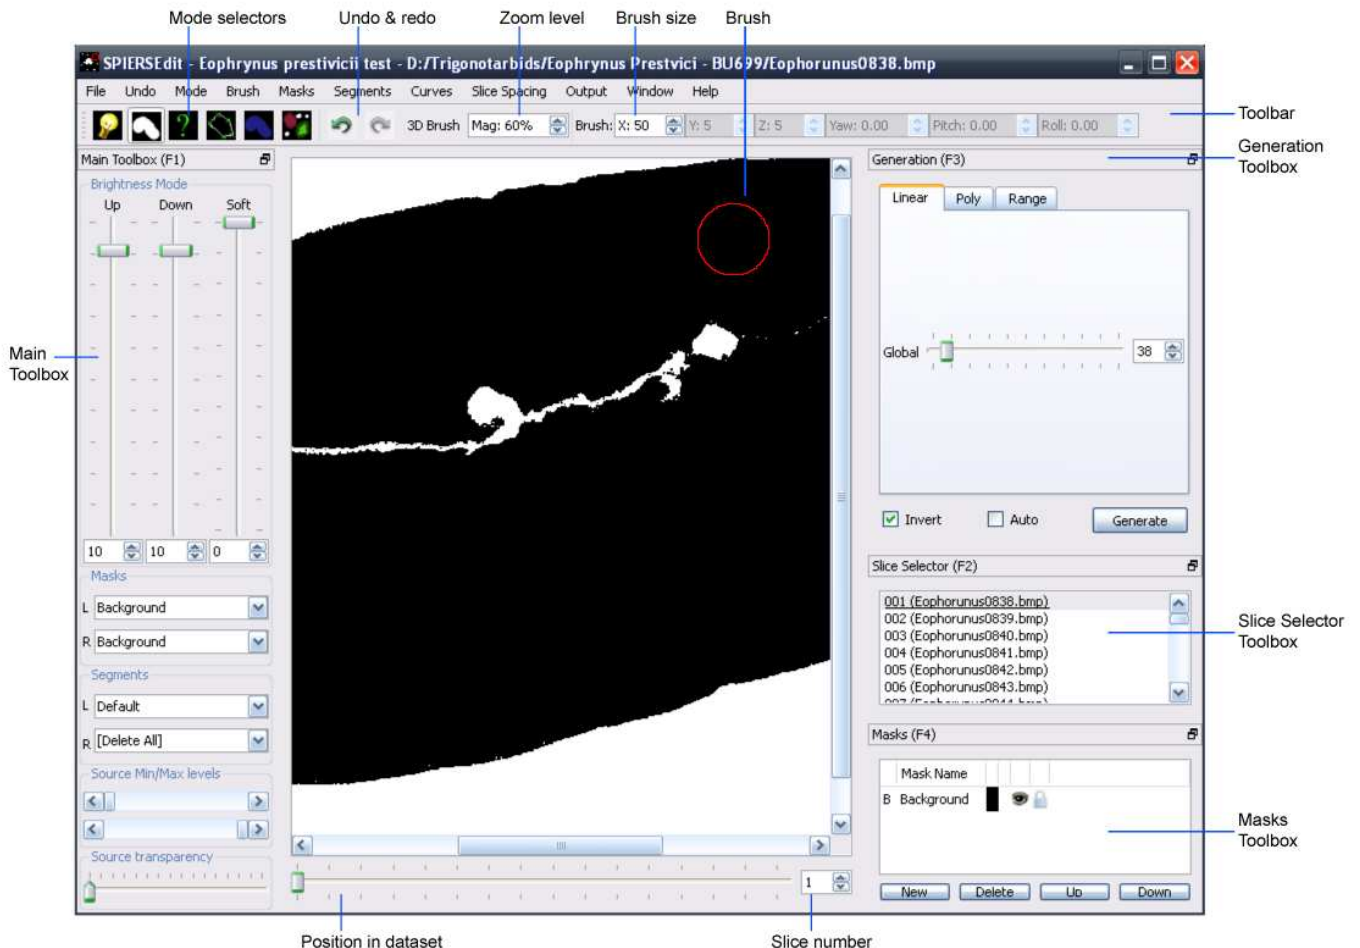


SPIERS Software Suite

SPIERSedit 2.0 - The Basics

Interface: The default layout of SPIERSedit is shown below:



Toolbox Shortcuts: F1 - Main Toolbox; F2 - Slice Selector; F3 - Generation; F4 - Masks; F5 - Segments; F6 - Curves; F7 - Output; F8 - Histogram; F9 - Info.

