered, if the zoom mode *True Size Zoom* (see Section 5.2) was selected on an uncalibrated monitor and the warning icon was clicked. Alternatively the monitor calibration can be invoked by opening the following URL in a browser from the client, which shall be calibrated:

`http://[HOST]:[PORT]/va/cmd/calibration.html?sId=CALIBRATE&ch=0`

The calibration page contains two lines (see Figure 3.12). To calibrate the monitor the vertical and horizontal line has to be measured. The measured length has to be entered into the fields in mm units. The calibration is completed by clicking the OK button.

4

RENDERING METHODS

The appropriate display of the data is automatically chosen by the *Hanging protocol* selection logic. If multiple rendering methods are available, different displays can be selected in the *Available Series* tab (Section 3.1.1).

4.1   Image Display

The Image Display is a conventional two dimensional depiction of the data. The simple image display is accessible via the hanging protocol *One up*, or it is part of specialized *Hanging protocols* for a specific set of acquisition images.

4.1.1   IMG – Image Rendering

The IMG Image Rendering results in a conventional display of radiological data as a stack of images. The images can be zoomed, panned, and - if the data are not RGB values - windowed as well (Figure 4.1). If spacing information exists in the data, a ruler (Section 3.2.2) indicating the real world size is displayed. The orientation indication (Section 3.2.2) depicts the orientation of the image with respect to the patient.

Image Localizer

Image Localizer depicts the spacial location of images with relation to each other. A common practice is to create a *Scout Image* during a Computed Tomography acquisition and store it together with the stack of sectional image data. The spatial location of the sectional image data is displayed in the *Scout Image* as localizer lines

Figure 4.1: IMG – Image Rendering

(see Figure 4.2). Thus a distinct and intuitive correspondence between the patient's body and the acquired images is possible.

For other modalities similar concepts exist. This especially holds for Magnetic Resonance Imaging, where the localizer lines are often referred to as *reference lines*. The application handles these cases the same way as long as the referencing images are in the same *Frame of Reference*. This means the reference images are spacial co-registered. In addition the referenced images must not be parallel to the *Scout Image*.

The spacial positions of the available images are sketched as faint lines, whereas the visible images are highlighted explicitly. When navigating through the stack of images the highlighted lines are updated automatically. It is also possible to use the localizer lines for navigation. When positioning the mouse cursor over the highlighted localizer lines, the mouse cursor changes it's shape to an arrows icon. Dragging the highlighted localizer lines will navigate through the image stack synchronously.

The application identifies images as being *Scout Images* and activates the localizer display by default. Alternatively the localizer display can be enabled using the *Localizer Visibility* tool (see Section 5.16).

4.1.2 WAV – Waveform Rendering

The WAV - Waveform Rendering displays a curve from ECG modality data (Figure 4.3). The time axis is depicted from left to right (horizontally), and the voltage axis from top to bottom (vertically). The major tick unit on the time axis is 0.2s and the major tick unit on the voltage axis is 500 µV. The spacing between major ticks is 5mm. For each tick a horizontal or vertical line is drawn, which results in a squared grid pattern.
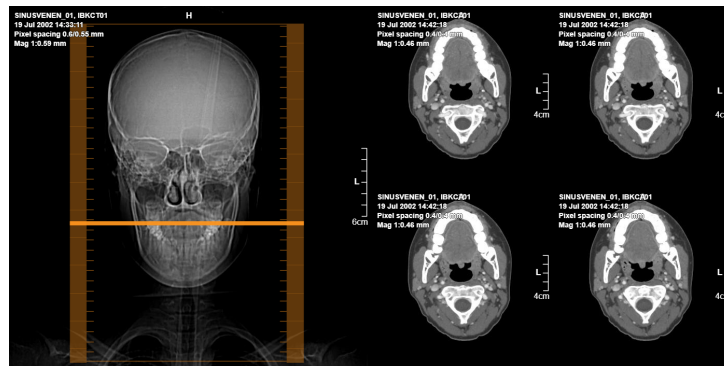
Figure 4.2: A scout image with image localizer lines

Figure 4.3: WAV – Waveform Rendering

4.1.3   DVH – Dose Volume Histogram Rendering

Dose Volume Histogram Rendering displays the statistical distribution of a radiation dose inside segmented objects (see Figure 4.4). The information on the segmented objects is taken from the referenced DICOM *RTStruct* object. For each segmented object, a separate curve in a distinct color is drawn. The color of the curve matches the color of the corresponding *RTStruct* object in the DVH rendering's legend. If no referenced *RTStruct* object is available all lines are drawn in white color and the segmented objects are named ''ROI n'' (n being the index of the object).

> **If the lines in the DVH Rendering are all white, no linkage to a valid *RTStruct* is available. No diagnostic reading must be performed in this case.**
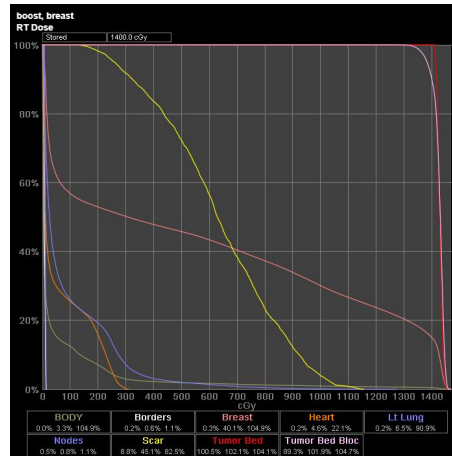
Figure 4.4: Dose Volume Histogram Rendering

The curve depicts the minimal radiation dose exposure in centigray (cGy) of an object's fraction in percent. The yellow curve of the object "Scar" in Figure 4.4 roughly passes through the point of 600cGy on the x-axis and 60% on the y-axis. This means that 60% of the object "Scar" is exposed to a radiation dose of 600cGy or more.

