
SPIERSalign Documentation

Release 3.2.0

Russell J. Garwood, Mark D. Sutton, Alan R.T. Spencer

Feb 10, 2024

CONTENTS

1	Relevant references	3
2	Table of Contents	5
2.1	Introduction	5
2.2	Requirements	5
2.2.1	System	5
2.2.2	Data	7
2.3	What SPIERSalign Produces	7
2.4	Interface	7
2.5	Basic Aligning Procedure	7
2.6	Using SPIERSalign With CT Data	10
2.7	Starting SPIERSalign	10
2.8	Saving, Exiting, and Restarting a SPIERSalign Session	10
2.9	Navigating Datasets	10
2.10	Zooming the Image	10
2.11	Shifting, Rotating, And Scaling Images	10
2.11.1	Shifting	11
2.11.2	Rotating	11
2.11.3	Scaling	11
2.11.4	Reset	11
2.12	Panels	11
2.12.1	Info Panel	12
2.12.2	Markers Panel	12
2.12.3	Crop Panel	12
2.12.4	Auto Align Panel	13
2.13	Automated Alignment	13
2.13.1	Semi-automated Align	13
2.13.2	True Automated Align (Auto Align)	13
2.14	Advanced Navigation	15
2.14.1	Hiding Images	15
2.14.2	Propagation	15
2.14.3	File Locking	15
2.14.4	Swapping Files	16
2.15	Advanced Settings File Usage	16
2.15.1	Loading Backups	16
2.15.2	Constant Autosave	16
2.16	Update System	16
2.17	Keyboard Reference	16

[S]erial [P]alaeontological [I]mage [E]diting and [R]endering [S]ystem: Alignment and data-preparation utility

Main Coding: Russell J. Garwood Additional Coding: Mark D. Sutton, Alan R.T. Spencer Documentation: Russell J. Garwood, Mark Sutton, Alan R.T. Spencer

SPIERSalign is a stand-alone program which enables the user to align (or register) and crop a sequence of images (a tomographic dataset) prior to editing and/or reconstruction with SPIERSedit.



t:@palaeoware

w:<https://github.com/palaeoware>.

RELEVANT REFERENCES

Sutton, M.D., Garwood, R.J., Siveter, D.J. & Siveter, D.J. 2012. Spiers and VAXML; A software toolkit for tomographic visualisation, and a format for virtual specimen interchange. *Palaeontologia Electronica* 15(2): 15.2.5T

TABLE OF CONTENTS

2.1 Introduction

SPIERSalign is a stand-alone program which enables the user to align (or register) and crop a sequence of images (a tomographic dataset) prior to editing and/or reconstruction with SPIERSedit. Although correct alignment is always required for reconstruction, some tomographic datasets (e.g. CT scans) come pre-aligned, and hence this step is not always necessary. Nevertheless, SPIERSalign can be useful for cropping and converting the image format of such datasets. SPIERSalign is primarily a manual alignment tool, which allows the user to move (shift), rotate and rescale each image individually. Automatic alignment tools are provided, but will not always successfully register a dataset. SPIERSalign also provides a facility to define a ‘crop box’ (a region of interest), which will produce a set of cropped images for use in SPIERSedit. Note that alignment is considerably easier and more accurate if there are fiducial marks in the image whose position is invariant between sections. Such marks include edges cut near the specimen or holes drilled through it (Fig. 1).

2.2 Requirements

2.2.1 System

SPIERSalign has no minimum requirements as such, and will run on most systems; it has, however, not been tested on versions of Windows older than Windows XP. The Mac OSX version requires OSX 10.5. Performance will benefit from high processor speed, a fast disk system, and a large high resolution display. A minimum of 512Mb of memory is recommended on Windows XP, 1Gb+ on Vista/7. Memory in SPIERSalign will act as a file-cache, enabling images to load almost instantaneously after the first time they are viewed; A large amount of RAM will enable more files to be cached, which may speed aligning work. A 64-bit version is available; unlike SPIERSedit and SPIERSview however, there are unlikely to be performance or memory usage problems with the 32-bit version on any dataset; the 64-bit version is provided simply to allow a clean 64-bit install of the entire SPIERS suite. Several gigabytes of free space is recommended. A working dataset can be large; for example, high resolution serial photograph datasets can reach 3 Gb in size, and modern CT datasets can be in excess of 15 Gb.