Mode Shortcuts: Ctrl + B - Brightness; Ctrl + S - Segment; Ctrl + R - Recalc; Ctrl + C - Curve; Ctrl + L - Lock; Ctrl + M - Mask.

Other shortcuts: These can be found next to the associated menu commands.

SPIERS is custom built software for producing 3-D models from tomographic datasets. These instructions are aimed at CT users, but the software can also reconstruct serial grinding and acetate peel datasets.

Datasets should be image (bmp/png/jpg) sequences. SPIERS data is stored in a .spe file and an associated subfolder with the same filename. The dataset will only work when the .spe is in the source images folder with the similarly named subfolder.

Creating a new dataset: File → New. Navigate to folder. Select images → Open. Give file name and a description. *Downsamples:* These allow downsampling of the data. A high number will allow the computer to run faster, and each dataset will require less work, however the resolution will be reduced

by the downsample factor. Downsampling can be changed at a later date. All images need to be the same dimensions.

Following this a progress bar a black rectangle should appear. This is the 'threshold' image - every white pixel represents fossil, and every black one matrix. At the moment all are black. Pressing space toggles between the threshold and the original. The appearance of the original can be altered using the 'Source Min/Max levels' sliders in the Main Toolbox - these adjust the brightness and contrast. The 'Source Transparency' below these sliders allows the threshold to be superimposed on the CT image. Zoom can be changed using the "Mag" option in the toolbar, or by shortcuts 'q' and 'a'.

To create a model, a threshold needs to be created. First, however, there are numerous options in SPIERS to customise the program for different PCs and different datasets.

Settings: File → Settings give access to SPIERS' settings. The cache memory can be altered here to improve performance (larger=better performance, but greater use of system resources). The source files and working files can be compressed and uncompressed. Uncompressed (default) will take up more disk space, but run faster. Compressing makes the datasets smaller, but also slower, especially when rendering. Setting low compression saves the files using a lossless format, while high saves them with a lossy algorithm (but can compress a dataset by a factor of ~5). Most systems will run best with both set to off (the default), but if disk space is limited compression is an option. When the compression is set any changes will be resaved with the new compression settings. Other files are not altered. The 'Re(compress) source/working files' buttons apply the compression to the dataset in its entirety.

Toolboxes: All toolboxes (shortcuts above) are dockable windows. These can be dragged and placed anywhere on the screen, resized, docked in the user interface on either side (in any desired arrangement), and docked on top of each other to create tabbed side panels.

Generation: F3 opens the Generation Toolbox which allows the generation of a threshold image. For CT data the linear tab of this toolbox is usually sufficient. Choose this tab and click generate. This creates a rough threshold which needs to be adjusted to match the fossil. If the fossil is darker than the matrix the invert option should be checked. The global slider adjusts the grey level cut-off between 'fossil' and 'matrix'. Clicking generate after each change or using the 'Auto' option updates the threshold image. The global slider should be varied until floating pixels are minimised, yet as much fossil as possible is correctly assigned. This generation rule can then be applied to the dataset - the slices to which the generation should be applied can be selected in the Slice Selection Toolbox (F2). Clicking 'Generate' creates a threshold image for each selected slice. *Polynomial: This is only necessary for colour images. Range: This allows the user to define a range of grey levels rather than a single cut-off. See 'segments'.*

Cleaning: The threshold is used to create the model. Any 'on' pixel - fossil - creates a voxel (3-D pixel) for that slice. When the 'on' pixels in each slice are rendered next to each other as voxels a three dimensional volume - the model - is created. Some floating pixels and other noise (caused by cracks, ring artefacts, or other imperfections) are likely to remain. Without cleaning these will render as floating points, or unnecessary surfaces. Also, fine structures can be omitted by the automatic generation, and edges can need manual cleaning. All such imperfections can be fixed in several ways using the different modes (tools).