The percentage values in the object legend describe the *minimal, mean,* and *maximum* exposure with respect to the maximum prescription dose. The maximum prescription dose is displayed in the top left corner of the DVH display.

4.2   Sectional Display

The Sectional Display extracts two dimensional images from volumetric reconstructions. A volumetric reconstruction is based on the interpretation of a stack of images as one continuous spacial region. If the displayed data are suitable for volumetric reconstruction, then *Hanging protocols* for cross-sectional image reformatting will be available.

4.2.1   MPR – Multi Plane Reformation

Standard Multi Plane Reformation (MPR) provides two dimensional reformations of spatially connected orthogonal planes. The planes can be positioned interactively with arbitrary (oblique) orientations. By default the display layout comprises of three orthogonal MPR views and one additional display arranged in a 2x2 image setup (Figure 4.5). The fourth display is filled either with a volumetric display, or with a stack of images. The choice depends on the configuration of the system.
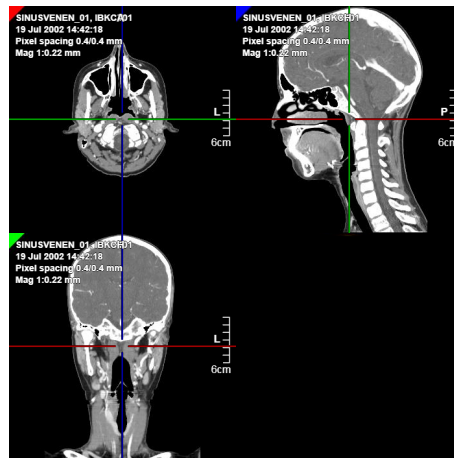
Figure 4.5: Multi Plane Reformation

MPR Lines

The standard MPR display contains three image annotations. In the top-left corner a triangle indicates the color of the displayed plane. Two orthogonal lines in different colors represent the remaining MPR displays. These lines are refered to as *MPR lines*. In total these annotations indicate the spatial relations of the MPR planes to each other (see Figure 4.5).

The *MPR lines* enable interaction with the MPR display. Positioning the mouse pointer at the center of the MPR lines, the cursor changes to indicate available movement interactions. Dragging the center of two crossing MPR lines will reposition interactively the intersection point (*MPR center*) and thus the resampling location of the two affected MPR planes.

In the close vicinity to the MPR center the mouse cursor changes to an arrow icon, if positioned over an MPR line. Dragging the MPR line from this position will reposition the selected MPR plane, keeping it orthogonal to the other two MPR planes.

Towards the end of the MPR line the mouse cursor changes to a hand symbol. Dragging the MPR line will rotate the selected MPR plane. The rotation is restricted to be in-plane with the plane, where the rotation is performed.

All types of interactions can be performed in arbitrary order, any number of times. The *Reset Display* tool (see Section 5.7) restores the initial state of the MPR.

Thick-slab MPR

A standard MPR display resamples a single plane from the volumetric data – in mathematical terms, an infinitely thin slab. The *Thick-slab MPR* resamples a plane with a certain thickness. This means adjacent data from the volume, i.e., orthogonal to the MPR plane, contribute to the final display. A common use case
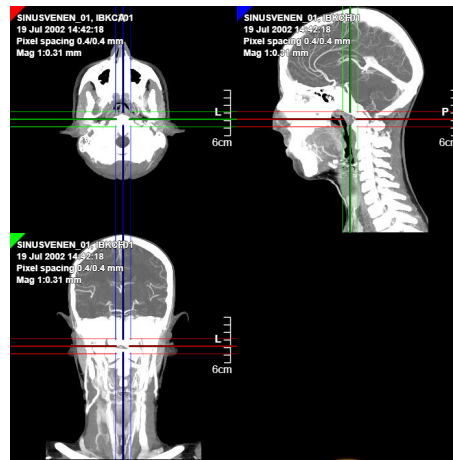
Figure 4.6: Thick-slab MPR

for this tool is lung nodule reading. After setting a certain thickness, the MPR plane is moved through the volume and the lung nodules become perceivable.

The *Thick-slab MPR* can be activated using the *Thick-slab MPR* tool (see Section 5.19). The interaction with the MPR display remains the same, as with the standard MPR display. In addition to the standard MPR the boundary lines of the resampled slab are indicated in the display (see Figure 4.6). The boundary lines are parallel to the MPR line and denote the limits of the resampled slab. The thickness of the resampled slab can be modified by dragging the boundary lines in the *Thick-slab MPR* display directly, or by using the tool *Thick-slab Settings* (see Section 5.20).

The resampling method may be selected via the *Thick-slab Settings* tool, or directly from the *Context Menu*'s main tab from the button group in the *Thick-slab Mode* (see Section 5.21). The resampling method defines how adjacent values orthogonal to the MPR plane are combined. Each pixel on the MPR plane is computed by processing the surrounding pixels along the plane's normal vector – the resampling ray. The available resampling methods are:

MIP refers to selecting the maximum intensity value along the resampling ray.

