CTvox Quick Start Guide

For Software Version 1.0

In order to work comfortably with CTvox, we suggest the following configuration:

- a display adapter with 1 GB or more of *on-board* memory
  (installation of the latest drivers is recommended, for maximum OpenGL compliance)
- 1024 x 800 screen resolution or better

Online help is available through **Help Topics** in the **Help** menu.
Context-sensitive help is also provided:

- Clicking  on the toolbar first and then any element of the user interface will pop up help on that specific item.
- Pressing **F1** in any dialog will pop up help specific to that dialog.

To load the dataset of your choice, select **Load Volume Data...** in the **Actions** menu, or click the  button on the toolbar. Alternatively, you can drop a file from the dataset on (a shortcut to) CTvox: the application will start and load the dataset without further prompting.

On the right, the *Load Dataset* dialog offers several options to resize your dataset:

- reducing the dataset in all directions
- loading only a central portion of the dataset, either in-slice (XY) or in the z-direction

Once a dataset has been loaded, the transfer function plot area (in the *Transfer Function* pane, in the upper left corner) will also display a histogram for the data. Since low intensities dominate the distribution (due to noise and other contributions), the histogram is best displayed using a log scale: right click in the transfer function window and select **Log-scaled histogram**.
Likewise, these low-intensity contributions (which dominate the periphery of reconstructed images) tend to obscure the core of the volume data. Therefore, it is usually useful to suppress these low intensities in the rendering to some extent, by making them highly transparent:

- Make the *alpha channel* active in the transfer function editor, by selecting *A* from the *Channel* drop down box. The markers for the purple alpha curve (initially a straight line, with just two markers, one at each end) appear, to indicate it is active.
  Alpha stands for opacity (the opposite of transparency). Hence, we will need to map low intensities to low opacities (= high transparency).

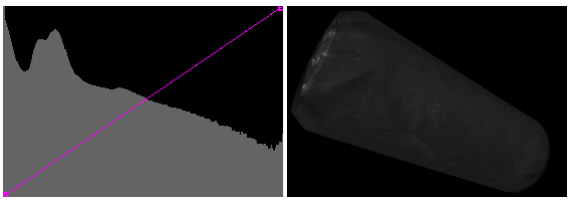

- Add a knot to the alpha curve by clicking in the lower left corner of the transfer function window: the curve changes to smoothly pass through the new and existing markers.

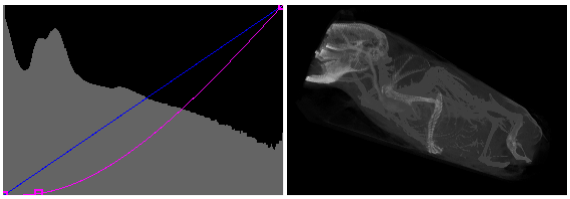

- Try moving the marker to smoothly change the transfer function and obtain the desired effect: left click the marker and drag the mouse with the left mouse button held down.
  One can switch between a *spline* (smooth) curve and a *polygonal line* by toggling **Use splines** in the context menu (invoked by right clicking in the transfer function plot area).
  The context menu also includes entries to reset the currently active or all curves and to delete a marker (when invoked with the mouse pointer over a marker).
  Markers can also be deleted by double clicking them or dragging them outside the plot area.

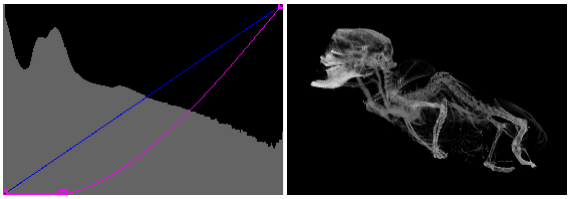

Often, the bitmaps making up the dataset will contain overlayed scales or textual information, which also hide the actual object from view. To remove these, proceed as follows:

- The wireframe box surrounding the object acts as a *clipping box*: its faces can be moved independently, to selectively clip away portions of the object.