Modes: All modes can be accessed using the shortcuts above, on the toolbar, or in the Mode menu. All are applied with a brush. Brush options are found in the brush menu and on the toolbar (note: 3-D brush is currently experimental). The brush is arbitrarily resizeable, and has ten preset sizes (shortcut ctrl + 1 → ctrl + 0) *Floating voxels: While these can be removed manually on each slice, floating pixels (which cause floating voxels) are very time consuming to remove by hand on every slice. There is a filter in SPIERSview to automate this for small floating voxels, so only large ones need be removed manually.*

Brightness mode: This allows manual cleaning of data by adjusting brightness level from the original image (this is applied to a second copy of the image ('working image') stored in the .spe's subfolder, so will not affect the originals - Ctrl-T displays the working image rather than the threshold image). Brightening (left click) will push certain pixels above the threshold level, i.e. turn them 'on', and darkening (right) will do the opposite. Both are applied with the brush, and the amount of brightening or darkening can be modified in the Main Toolbox (F1). The feathering on the edge of the brush can be changed with the 'Soft' slider.

Segmentation mode: With only one segment (see 'segments') this acts much like an 'on' and 'off' drawing tool.

Recalc mode: This recalculates the pixels to which it is applied using the current generation settings in the Generation Toolbox. This manual application of different generation rules to isolated areas.

These three tools allow each slice to be manually corrected. Often using a combination of these and masking (below) fairly automated cleaning can often be devised.

Masking: Masking allows separate rendering and false colouring of different part of the fossil's anatomy. For example, if the fossil possesses a carapace this can be 'masked off' and then removed from the model allowing what is inside to be rendered. Separate parts of the internal anatomy can then be masked and removed, and through iterative improvements an accurate virtual dissection can be performed. Masks are also applied to individual slices using the brush tool.

Creating new masks: Open the Masks Toolbox (F4). Click new, and a new mask is created. Double click on the name to edit, the left colour patch to alter contrast and right to alter the colour itself, the eye to make the threshold within that mask invisible, and the padlock to lock to the mask to further changes. The far left column shows which mouse button applies which mask - clicking here for the desired mask with the desired button assigns this (also possible via drop-down menus in the Main Toolbox). L in the left column indicates the left button applies the mask, R right, and B both. A maximum of 255 masks can be created.

Applying masks: It is often best to begin masking an isolated part, such as a limb. Find such a structure, and adjust the brush size, and then click on the structure. The mask for the button clicked is then applied to this structure, and the threshold changes colour to show this. Moving forward or backward a slice (< and >) then allows this process to be repeated, following structures throughout the dataset. When isolated structures - for example, limbs - join a larger structure - e.g. the body - it is often necessary to define an arbitrary boundary between the two.

Copying masks: This can be done in numerous ways, and all have shortcuts (shown in the masks menu). The selected mask (highlighted in dark grey in the Masks Toolbox) can be copied from the previous slice to the selected, or from the next to the selected (allowing the user to move forwards/backwards a slice, and then use the shortcut to automatically place the current mask on the current slice, to be edited). Large datasets often make it preferable to copy masks ~10 slices at a time. This can be done by selecting the desired slices in the Slice Selection Toolbox, and using the 'Copy from current slice to selected' option. Note this copies only the selected mask, which is not necessarily the 'left click' mask.

It is advisable to create one mask at a time and apply it to all required slices, rather create all masks at once and apply these concurrently, slice by slice. The former often proves quicker, and allows for easier iterative improvement of the mask.

Locking masks: Masks can be locked as described above. This prevents subsequently accidentally overwriting them. Once a mask is complete locking will prevent the mask being modified by either the brush, or by copying masks across slices.

Hiding masks: Masks can be hidden as described above (double clicking on the eye). Choosing 'Hidden masks locked for generation' ensures that any mask which is not shown cannot be altered by the recalc brush or by regenerating the entire slice. This allows one set of generation rules to be created for one structure, which can then be masked off and hidden, and then the rest of the dataset regenerated using another rule - often useful if there are preservational differences between different structures.