MinIP picks the minimum intensity value along the resampling ray.

Average sums up all values along the resampling ray and divides the sum by the number of sampling points.

4.3   Volumetric Display

The rendering methods described below are easily accessible via the Live Preset Gallery. The initially applied rendering method depends on the image type and is

Figure 4.7: Maximum Intensity Projection

configured in the *Hanging protocol*. All rendering methods are fully interactive and computed using the original resolution of the dataset.

4.3.1 MIP – Maximum Intensity Projection

The traditional Maximum Intensity Projection computes an X-ray type image by displaying the maximum greyscale intensity value for each image position. Each data value is mapped to its displayed luminance with a greyscale ramp from black to white. VISUAPPS MIP provides in addition a color mapping between intensity value and displayed color.

4.3.2 DVR – Direct Volume Rendering

Direct Volume Rendering displays the acquired volumetric data taking a color and translucency map into account (see Figure 4.8). Effectively this is the display method of choice when 3D renditions are referred to. Surface Shaded Display (SSD) is just a special case of DVR, which was introduced by software vendors to overcome the computational complexity of DVR. As DVR can handle all SSD functionality, no separate SSD rendering method is provided.

4.3.3 HVR – Hybrid Volume Rendering

Hybrid Volume Rendering is one of the outstanding functionalities of the VISUAPPS viewer (see Figure 4.9). HVR combines the advantages of MIP and DVR in a single picture. Adjusting a single parameter allows the user a smooth transitioning between MIP and DVR. Highlighting positions of implants and stents is quick and straightforward.
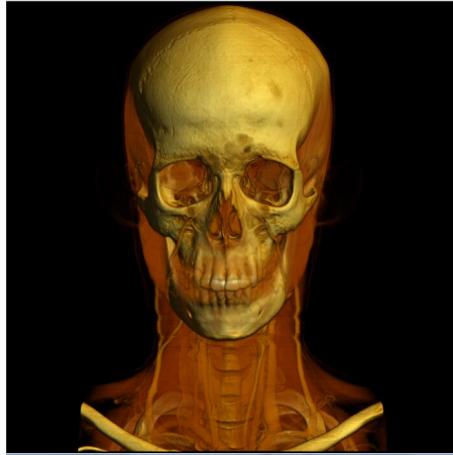
Figure 4.8: Direct Volume Rendering
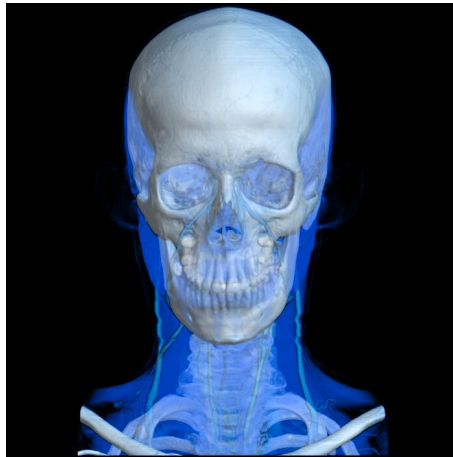
Figure 4.9: Hybrid Volume Rendering
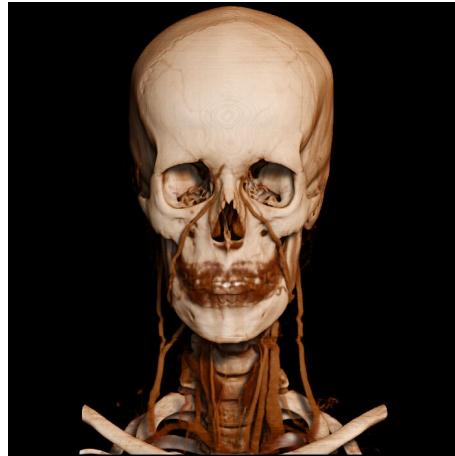
Figure 4.10: Shadow Volume Rendering

4.3.4   SVR – Shadow Volume Rendering

Shadow Volume Rendering (SVR) introduces shadows to 3D renditions. Shadows are an important visual cue of the human perception to interpret projective images (see Figure 4.10). Thus SVR enhances the volumetric perception and supports untrained users in understanding the images.

4.4   Fusion Display

Fusion Display renders co-registered data in a combined rendering. Term *reference data* denotes data that provide the spatial reference coordinates. Usually these are data that contain the anatomical structures (e.g., CT, MR). The *superimposed data* are data that are co-registered to the *reference data*. These data normally contain functional information (e.g., PET, NM).

Image fusion is available for image display, sectional display, and volumetric display (see Figure 4.11). The prerequisite for Fusion Display is either both data have the same *Frame of Reference*, or a DICOM *Spatial Registration SOP class*, referencing both data. The VISUAPPS Viewer does not provide functionality to automatically coregister data with different *Frames of Reference*.