  (The visibility of the clipping box can be toggled using the  button on the toolbar.)
- If necessary, rotate the object by *dragging the mouse* with the *right button* held down (see further below) until the mouse pointer can be brought *directly* over the cube face that needs to be moved inward.
- With the mouse pointer directly over the cube face, click and hold down the *left mouse button* while holding the **Shift** key pressed. This will *'pick'* the clipping plane: to reflect this, the face's edge turns from purple to yellow.
- By dragging the mouse, the cube face can be moved along its perpendicular, clipping away portions of the object that fall outside the clipping box.
- Release the mouse button when the desired position is reached: the face's edge shifts back to purple, to indicate it is no longer picked.

Two modes exist for navigating the scene:

-  *Object movement*: the orientation (rotation) and position (translation) of the object can be adjusted using the right and left mouse button respectively; the camera can be moved closer to or farther from the object using the mouse wheel.
-  *Camera movement*: the object is held in place; the camera can be rotated around the object (right mouse button) and the camera viewing angle can be adjusted (left mouse button); again, the mouse wheel moves the camera forward or backward.

The following table summarizes the navigation controls:

Mouse interaction	Object movement	Camera movement
Left button down	Object translation	Camera viewing angle
Left button + Ctrl	Clipping box only translation	-
Left button + Shift	Pick clipping plane	Pick clipping plane
Right button down	Object rotation	Camera rotation
Right button + Ctrl	Clipping box only rotation	-
Mouse wheel	Camera forward/backward	Camera forward/backward
Double click	Toggle interaction mode	Toggle interaction mode

The right mouse button controls rotation in both modes.
In object movement mode, the left button controls the *camera viewing angle* (acting as a zoom, but introducing perspective distortion when the camera is close to the object; the viewing angle can also be set in the *View Options* pane); in camera movement mode, the left button controls the translation.
By holding the **Shift** key down while left clicking, the face of the clipping box under the mouse pointer can be 'grabbed' and subsequently dragged to clip a portion of the object.
Holding the **Ctrl** key during rotation or translation in object movement mode will keep the object in place and move only the clipping box, allowing to position a clipping plane appropriately.
The mouse wheel controls the distance between camera and object in both modes.
Double clicking either mouse button in the imaging window toggles the movement mode.

Precise motion control is available through the **Movement > Numeric...** command in the **Actions** menu.

The orientation and position of the camera and object can be reset to their starting state using **Actions > Movement > Reset Camera** () and **Actions > Movement > Reset Object** () respectively.

In *volume rendering*, every voxel in the volume data is assigned an emission color (the color the voxel emits, determined by its red (R), green (G) and blue (B) components) as well as an opacity (represented by A, for *alpha*; opacity is the opposite of transparency). Both of course depend on the original volume data and the mapping is governed by the so-called *transfer function*, displayed on the left side of the

CTvox window. The horizontal axis in the plot represents the original scalar data (the x-ray attenuation), the vertical axis represents one of the transfer function components (R, G, B, RGB, A).
Modifying the opacity controls the visibility of the corresponding voxels and how much they obscure more distant voxels. By setting the opacity for a given intensity range to zero (full transparency), the corresponding voxels are effectively made invisible.
The color channels can be linked (yielding the RGB channel), in which case only gray values will appear in the image, or controlled independently (R, G and B channels), allowing to introduce color in the image .

Channels	Transfer function plot	Resulting image
RGB, A	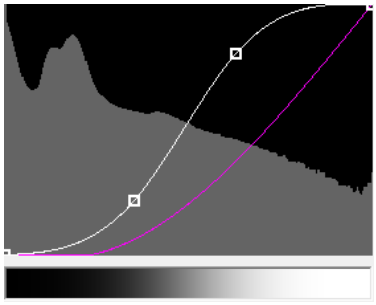	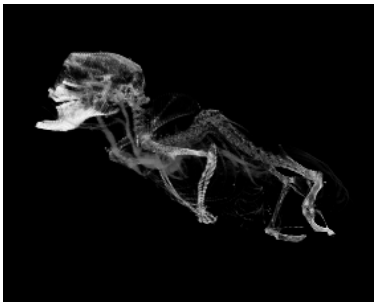
R, G, B, A	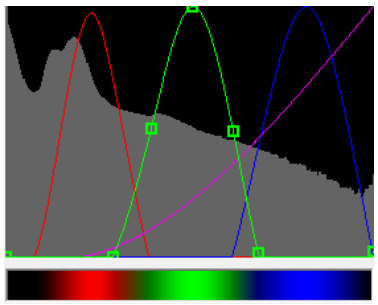	

After startup, the *blending mode* is set to *Volume* blending (this blending type can be selected at any time by selecting **Blending > Volume** in the **Actions** menu or pressing the  button on the toolbar), meaning the emission-absorption model as described above is used.
Other blending modes available are:

-  *Attenutation*
  This mode corresponds to digitally reconstructing a radiograph: a radiograph along the current viewing direction is generated.