Using these tools the entire dataset can be cleaned and an accurate model created. More advanced techniques (below) are sometimes necessary, but for many models the techniques outlined above are sufficient. The model then needs to be rendered.

Rendering: This requires the Output Toolbox (F7). Models are rendered as isosurfaces - meshes of triangles representing surfaces 'draped over' the voxels in a volume. First the settings need to be adjusted (far left tab). Slices/mm and pixels/mm depend on the size of the fossil, and can be calculated based upon a slice image, or taken directly from the xtekt file. In most CT scans, however, these are the same, and the model will render fine with 1.00/1.00 for both of these. KTr limit (not implemented in

the current build of the program) will stop the model from outputting when the defined number of kilotriangles is reached to prevent unfeasibly large models being created. For old PCs 3000KTr (=3,000,000 triangles) is a sensible limit. Higher performance machines can handle 12,000KTr with ease.

First and last file selectors allow definition of the starting and finishing slices for a model. This makes it possible to render just a section of the dataset (e.g. that currently being edited) for faster rendering times. Bin XY and Z allow the final model to be downsampled (note: this is different from dataset downsampling - it purely means the rendered model is reduced in resolution, not the underlying data). This allows for smaller triangle counts and again faster rendering for preliminary models. In downsampled models a single voxel could, for example, contain eight 'full resolution' voxels. Pixel sensitivity dictates for how many such voxels are required for the downsampled voxel to render as fossil.

Objects: The middle tab allows the definition of objects to be rendered; these could be different masks, or segments (see advanced techniques). The user can select the desired mask in the Masks Toolbox, and then click new in the objects tab. This creates an object to be rendered in the 3-D model. Double clicking on the name allows it to be renamed. The colour can also be changed with a double click, and the key used to show and hide the object is similarly edited. Double clicking on the eye stops the object being output. Fidelity dictates how triangles of the isosurface are rendered (e.g. 10% = 1/10 triangles) allowing fast rendering and a crude form of transparency. Different objects can be rendered in a group by selecting them and using the group button. Double clicking on the merge icon on the far right of the object list merges all the objects in a group, rendering them as a single object with the same settings and colour.

Rendering: Once the settings are correct, and all objects have been defined, a model can be output. Shift + F12 creates an .spv file, which is the format for the viewing program SPIERSview. This gives an option of where to save it, and then outputs each object, as shown in the progress bar. Ctrl + F12 automatically loads the output in SPIERSview once the spv has been created. F12 creates a temporary .spv in the source folder and launches SPIERSview automatically.

SPIERSedit 2.0 - Advanced Techniques:

Segments: Creating multiple segments allows the user to create more complex automated generation settings. For example, if a fossil is preserved with both pyrite and void within a siderite nodule the lightest and darkest pixels will be fossil, while those in between will represent the siderite. Using multiple segments allows both these phases to be generated automatically.

Creating segments: The Segments Toolbox is required (F5). All datasets initially have a single segment, 'Default'. Clicking the 'new' button creates another segment. Each segment then needs its own generation rules - these are easiest devised using the range tab in the Generation Toolbox. Here the boxes represent the segments' generation rules - the range of greyscales each box occupies reflect the greyscales which will be assigned to each segment upon generation. The boxes can be moved to different grey levels, coupled (i.e. with no gaps between the different segments) or uncoupled (right mouse button), and any number of segments can be created. Segments can be automatically distributed by selecting all segments to distribute and using the 'Distribute over range' command (Segments menu). This requires the user to define the range (e.g. 0 to 255). Clicking Generate creates the threshold for the two segments, each of which appears as a different colour in segmentation mode. When masking these colours will not be displayed (unless 'Show segments in mask mode' is selected in the mode menu).

Working with segments: Once generation rules have been created and applied to each segment some of the tools behave somewhat differently. Brightening and darkening pixels can move the grey levels between the different segments. Using the segmentation mode allows pixels to be assigned to any of the segments manually, or none, the Segments drop downs in the Main Toolbox allowing the user to define which of these the left and right mouse buttons apply (or by clicking, as with masks, or by clicking on the segment in the 'range' generation tab). Recalc manually recalculates pixels based upon the generation settings for both segments. Segments can be locked in the segments menu, and copied between slices. Creating objects with more than one object requires the user to define which segment(s) to render for each mask. Thus it is possible to render only one segment within a mask, or all. Similarly different segments can be rendered separately as separate objects, even if they are in numerous masks.