If all requirements are met, *Fusion Display* hangings are available in the *Available Series Module*. To change the window/level of the *reference data* and *superimposed data* independently the tools *Window/Level Scope* (see Section 5.27) is available.
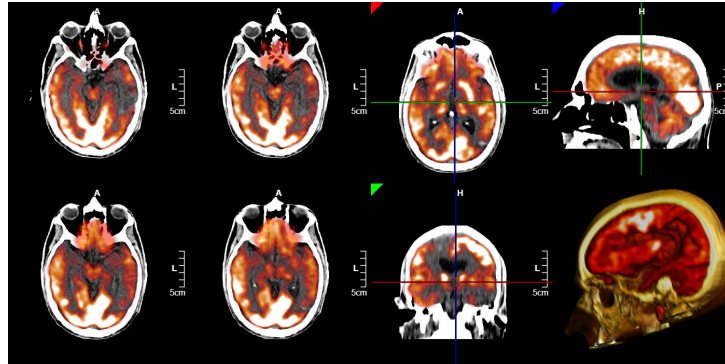
Figure 4.11: Image fusion for image, sectional, and volumetric display
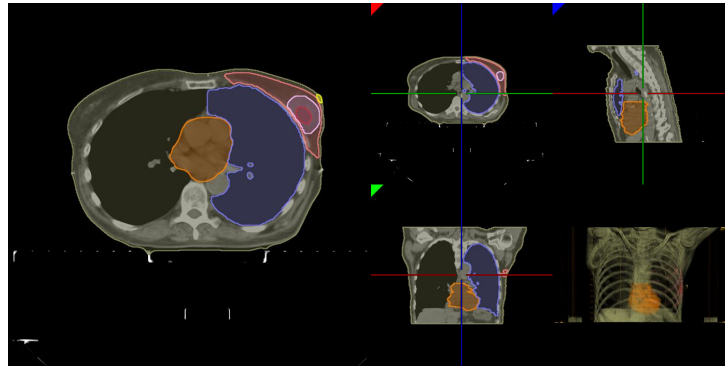
Figure 4.12: Rendering of masks in image, sectional, and volumetric display

4.5   Masks Display

Masks Display is possible if mask information is contained in the loaded imaging data (i.e., DICOM *RTStruct* data).   In this case specific *Hanging Protocols* are available in the *Available Series Module* that provide this display method. Displaying masks is supported by *Image Display*, *Sectional Display*, and *Volumetric Display*.

The visual appearance of the masks can be modified using the *Mask Editor* (see Section 3.1.4) or using the *Mask* tab in the *Context Menu* (see Section 3.2.3).

5

TOOLS

Tools are accessible via the context menu (Section 3.2.3). All tools are interactive and instantly provide user feedback. Only the currently applicable tools are presented in the user interface to reduce complexity. In this chapter the following terminology is used:

LMB denotes the left mouse button.

RMB denotes the right mouse button.

Drag is the process of holding a mouse button pressed while moving the mouse.

5.1 Image Zoom

Imaged Zoom enlarges or reduces the size of the displayed content. For an improved ergonomic experience the applied scaling factor is adjusted according to the zoom factor in use. The lower image scaling (downsizing) is limited to the minimum of either 50% of the display area, or 50% of the data size (2 data pixels per screen pixel). The maximum magnification is bound by a factor of 10 (1 data pixel per 10 screen pixels).

Image zoom

5.2 Zoom Mode

The VISUAPPS Viewer provides three different zoom modes: *Free*, *Pixel 1:1*, and *True Size*. Zoom modes restrict the possible image scaling to certain specific values. The zoom modes can be invoked via the context menu.

Zoom modes

Free Zoom   The scaling factor can be adjusted continuously from the minimum to the maximum possible scaling factor according to the user's discretion.

Pixel 1:1 Zoom   The scaling factor is bound to integer multiples from the original data to the displayed rendition. Switching to Pixel 1:1 Zoom mode initially maps one data pixel precisely to one display pixel. The scaling factor may be adjusted to alternative rational multiples using the *Image Zoom* tool: 0.5, 1, 2, 3, 6. In the special case of 1:1 no image resampling is performed, i.e., the image in its original resolution is displayed. The *Magnification* demographics (see Section 3.2.2) indicates this condition by appending the postfix "(Pixel 1:1)":

<div align="center">

`Mag 1:1.0 px (Pix 1:1)`

</div>

True Size Zoom   The scaling factor is adjusted so that the displayed data are presented in their true physical size. This zoom mode allows the user direct measurements of the displayed anatomy on the monitor. For this mode the monitor has to be calibrated, and the displayed data have to contain valid pixel spacing information.

> ⚠️ **When switching to *True Size Zoom* mode or performing geometric monitor calibration the user has to make sure that the browser zoom is set to 100%.**

If the monitor is not calibrated, the warning "Monitor not calibrated! (click to calibrate)" will be displayed. Clicking on the warning icon will invoke the monitor calibration functionality (see Section 3.3).

If the data does not contain correct spacing information, the warning "Data not calibrated! (click to calibrate)" will be displayed. It is possible to manually calibrate the image using the *Line Markup Tool* (see Section 5.8).

If active and all requirements for *True Size Zoom* mode are fulfilled the, *Magnification* demographics (see Section 3.2.2) indicates this condition by appending the postfix "(True Size)":

<div align="center">

`Mag 1:1.0 mm (True Size)`

</div>

5.3 Pan

The *Pan* tool positions the displayed content. Panning is aligned with the content which means that during panning the mouse's relative position to the content remains constant.

Image pan

The user can drag the image with the *LMB* to position it as desired.

5.4 Window Level

The *Window Level* tool modifies brightness and contrast of the imaging data. More precisely it defines a mapping from the original intensity range to the displayable luminance spectrum. Modifiying the window width results in images to appear sharper or smoother. The level parameter defines a certain intensity value. Adjusting the level, different structures may become visible (e.g. lung-tissue vs. liver-tissue). This tool is available for *volumetric display* and for two dimensional *image display*.