Fig. 1: Figure 1. Fiduciary markings. Top, blocks with edge (left) and holes for as fiduciary markings. Bottom; photos of blocks showing markings in images

2.2.2 Data

SPIERSalign works with windows bitmap format files (ending with a .bmp extension), PNGs (.png), TIFFs (.tif/.tiff) and jpegs (.jpg/.jpeg). Files can be either 24-bit RGB or 8-bit greyscale, and of any resolution. The dataset can also comprise images of different resolutions (although this is not recommended). There are no restrictions on filenames: an identifier followed by a three or four digit number (e.g. mydata0001.bmp) is recommended for clarity. Note that the program will order image files alphabetically, and hence that leading zeroes are required on filenames (e.g. image0001.bmp rather than image1.bmp). All images must be placed in a single directory.

2.3 What SPIERSalign Produces

After you first open a dataset with SPIERSalign it will create a file called settings.txt in the working directory. This file is intended for the program to use – while you can read it, you should not modify it. Settings.txt contains details of the files and the changes you have made to them, and is required until SPIERSalign has finished with the images. When the images have been cropped a new sub-directory will be created, into which the cropped images are placed. Their filenames remain unchanged. The settings text file and working (.xxx) images are retained until you compress the dataset. It is advisable to keep a backup copy of the dataset, and regularly backup the setting files. Currently SPIERSalign also constantly appends to a file log.txt stored in the same folder as settings.txt to provide information in the case of a crash. If this is getting large it can be deleted (in contrast to settings.txt – deleting the latter will wipe your dataset).

2.4 Interface

2.5 Basic Aligning Procedure

A SPIERSalign session in which you wish to align a set of 100 unregistered images will normally involve the following stages:

- 1) Open a series of images and select one (arbitrarily) to act as the datum – all other images will be rotated/shifted/scaled to line up with this one. Typically, the datum image is near the centre of the run; this example will assume that image 50 out of 100 has been selected.
- 2) Set up lines and/or circles on image 50 to mark the fiduciary points (see Markers Panel below). These markers don't actually do anything, but provide visual reference points that can be used as a guide for alignment.
- 3) Move to image 51 (see *Navigating Datasets*), and perform rotate, shift and if necessary rescale operations until the image position fits the fiduciary markings (see *Shifting, Rotating and Scaling Images*). It is advisable also to flick repeatedly between image 50 and 51 to check visually that alignment has been achieved.
- 4) Repeat step 3 for image 52 (checking against image 51), and continue until the end of the run (image 100) has been reached.
- 5) Return to the datum image (50), and repeat steps 3 and 4, this time going backwards through the sequence (aligning 49 to 50, 48 to 49 etc.)
- 6) When happy with alignment (it is advisable to run through the entire sequence to ensure that all movement is as smooth as possible), enter crop mode (see *Crop Panel*) and resize it, ensuring that the object of interest is contained within the box on every image.
- 7) Finalise this process by selecting crop. SPIERSalign will process the images, producing sub-directory with the aligned and cropped images ready for SPIERSedit.




Fig. 2: Figure 2. Interface in aligning mode




Fig. 3: Figure 3. Interface in cropping mode.

2.6 Using SPIERSalign With CT Data

Data from CT scanners is pre-aligned, but use of SPIERSalign to prepare CT data is nonetheless encouraged, as it allows you to crop the dataset easily for SPIERSedit, creating smaller working images and hence faster image processing. With CT or other pre-aligned data, steps 1-5 in previous section can be omitted.

2.7 Starting SPIERSalign

Start the program, and select *Open* from the File menu. Navigate to the folder containing your dataset and press ok. Alternatively you can choose a recently opened dataset from the list shown in Open Recent on the File menu. If this is a new dataset SPIERSalign will ask if you want to check that image dimensions are the same; this is recommended for all but expert users. After performing this check, the program will then display the first image of the sequence.

2.8 Saving, Exiting, and Restarting a SPIERSalign Session

To exit an (unfinished) SPIERSalign session simply close the program window; this saves all settings and exits. The settings file can also be saved at any point by choosing Save on the File menu (*Ctrl-S*). It is recommended the settings are saved every few images. To start again at a later date, open (as above) the same working directory. While working with a dataset, a backup settings file can be saved at any point with the Save Backup command on the File menu (*Ctrl-Shift-S*). This will create a time-stamped settings file of the form “settings backup - 16 Sep 2008 - 15.42.50.txt”. Such a backup can be loaded at any time and applied to the dataset (see *Loading A Settings File* below).