Curves: Linear features can be marked on with the use of curves - line drawing tools. These are controlled from the Curves Toolbox (F6). A new curve can be created by clicking new. This curve is given a colour, and by default belongs to no segment. It can be assigned to a segment by double clicking in the segment column, and choosing one. Once a segment is selected, the curve appears as fossil pixels for that segment. *Saving curves: Curves are stored within the .spe file - if many curves are being used it is recommended the user increases the autosave frequency or manually save (ctrl + v) more regularly.*

Placing curves: In curve mode the brush can be used to edit a curve. To initially place the curve the '=' key can be used, which creates two points and two splines to affect the curves between them. These can be moved to fit any shape required. Pressing the '=' key once a curve is created adds another point, and with the addition of multiple points complex shapes can be created. These can be closed, filled or a simple line. Toggling between these modes is achieved by double clicking on the column between colour and segment in the Curves Toolbox.

Copying curves: Curves can be copied between slices using the commands (and associated shortcuts) in the curves menu. Much like masks this allows a curve to be copied onto and edited for each slice. Also like masks, however, it is possible to copying curves over multiple slices. This involves curve interpolation - if a curve exists on slice one and slice ten, selecting slices 1-10 and selecting 'Interpolate over slices' in the curves menu will then copy the curve between the slices. Each slice will show a slight change in the curve so there is a smooth transition from the shape on the first slice to that of the tenth. *Masks from curve: It is also possible to create a mask from a curve - if the curve is not set to a segment the threshold will not be affected by its presence, but choosing 'Mask from curve' in the masks menu will assign pixels to the selected mask on the selected slices based upon the selected curve. Filled curves allow fairly automated definition of complex zones using interpolation, so this can be an efficient method of masking.*

Other tools - slice spacing: The Slice Selector Toolbox automatically displays a list of the images in a dataset, showing the number in the sequence and the filename. The slice spacing menu allows the position of the slice in terms of mm to be displayed. The slice spacing itself (found in the .xtekt file) can be adjusted in the Output Toolbox. This can be offset as necessary using the 'Set slice position' option. Altering the slice spacing can allow the stretching of compressed datasets if the slices are correctly orientated.

Measuring volumes: Volumes of selected masks or segments can be measured using the 'Measure volumes' option in the output menu. This gives the voxel numbers for each object in the Output Toolbox. Volume can then be calculated using the voxel size.

Histogram & Info: The histogram shows the greyscale file for the selected segment on the current image. This can be changed to show greyscale levels for a manually defined area by highlighting this area in Lock/select mode and then selecting 'Histogram shows selected' in the mode menu. The info toolbox shows information such as the position of the brush, which mask it is placed over, and data regarding the source and working images.

Undo & Open Recent: Undo operates as other window software. Ctrl-Z is undo, and Ctrl-Y is redo. Open recent can be found in the file menu and shows recently opened datasets. *Note: Undo only covers some actions – multi-slice operations cannot be undone, and reset the undo stack, so it is best not to over-rely on the undo function.*

Change downsampling: This can be accessed from the file menu (not to be confused with output downsampling and image compression). This is the factor defined upon creation of a new .spe. Often it is preferable to create a model downsampled by 2 on both XY and Z planes, which is significantly faster, and then upsample if greater resolution is needed. When changing the Z downsampling (from 2, for example, when every other slice is used, to 1 where every slice is used) the Lock/Select, Mask, and Curves data can be interpolated onto each newly incorporated slice.