The interaction functionality of this tool is configurable. By default dragging horizontally (left to right) adjusts the level. Vertical adjustments change the window width. Dragging upwards sharpens the image, dragging downwards smoothens the image.

Window/Level

5.5 Window Invert

By default high intensity values (normally corresponding to areas with high attenuation) are mapped to brighter values. The *Window Invert* tool inverts the image intensity mapping. Dark areas are now displayed bright and vice verse. If the image contains information to identify the background, the background is not inverted. The information about the background area may either be represented as pixel padding values or as geometric shutter.

A single click with the *LMB* on the button toggles between inverted and regular mode.

Window Invert

5.6 Rotate Display (2D)

Image rotation can be performed using the *Rotate Display* tool. The rotation center is located in the display center. By using the *Pan* tool an arbitrary image rotation can be achieved.

While dragging the image rotates in a way that the mouse remains, as close as possible, over the same image location.

Rotate Display 2D

5.7 Reset Display

Reset Display brings back the original positioning and coloring of the display. The initial setup depends on the acquisition modality and is defined in the *Hanging protocol*.

The action is triggered by clicking on the button.

Reset

5.8 Line Markup

Using this tool, a length measurement markup can be applied. The markup displays the length of the drawn line. The unit depends on the image pixel spacing information (see Section 3.2.4).

Line Markup

```
???.?mm
```

The line markup drawing is started by a single click on the icon. For the markup drawing itself two interaction procedures are possible: *Drag-mode* and *2click-mode*. After the mouse was charged with the markup the procedure continues as follows:

Drag-mode The desired starting point of the line is placed on the image by pressing the *LMB*. While keeping the mouse button pressed, the mouse cursor is moved to the desired endpoint of the line. Then the mouse button is released.

2click-mode A single click at the desired starting point begins the line drawing procedure. A second click at the desired endpoint finishes the line markup drawing.

In both cases the markup line is continuously updated on the screen. The line can be modified by dragging the starting point and endpoint. Clicking in the central area of the line markup removes the markup again.

Image Pixel Calibration   If no pixel spacing information exists, a calibration tool is automatically added to the image on line markup creation. The intention is to align this calibration tool with a real-world calibration reference.

When calibrating the image's pixel spacing, the real-world calibration reference has to be placed parallel to the modality's projection plane. Furthermore it has to be in plane with the pathology of interest. If these requirements are not fulfilled, no calibration is possible. The calibration tool has to be removed manually in this case, as the pixel spacing is not valid.

The calibration tool can be placed by dragging its starting point and endpoint. Clicking on the reference unit the calibration tool's length scale toggles between 1cm and 10cm (default). Clicking on the calibration tool itself removes it and resets the measurement unit to *px*.

5.9 Circle Markup

Circle Markup

A circle markup can be applied. From the drawn circle various values are derived: radius *r*, circumference *U*, area *A*. The minimal, average, and maximum value of the pixels inside the circle are depicted in the schema *min < avg < max* plus the intensity unit. In addition the standard deviation σ and the discretization error is presented (see below). The geometric units depend on the image pixel spacing information (see Section 3.2.4).

```
r= ???.?mm
U= ???.?mm
A= ???.?mm
???.? < ???.? < ???.? HU
σ= ???.?  err= ??.?%
```

Placing the circle markup works analogously to the placement of the line markup tool (see Section 5.8). The first position defines the center of the circle. The second position specifies a point on the circle. Repositioning of the circle can be done by manipulating the center point. The radius of the circle is adjustable by modifying the circle in the right-top and right-bottom quadrant. Clicking in the circle's left-top or left-bottom quadrant removes the markup.

Discretization Error   The standard deviation σ is computed from the pixels inside the circular area defined by the circle markup. The difference between the analytically exact area and the area covered by the actual pixels is referred to as *discretization error*. This error is computed according to the following formula, where $x_{mm}y_{mm}$ depicts the area one pixel covers in real world:

$$err = \frac{\left| r^2\pi - \sum_{pixels} x_{mm}y_{mm} \right|}{r^2\pi} 100 \tag{5.1}$$

If the error exceeds 25% the output of the standard deviation σ is suppressed.

5.10   Angle Markup

An angular annotation can be applied using the *Angle Markup* tool. The angle between two lines is displayed in degree between 0° and 180°.

```
??.?°
```


Angle Markup

The placement of the two lines can be done either in *Drag-mode* or *2click-mode* (see Section 5.8). Each line can be modified by repositioning the endpoints. A single click on either angle line removes the angle markup from the display.

5.11   Navigation Tool

The navigation tool is available if multiple renditions are contained in a display. For a *2D image display* this is the case if multiple images are available. If the images differ in their spacial location, the navigation tool browses through the image stack. If the images represent a time series, the navigation tool browses through the time steps. For a *3D display* the navigation tool always iterates over time series.

Navigation Tool

A single click on the *Navigate forward* button iterates to the next image, the *Navigate back* button iterates one image backwards. With the button *Navigate through images* (in the middle) jumping directly to certain image positions is
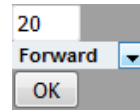
Figure 5.1: Cine Settings

possible. Dragging the mouse to the right moves forward, dragging the mouse to the left moves backwards through the image stack. *Navigate through images* drops images during interaction in case the update frequency can not be sustained.