2.9 Navigating Datasets

To move forward or backwards through the dataset, press ‘.’ (forward) or ‘,’ (backward) - these are on the ‘<’ and ‘>’ keys on a UK keyboard. You can also drag the horizontal slider at the base of the window (see Fig. 2) to move through the image set, or type the number image you want to skip to into the spin box on the bottom right.

2.10 Zooming the Image

Ctrl = allows you to zoom in, and *Ctrl -* is zoom out. *Ctrl-R* displays the working image at 100% (one pixel in image = one pixel onscreen), and *Ctrl-Shift-R* fits it to the window (normally the correct thing to do). These zoom commands are also all available from the *Zoom* menu. They do not rescale the underlying image, merely change the magnification with which it is displayed onscreen.

2.11 Shifting, Rotating, And Scaling Images

Registration work consists of moving images so they line up properly; this is done by Shifting, Rotating and Scaling. Commands to do this exist on the *Transform* menu, though normally the keyboard shortcuts (given below and on the menu) should be used rather than the menu commands.

2.11.1 Shifting

Each individual image can be shifted (moved) left, right, up or down by using the arrow keys. By default a single key press will move the image by one pixel. Pressing *Ctrl* and the arrow key will move the image by half a pixel. *Ctrl* + *Shift* an arrow key will move it ten pixels.

2.11.2 Rotating

Each individual image can also be rotated about its centre either clockwise or anticlockwise. To rotate use either '0' (1° clockwise) or '1' (1° anticlockwise). Using *Ctrl* and these keys will rotate by 0.1°, and *Ctrl* + *shift* will make the rotation 10°.

2.11.3 Scaling

Each individual image can be rescaled (made larger or smaller). Ideal datasets should not need scale correction, but in the experience of the authors small scale errors of the order of 1% often occur in some forms of data. Normally the image should be rotated and shifted first, and only then rescaled if it is clear that this is necessary. To increase the size of the image use ']'. To decrease it, use '['. Again, *Ctrl* will rescale by a smaller factor, and *Ctrl-Shift* by a larger one.

2.11.4 Reset

The Reset Image command on the Transform menu (Ctrl-Shift-Z) resets all rotate, scale and shift operations on the current image. A Reset Image button can also be found on the Auto Align panel (see below). A Reset Scroll Bars around Current Image command is also provided, which may be useful in datasets comprising numerous resolutions; if used frequently this may impede alignment by hampering comparison between slices.

Note: The resampling process involved in rotating/scaling/sub-pixel shifting images will produce slight blurring. This is unavoidable, and is very unlikely to cause a problem. However, repeated rotating, scaling etc. of a single image can in theory cause a build-up of this blur effect. SPIERSalign normally avoids this problem; every time a rotate, scale or shift is performed the program goes back to the original image and does a single operation combining all rotates/scales/shifts specified so far. However, once a cropped dataset has been created (see below) the changes are permanently applied to these images. For this reason, it is NOT good practice to realign images in a dataset which has already been cropped – this may result in some images being rotated/scaled/shifted twice, introducing more substantial blurring. If a user feels that they may want to go back to a dataset to re-align it, compress it and then if necessary reload the settings file. Note that this does not apply to whole-pixel shifting (using the arrow keys), which does not introduce any blur and can be done as often as required.

2.12 Panels

SPIERSalign uses *Panels* (dockable windows) for much of its interface (see SPIERSedit and SPIERSview manuals). Available panels are *Markers*, *Crop*, *Auto-Align* and *Info*, shortcut keys F1, F2, F3 and F9 respectively.

2.12.1 Info Panel

The *Info* panel (see Figs 2,3) shows the user what mode they are in - e.g. *Normal Align*, *Automated Align*, *Semi-automated Align*, *Propagate*, *Lock*, as well as the cursor's position in pixels from the top left corner, the file name of the current image, and the image's dimensions. It also provides a 'text box' in which the user can enter notes regarding the aligning process; these notes are automatically saved along with the settings file.

2.12.2 Markers Panel

See Figure 2. The Markers panel provides a list of all markers, plus facilities to add, remove and edit them.

The *Add* and *Remove Marker* buttons allow you to vary the number of markers (there is a lower limit of five). The *Red*, *Green* and *Blue* dials dictate the colour of the markers; this should be chosen to maximise contrast with the images. Note that the currently selected marker is always the inverse colour of the other markers. The marker size spin box allows you to change the size of the circular markers, while the line thickness changes the width of the line for both linear and circular markers.