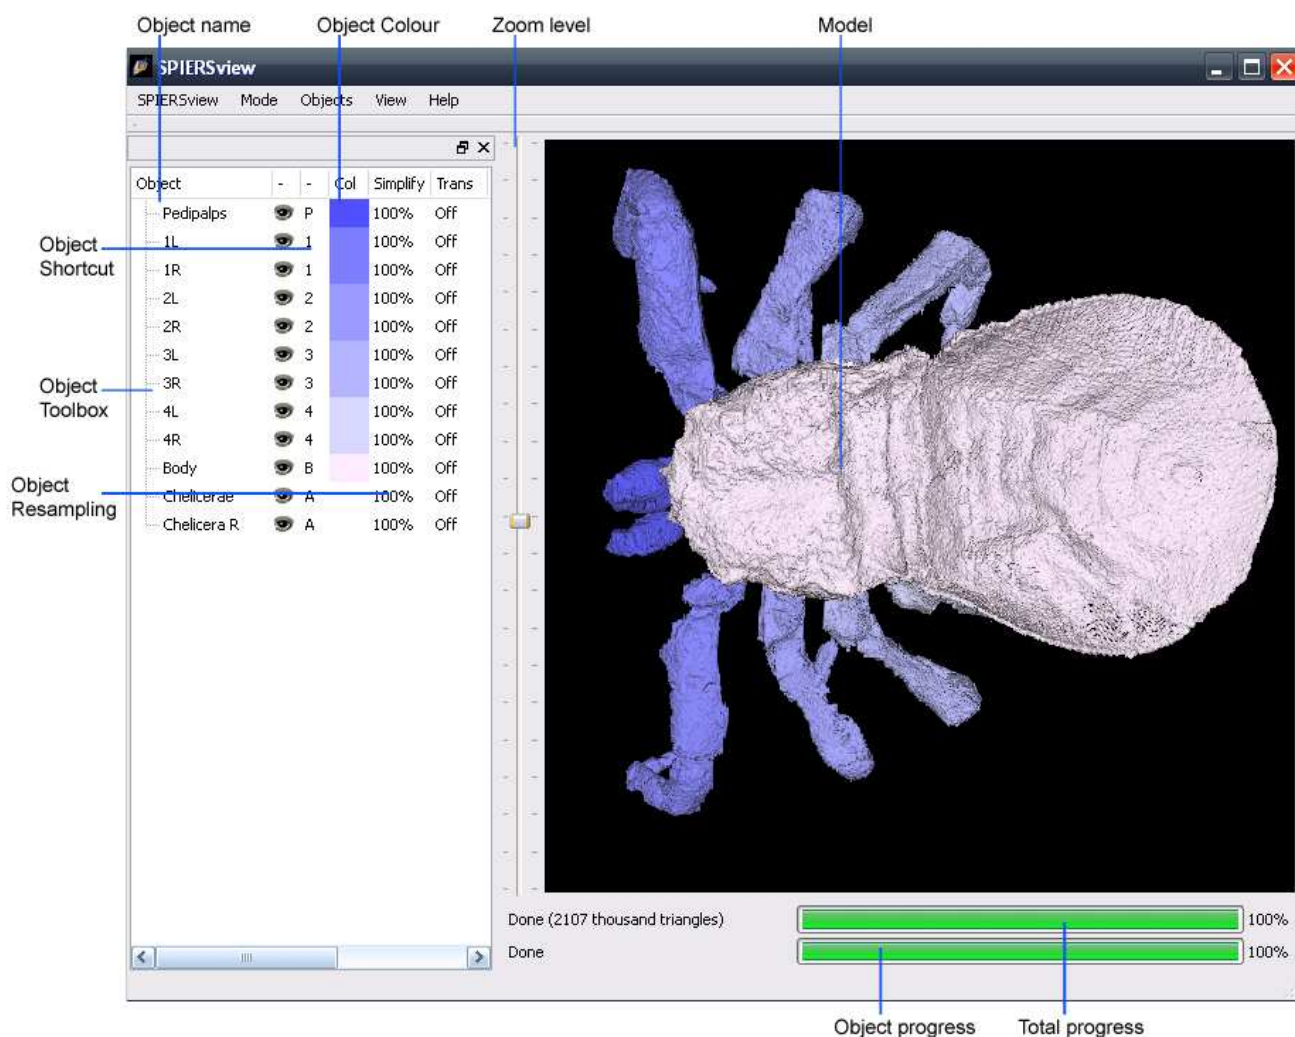
Save as: This allows a second copy of a dataset to be created for backup or any other purposes. *Note: SPIERS stores much of its information as individual files within the .spe subfolder. If a dataset behaves unusually or is corrupt, often recovery can be as simple as creating a new .spe with the correct number of masks, curves and segments, then copying and pasting the contents of the old .spe subfolder to the new one.*