5.12   Cine Tool

Cine Tool

The cine tool enables and disables to automatically loop through the images. The tool is available if navigation is possible (see Section 5.11). The playback parameters can be modified through the *Cine Settings* button (rightmost icon). With the parameters dialog (see Figure 5.1) the *frame rate* and the *playback mode* can be specified. The playback mode provides three options:

Forward The animation is played from the first frame to the last frame. When the last frame is reached, the animation restarts with the first frame.

Backwards The animation is played from the last frame to the first frame. When the first frame is reached, the animation restarts with the last frame.

Toggle The animation starts from the first frame and continues to the last frame. If the last frame is reached the animation is reversed and continues from the last frame to the first frame.

Clicking on the *Play* button (leftmost icon) starts the cine mode. The button changes its icon to a *Stop* icon, as an indication that the cine mode is active. When pressing the *Stop* button, the cine mode stops and the button is changed back to its initial state. Clicking on the *Cine Settings* button opens the parameters dialog. The frame rate can be adjusted by typing a new value. The playback mode can be changed via a drop down list. Clicking on the *OK* button sets the new values. Clicking anywhere else in the display closes the parameters dialog without modifications.

5.13   Presentation Group Switch

Presentation Group Switch

If multiple *Presentation Groups* are available in the current *Hanging protocol* the *Presentation Group Switch* tool is available. The current content of the *Image Area* will be entirely replaced with the content of the next/previous *Presentation Group*.

A single click on one of the two tool buttons will trigger the *Presentation Group* switch.

5.14 Demographics Visibility

The visibility of patient demographics and technical information can be changed using the *Demographics Visibility* tool. By default only the patient demographics are visible. The patient demographics information is located in the upper corners of the image. The technical information is displayed in the lower corners. The technical information is not always present, depending on the data type (see Section 3.2.2).

Demographics Visibility

The *Demographics Visibility* tool is a tri-state button. The first click adds the technical information, and the button's icon changes. A second click on the button disables all image information annotations, again the button's face changes. The third click brings back the default configuration, which is patient demographics only display.
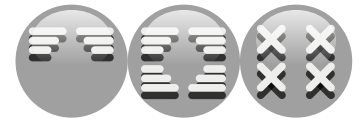
5.15 Annotation Visibility

The visibility of markups (e.g., lines, circles, etc..) can be turned on and off with the *Annotation Visibility* tool. If not defined differently by the *Hanging protocol*, markups are initially visible by default. A special case are *Mammography CAD Marker*, as initially there is only an indication about their presence visible.

Annotation Visibility

Clicking on the button *Annotation Visibility* toggles the visibility of markups.
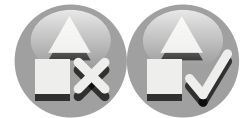
5.16 Localizer Visibility

An image may be set to act as *Scout Image* by the *Localizer Visibility* tool. The visibility of localizer lines is activated or deactivated. If images in the same *Frame of Reference* exist, the sectional lines will be displayed as localizer lines. Further interaction may be performed, as described in Section 4.1.1.

Localizer Visibility

Clicking on the button *Localizer Visibility* toggles the visibility of localizer lines.
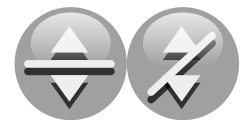
5.17 Capture Jpeg

Capture Jpeg creates screen captures of the currently displayed content. The capture is presented in a new tab. The image can be stored locally using the browser's "save" functionality. Along with the regular jpeg image file, additional information is stored to recreate the captured state of display. To apply the stored state, the jpeg file has to be dropped onto the display.

Capture Jpeg

A single click on the button opens a new tab in the browser with the captured content. Using the "save as" functionality of the browser, or using the drag and drop functionality allows the storage of the image directly to the desired file-system location. To restore a previous display, the previously stored jpeg file can be dragged and dropped onto the display.

5.18   Add Key Frame

Add Key Frame

The *Add Key Frame* tool creates a new *key frame* from the currently displayed content. The *key frame* is added to the current video. The video creation itself is done with the *Video Creation Module* (see Section 3.1.3).

A single click on the button creates a *key frame*. The key frame will appear in the *Video Creation Module*'s key frame tray.

5.19   Thick-slab MPR

Thick-slab MPR

The *Thick-slab MPR* tool enables the MPR display to resample a slab of a certain thickness in the volume instead of a simple plane (see Section 4.2.1). Using this mode, objects in the close vicinity to the MPR plane become visible in the display.

A single click on the button toggles the *Thick-slab MPR* mode on and off.

5.20   Thick-slab Settings

Thick-slab Settings

The *Thick-slab Settings* tool opens a dialog, which allows the user to change *(a)* the thickness of the slab and *(b)* the resampling mode. The thickness of the slab is given in mm. The thickness is the distance from border to border with the MPR plane in the center. The resampling method can be chosen from one of the three options available: *MIP*, *MinIP*, and *Average* (see Section 4.2.1).

A single click on the button opens a dialog window. Inside the dialog window a text field allows the user the specification of the slab thickness in mm, and a drop-down chooser to pick a resampling method.

5.21   Thick-slab Mode

Thick-slab Modes

The *Thick-slab Mode* buttons provide a shortcut to quickly switch to a specific resampling method. The available methods are *MIP* – Maximum Intensity Projection, *MinIP* – Minimum Intensity Projection, and *Average* – Average intensity computation (see Section 4.2.1).