-  *Maximum Intensity Projection*
  For every ray hitting the screen, only the most intense voxel along the ray is retained. Hence, this mode automatically highlights the structures of maximum intensity.

The *View Options* pane, on the left, groups settings that determine the general appearance of the scene.
The *camera viewing angle*, described above, can be set using the corresponding slider.
The *background* against which the scene is displayed can be specified:

- *Color fill*
  An homogeneous background of the user-specified color. In this mode, the camera position is unbounded.
- *Color cube*
  A cube with faces of slightly varying color, based on a user-specified color. The *cube size* can be specified. In this mode, camera movement is limited, such that the camera remains within the cube.
- *Image cube*
  A cube whose faces are painted with user-specified bitmaps. A default set of bitmaps is provided to generate the coastline scenery depicted below. Here, too, the camera position is bounded by the cube and the cube size adjustable.

The three types of background are illustrated below:

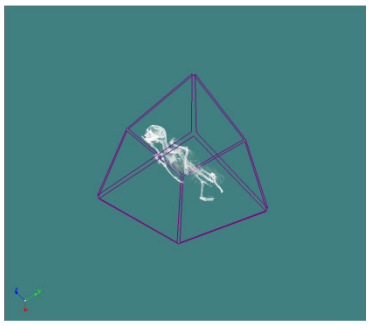
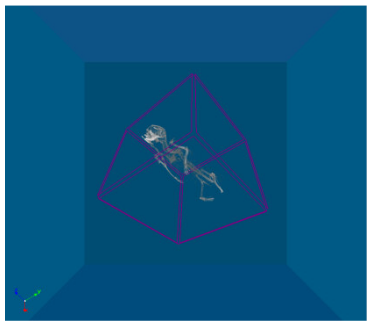
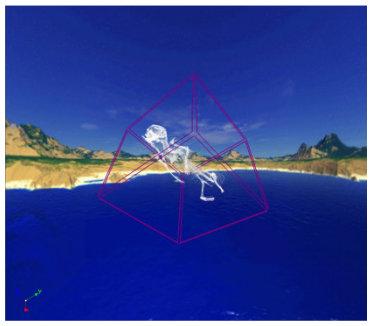

The Flight Recorder functionality can be used to produce fly-through animations for your 3D data. To invoke/hide the *Flight Recorder* dialog, select the **Actions > Flight Recorder** command or click  on the toolbar.

The idea is to set up the scene and to tag it as a *key frame* by clicking the *Add* button. When eventually creating the movie, the program will generate the specified number of intermediate frames by appropriately interpolating these key frames.

Use the following controls to navigate the sequence of key frames:

|<	show the previous key frame
<<	cycle backward trough the key frames
|>	Start preview
||	pause preview/cycling
>>	cycle forward trough the key frames
>|	show the next key frame

The *Hide clipping box* checkbox allows to suppress the clipping box in the generated movie, even though visible in the rendering.
When ready, the movie can be generated by clicking *Save movie*. The output is either an AVI file or a series of bitmaps (BMP or JPG).

At any moment, a screen capture of the rendering can be saved (as BMP or JPG), using the **Save Image...** command in the **Actions** menu or the  button on the toolbar.

CTrender supports *stereo viewing*, to further enhance the 3D impression. To switch on/off stereo viewing, use the **Stereo Mode** command in the **Options** menu or click the  button on the toolbar.
The *stereo angle* required for an ideal stereo image depends on the specific setup (i.e. the viewer's interocular distance and the distance from eyes to screen) and can be adjusted in the *Preferences* dialog, accessible through **Options > Preferences...**