Markers can be selected by clicking close to them on the image. They can then be dragged into position. They can also be selected by clicking on the markers list in the markers dialogue, which shows the selected marker highlighted in grey. The toggle shape button transforms the marker from a circle to a line or back again. Lines can be moved by clicking and dragging near the middle of the line; clicking and dragging towards the end of the line allows you to position the endpoints as necessary to match your fiduciary markings.

The *Lock Markers* button stops accidental movement of markers; they can be unlocked with a second click of the button.

Automarkers buttons (*Auto Markers*, *AM: Align* and *AM: Edit Grid*) are dealt with under the *Semi-Automated Align* section below.

2.12.3 Crop Panel

Crop mode: Bring the *Crop* Panel up puts SPIERSalign into crop mode (see Fig. 3). A white rectangle is overlaid on the image; this represents the 'crop box'. A message will appear in the status bar with the dimensions of the crop area; this is also provided in spin-boxes on the Crop panel, where it can be entered precisely. The crop-box defines the region of interest. Everything outside of this box will be removed from the finalised files that SPIERSalign creates in the 'cut' directory (i.e. the ones that SPIERSedit will work with). Users should ensure that all object(s) of interest are contained inside the crop-box **on every slice**. To check this navigate backwards and forwards in the normal way whilst in crop mode. The crop box can be both resized and moved using the mouse, but for fine adjustments shortcut keys can be used. Crop-box colour can be set in the Crop panel to maximise contrast with the image.

Cropping a Dataset: Once all images have been aligned and a crop-box has been set, the Crop button on the Crop panel initiates the cropping procedure. SPIERSalign will ask for first the starting and then the finishing slices of the crop (this allows users to exclude regions at the start or end of a dataset that do not contain the object(s) of interest). SPIERSalign then asks for the file format to generate; this can be JPEG, PNG, BMP or TIFF. Use of PNGs or BMPs is recommended in most circumstances; PNG files are relatively small and are compressed with a lossless algorithm. Saving as a JPEG produces extremely small files but involves 'lossy' compression which may subtly damage data; while artefacts are likely to be minimal, use of JPEGs is not recommended unless disk-space is an overriding concern. BMPs are saved in uncompressed format, and are typically over three times the size of PNGs; they may, however, load faster in SPIERSedit on systems with high-speed storage systems as they do not need to be decompressed. Thus they are the best choice in some circumstances. TIFFs cannot be imported into SPIERSedit, but may be required by other tomographic software. Once the format has been selected the program will work through the dataset, cropping the images to the required size and placing them in a 'cut' folder in the working directory. This process may take a few minutes. The images can then be checked and loaded into SPIERSedit as a new dataset. If there are any problems with the images, crop can be performed again until you are happy with the dataset. *Compressing A Dataset:* Cropping will not alter the original image stack, remove the temporary .xxx images, or delete the settings file. This allows a revised alignment and second crop, or definition of a second region of interest. If you wish to remove these files once you have

finished with the dataset (the temporary images can take up a lot of disk space), use the Compress Dataset command on the File menu. This command removes the temporary images (but not the settings file) and also converts the original images to PNG format. The PNG images will be exactly the same resolution and quality as the originals, but occupy less storage space. The settings file can be reapplied at a later date if further changes are required.

2.12.4 Auto Align Panel

See *True Automated Align* section

2.13 Automated Alignment

2.13.1 Semi-automated Align

The *Semi-automated Align*, or automated markers system is a fast way of aligning images using the markers system. To do this markers should be placed to pick out fiduciary markings and features of interest, as outlined above, on the datum image. Automated markers should then be toggled by clicking the *Auto Markers* button in the *Markers* panel. This places a grid over the markers and underlying images (Fig. 4A). The user can then navigate through the dataset. When viewing an image requiring alignment, clicking and dragging the mouse can manipulate the automarker grid. Dragging the corners will rotate the grid, and anywhere near the centre will physically drag it (Fig. 4B). The markers will move with the grid, allowing them to be aligned with the features of interest or fiduciary marks in the underlying image. When the markers lie in the correct position relative to the fossil, the button *AM:Align* (keyboard shortcut Alt-A) aligns the slice (Fig. 4C). This is done by returning the grid and markers to their original positions – which were defined on the datum image – and applying the same ‘correcting’ transformation to the current image. This process can then be repeated moving through the dataset.