Clicking on the respective button changes the resampling method to the selected one. The MPR display is immediately updated, if *Thick-slap MPR* display is active.

5.22   Rotate Display (3D)

Rotate 3D

Rotation is performed by simulating a virtual trackball. This is a common solution for controlling 3D interactions with a 2D input device. The interaction is intuitively similar to rotating a tennis ball with the forefinger. This tool is only available for the volumetric display.

When applying the tool the volumetric display's orientation follows the mouse position.
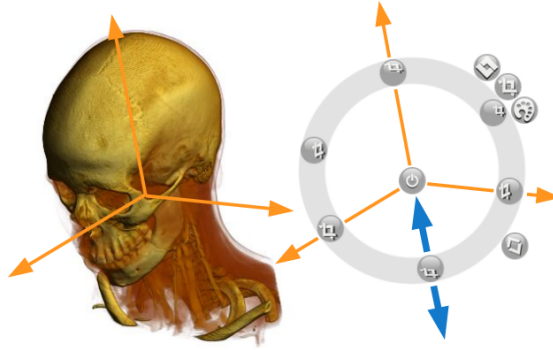
Figure 5.2: Volume Cropping tool

5.23 Volume Cropping

Volume cropping is only applicable in the volumetric display (see Section 4.3). Each button represents an axis aligned plane with can be used to restrict the displayed volume. In total six cropping planes exist: *left* and *right*, *head* and *feet*, *anterior* (front) and *posterior* (back). The main purpose of this tool is to quickly remove obstacles that obscure the pathology of interest. The button's location depends on the orientation of the displayed volume (see Figure 5.2).

To apply a cropping plane, the corresponding button has to be dragged towards the center (see Figure 5.2 - blue arrows). To make the cropped parts visible again, the button has to be dragged away from the center. If two cropping buttons are too close to each other to distinguish the corresponding planes, it is helpful to rotate the volume a little bit.
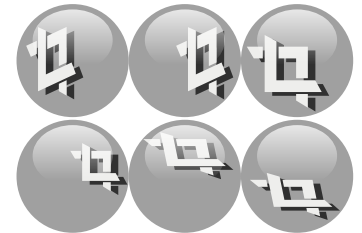
Cropping Planes

5.24 Slab Slicing

Slab slicing is only available for volumetric display (see Section 4.3). If the volume is cropped to a to thick slab (see Figure 5.3), this slab can be moved through the dataset using this tool. The slab provides a crossectional insight to the dataset with adjacent information context depicted. A possible application is the search for nodules in the lung. The cropping planes cut away the distracting parts of the tracheal system and the pulmonary arteries for a specific focal plane. While slicing the slab through the dataset the lung nodules become visible one after another.

As a prerequisite for this tool the volume has to be cropped using the *Volume Cropping* tool (see Section 5.23) beforehand. Dragging the *Slab Slicing* tool button up or down displaces both cropping planes most perpendicular to the viewing direction. If the slab slicing should be applied to a different pair of cropping planes,

Slab Slicing
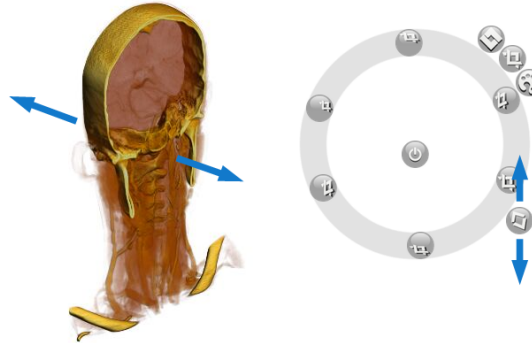
Figure 5.3: Slab Slicing tool

the volumetric display has to be rotated accordingly. This can be done using the *rotate display (3D)* tool (see Section 5.22).

5.25   Hybrid Display Weighting

Hybrid Display Weighting

The *Hybrid Display Weighting* tool is only available for *hybrid volume rendering* (see Section 4.3.3). This tool modifies the combination weighting of the *DVR* and *MIP* contributions to the final display. By default *DVR* and *MIP* contribute equally.

   To increase the contribution of *MIP* rendering to the final display, the button has to be dragged upwards. When dragging the button downwards the *DVR* contribution is increased.

5.26   Background Color

Background Color

The *Background Color* tool is only available for *volume display* (see Section 4.3). This tool toggles the background color between white and black.

   A single click with the *LMB* on the button toggles the background color.

5.27   Window/Level Scope

Window/Level Scope

The *Window/Level Scope* specifies which data will be modified when applying the *Window/Level* tool. This tool is only available for *Fusion Display* (see Section 4.4).

   The buttons for specifying the window/level target are located in the *Presets* tab of the *Context Menu* (see Section 3.2.3). A single click on one of the buttons makes this target active until a different member of the button group is pressed.

APPLICATIONS

The following chapter describes special use cases, which the application supports.