SPIERS Align: This portion of the SPIERS software suite is intended for the alignment and cropping of serial ground datasets. However it can also be a useful tool for preparing CT datasets. It gives the user the ability to load image sequences, and crop them to the desired size, while choosing an appropriate output format. Well-cropped images will ultimately be smaller than their uncropped counterparts, and improve SPIERS' performance. A manual for SPIERSalign can be found in the installation directory.

Using the above instructions complex models can be created using SPIERSedit, to be viewed in SPIERSview, and exported for other software (below).

SPIERSview

Interface: Below the default layout of SPIERSview is shown:



SPIERSview is custom built software for viewing 3-D models from SPIERSedit. It allows the inspection and manipulation of a model, as well as the application of certain filters and ability to export.

Opening an SPV: Double clicking on an SPV file opens it in SPIERSview. Loading progress is shown in the lower two progress bars. As soon as the first object is loaded this can be manipulated while the rest load.

Manipulating models: Models can be manipulated using the mouse - left click and drag translates the model, while right click and drag rotates it. Shift and right click rotates the object in a different plane. The bar on the left can be used to zoom in and out.

Dissections: Objects can be hidden by pressing the letter assigned to them in SPIERSedit, or double clicking on the eye icon in the Object Toolbox. The letter can be edited by double clicking on the next

column, and the colour can similarly be edited. The other columns display settings for filters specified through the Objects menu (below).

Modes: Different modes and their shortcuts can be found in the modes menu. Options here allow the user to view models in anaglyph stereo (red-blue), and split stereo (for stereo-viewers). Strong colours can be muted to create more effective anaglyph models. The model can be set to autospin in this menu, and rotate lock can be applied to prevent the model from being translated (largely for Mac users who only have a single mouse button).

Objects: The objects menu allows the user to group and reorder the different objects shown in the Objects Toolbox. Options include inverting displayed objects, and showing/hiding all (not currently implemented in SPIERSview, but should be in the next release). Further the transparency of objects can be set in this menu, and several filters can be applied. After choosing any filter the model needs to be resurfaced (Mode menu) unless the Auto Resurface option is selected (also Mode menu).

Set resampling: This filter reduces the triangle count of the selected object to the desired percentage (also accessed by double clicking on the 'Simplify' column in the Object Toolbox). This results in loss of resolution, but makes the models smaller, easier to manipulate and easier to handle in other software once exported.

Set smoothing: This smooths the surface of the selected object to the requested degree. This makes the slices from which the model is constructed less apparent, but removes detail from the model.

Set island removal: This removes floating voxels (islands) from the model to the desired degree on the selected object. This filter removes the need to fix floating pixels while slice editing.

View: This menu contains commands for several of the features mentioned previously. It also allows the stereo-separation of anaglyph 3-D images to be adjusted to increase, or decrease the 3-D effect.

Exporting: Using the commands in the SPIERSview menu files can be exported as meshes in two common 3-D formats - DXF and STL. Both can be loaded into most 3-D software. This allows the ray-tracing (creation of images including shadows and surface textures using a light-modelling algorithm) of models in software such as Blender (open source raytracing software) and VG Studio Max (available on the NHM computers). STLs and DXFs can also be loaded into mesh editing software such as Meshlab.

Save presurfaced: This menu also provides the option of saving the SPV with meshes presurfaced. This means the surfaces do not have to be calculated when loading, drastically cutting the processing time on loading. The resulting SPV file is, however, much larger – a 10 million triangle model may be as large as 250Mb.

Screen capture: This allows the user to save what is currently displayed as a tiff, bmp, jpg or png image for use in presentations or publications.

Please feel free to email russell.garwood03@ic.ac.uk or m.sutton@ic.ac.uk with unanswered questions or if any aspects of the software remain unclear. Good luck, and happy model making!