Semi-automated alignment is best suited to datasets with drilled fiduciary holes or no fiduciary features; markers can be placed over each hole or parts of the specimen and the slice corrected; an entire dataset can be rapidly in this way, although subsequent finer adjustments are likely to be required.

Clicking *AM>Edit Grid* in the *Markers* panel allows the user to change the size, position and line thickness of the grid, in addition to the number of cells. Clicking the *Auto Markers* button in the *Markers* panel will turn off automated markers, allowing markers to be readjusted and used as usual. While using *Auto Markers* it is not possible to add or remove, or adjust the position, colour thickness or shape of the markers.

2.13.2 True Automated Align (Auto Align)

SPIERSalign provides a facility for the automated alignment of images based upon straight edges within the image. This algorithm will only work on datasets which possess two non-parallel fiduciary edges throughout the zone to be aligned. The system has its own panel (*Auto Align*, F3) showing the options for auto aligning - the panel should open in the left docking area when a dataset is loaded. Before alignment can occur, the user should click the *Setup* button; this displays two boxes on the image, one horizontal and one vertical. These should be placed over the edges upon on which the alignment is to be performed by resizing, dragging and rotating the boxes. These ‘edge zones’ should cover as much of the fiduciary edges as possible, on every image to be aligned, but without overshooting the edge on any image. Neither box should be rotated more than sixty degrees from their original position. They can go over the edge of the image (i.e. the top, bottom left or right boundary of the file) but should not be placed entirely off the image (Fig 5).

Once these zones are correctly positioned, the user can exit by clicking setup again. An instruction box will appear, explaining the next steps and giving the option of a *Verbose Mode*. This is usually not necessary – if chosen it will display the line equations for the detected edges on both datum and subsequent images, and the required adjustment of the current image. The user can then highlight the slices to be aligned in the *Auto Align* selection box. Clicking *Align* in the *Auto Align* panel will then align selected slices against the currently displayed, datum image. Often the setup boxes