6.1   Mammography

The application fulfills the IHE Mammography specification[1]. See Figure 6.1.
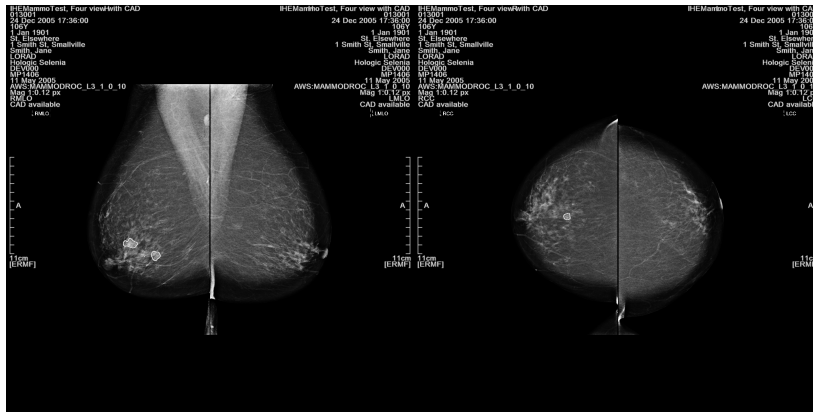
Figure 6.1: Mammography application

---

[1]http://www.ihe.net/Resources/upload/IHE_Mammo_Handbook_rev1.pdf
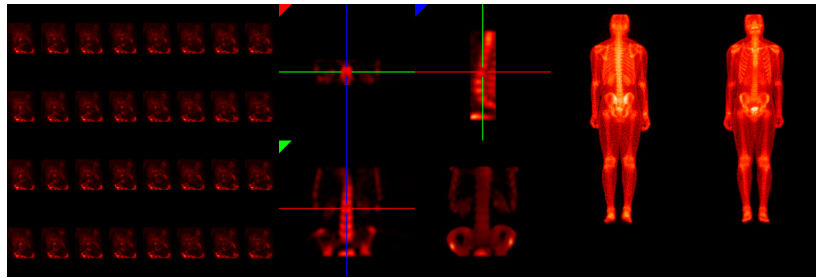
Figure 6.2: Nuclear Medicine application

6.2 Nuclear Medicine

The application fulfills the IHE Nuclear Medicine specification[2]. See Figure 6.2.

---

[2]http://www.ihe.net/Technical_Framework/upload/IHE_RAD_TF_Supp_NMI_TI_2007_05_17.pdf

TROUBLE SHOOTING

7.1 A tool is not available

If a certain tool does not appear in the context menu, the cause may be either that *(a)* the tool is not applicable in the current display, or *(b)* the tool was disabled by configuration.

7.2 The application responses slowly

If the application starts lagging, behaves very slowly and hardly responses, the cause might be that too many browser tabs or windows are opened. The application starts lagging if more than 2-4 (depending on the browser in use) tabs or windows are opened concurrently. To solve the issue close the previously opened tabs or windows.

CHAPTER **8**

REFERENCE

8.1 List of Tools

Icon	Tool	Interaction	Effect
	Close CM	Single click.	Close the context menu
	Image zoom	Drag up/down	Re-size the image relative to the initial mouse position
	Zoom modes	Single click	Sets a specific zoom mode
	Image pan	Drag in any direction	Re-positions the image relative to the mouse position
	Window/Level	Drag up/down for window Drag left/right for level	Changes contrast and brightness.
	Window Invert	Single click	Inverts the image luminance values
	Rotate Display 2D	Drag in any direction	Rotates the image
	Reset	Single click	Resets the display to the initial state
	Line Markup	Single click	Drawing of a measurement line
	Circle Markup	Single click	Drawing of a circle

	Angle Markup	Single click	Drawing of an angle
	Navigate back	Single click	Navigates to the previous image
	Free navigation	Drag left/right	Free positioning within stack of images
	Navigate forward	Single click	Navigates to the next image
	Start cine	Single click	Starts automated image looping
	Stop cine	Single click	Stops automated image looping
	Cine settings	Single click	Changing frames per second and looping settings.
	Previous PG	Single click	Switch to previous Presentation Group
	Next PG	Single click	Switch to next Presentation Group
	Demographics Visibility	Single click	Rotates through the visibility settings
	Annotation Visibility	Single click	Show/hides annotations
	Localizer Visibility	Single click	Show/hides localizer
	Capture Jpeg	Single click	Opens a jpeg capture in a new tab
	Add Key Frame	Single click	Creates a new key frame
	Thick-slab MPR	Single click	Enables/disables the Thick-slab MPR mode
	Thick-slab Settings	Single click	Settings for Thick-slab MPR mode
	Thick-slab MIP	Single click	Thick-slab MPR MIP resampling
	Thick-slab MinIP	Single click	Thick-slab MPR MinIP resampling
	Thick-slab AVG	Single click	Thick-slab MPR average resampling
	Rotate 3D	Drag in any direction	Rotates the volume
	Crop Left	Drag towards center of the context menu.	Applies left cropping plane

	Crop Right	Drag towards center of the context menu.	Applies right cropping plane
	Crop Head	Drag towards center of the context menu.	Applies head cropping plane
	Crop Feet	Drag towards center of the context menu.	Applies feet cropping plane
	Crop Anterior	Drag towards center of the context menu.	Applies front cropping plane
	Crop Posterior	Drag towards center of the context menu.	Applies back cropping plane
	Slab Slicing	Drag up/down	Move the slab through the data
	Hybrid Display Weighting	Drag up for MIP and down for DVR.	Seamless transition between MIP and DVR
	Background Color	Single click	Toggle the background color black/white
	Window/Level superimposed	Single click	Set window/level scope to superimposed dataset
	Window/Level both	Single click	Set window/level scope to both datasets
	Window/Level reference	Single click	Set window/level scope to reference dataset

INDEX