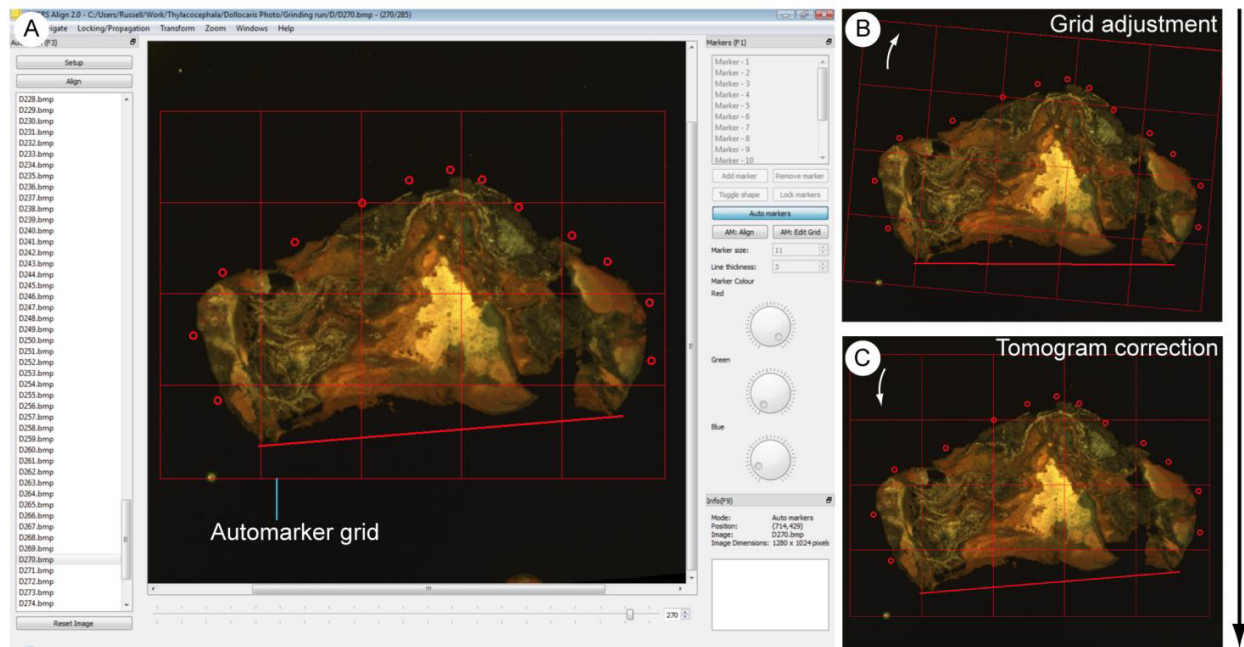


Fig. 4: Figure 4. The default user interface of SPIERSalign with semi-automated alignment toggled. Note the grid in addition to the markers; this is shown superimposed on the reference tomogram. B. An unaligned tomogram, with the grid rotated (white arrow) to align the markers with the boundaries of the fossil. C. The same tomogram, aligned by resetting the grid and transforming the underlying image.

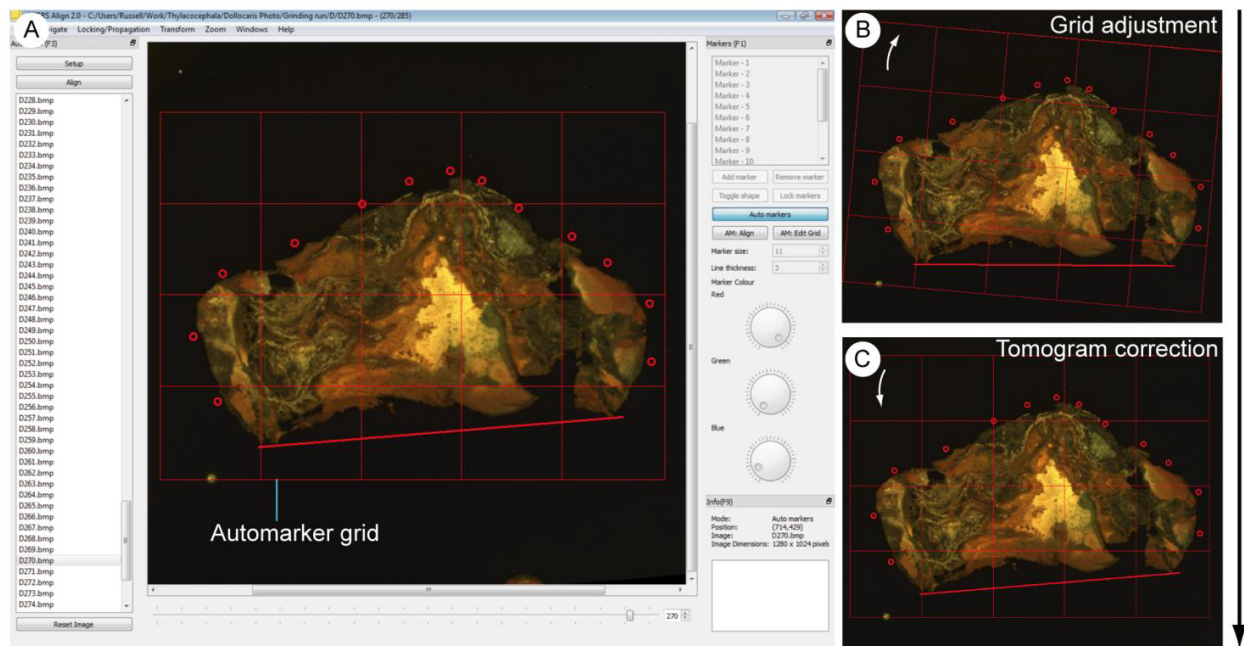


Fig. 5: Figure 5. Auto align in setup mode showing the correct placement of setup boxes.

will have to be quite large on the initial align to satisfy the conditions above, providing sub-optimal aligning. *Auto Align* can be run repeatedly, however, with increasingly tighter setup boxes around the fiduciary edges. This will allow a more accurate alignment, and can align many datasets so they only need minor corrections after three such iterations.

This algorithm only works on straight edges (more details on how it works are available: <https://spiral.imperial.ac.uk/handle/10044/1/6334>). It does not take into account any scale errors in images. If any of the images are very poorly aligned, the *Reset Image* button will return the image to its original position. Very noisy images could hinder its performance, as can poorly positioned setup boxes – i.e. those which feature a large area without an edge, and those at a low angle to each other.

2.14 Advanced Navigation

2.14.1 Hiding Images

Occasionally users may find an image in the sequence that they wish to hide temporarily (for example they wish to visually compare images 10 and 12 by flicking between them, without 11 always appearing in the way). To hide the currently viewed image use the *Hide Image* command on the *Navigate* menu. The image will remain on screen (the status bar notes that it is hidden), but if the user navigates forward and then back again with the ‘.’ and ‘,’ keys they will find that the hidden image is skipped. Using *Ctrl-Shift-‘.’* and *Ctrl-Shift-‘,’* to navigate will move to the next/previous image whether or not it is hidden; hidden images can also be accessed using both the spin box and horizontal file slider. To unhide the currently viewed image use the *Hide Image* command on the *Navigate* menu. An *Unhide All* command is also available on the *Navigate* menu.

Hidden images are still processed during cropping and propagation; it is advisable to perform an *Unhide All* operation before cropping to ensure that objects of interest do not extend outside the crop-box on hidden images.

2.14.2 Propagation

SPIERSalign provides a facility to record changes made to one image, and then apply those changes to all subsequent or previous images in the sequence. This feature is useful, for instance, where one image has been omitted during alignment and all of the changes made to this need to be applied to the subsequent (aligned) images. To use this mode select the *Propagation Mode* item on the *Locking/Propagation* menu **while viewing the image from which you want to apply the changes**. SPIERSalign will ask you whether you want to propagate forwards (to the end of the dataset) or backwards (to the beginning). Propagation mode limits navigation to the propagation image and one either before it (forwards propagation) or after it (backwards propagation). Only the propagation image can be rotated, scaled or shifted. Once all changes to the propagation image have been made, the *Apply Propagation* command on the *Locking/Propagation* menu will then apply the changes to the required images.

Propagation mode (forward flavour) can be used on the first image of a dataset to make changes to the entire dataset; for instance to adjust size or orientation.

2.14.3 File Locking

SPIERSalign provides a ‘file locking’ feature (*Lock Forward* or *Lock Backward* on the *Locking/Propagation* menu). When activated, this mode restricts navigation to just two images, and only enables rotations, scales or shifts to be applied to one of these. *Forward Locking* allows changes to be performed on the second image of the pair, and *Backward Locking* on the first. The *Move Forward/Back* command on the *Locking/Propagation* menu moves the pair either forward or back in the dataset, depending on the mode. This mode is intended as an aid to alignment, enabling a user to work through the dataset one slice at a time without accidentally moving the wrong image. Note that when viewing the locked image of the pair, the *Lock File* option in the *Locking/Propagation* menu will be ticked; unticking this temporarily overrides the locking.

2.14.4 Swapping Files

Occasionally you may find that two images are the wrong way round in the sequence, and need to be swapped over. The *Swap Image* command on the *Transform* menu will swap the current image with the next one. Note that it is possible to move images to any desired position in the sequence by successive swap operations.

2.15 Advanced Settings File Usage

2.15.1 Loading Backups

A backed-up settings file (see above) can be loaded and applied to the dataset using the Load Settings File command on the File menu. This operation overwrites the current settings file (which should be backed-up beforehand), and applies the settings to the image list. This will only work with the original dataset; if there have been any changes to the file names or any additions to or deletions from the dataset, it will fail.

2.15.2 Constant Autosave

If SPIERSalign crashes after a long period of work, it is possible that changes since the last save will be lost. This loss may not be apparent until you next try and transform any given image, at which point the new transformation will be applied to the original image, not the altered .xxx file which will have been loaded on initial viewing. To guard against this, the Constant Autosave option in the File menu is provided; with this selected SPIERSalign will resave settings.txt after every transformation, removing the risk of losing unsaved work. This should make no noticeable difference in performance on all but the oldest PCs, and so is checked as standard.

2.16 Update System

SPIERSview incorporates an auto-updating system (as do the other two SPIERS applications). If connected to the internet, it will check for updates every time it is run. If it finds an update it will give you the option of downloading and installing it. Agreeing to this will close SPIERSview once the updater has been downloaded.

Note: this part of SPIERS is under review as it does not necessarily work across all operating system. This functionality may change or be removed in the future.

2.17 Keyboard Reference

Keyboard shortcuts are given in the menus for all commands that have them. The following table is intended as a handy reference to augment these.

Ctrl + O	Open file
Ctrl + S	Save settings file
Ctrl + Shift + S	Save time-stamped backup of settings file
.	Next image
,	Previous image
Ctrl + .	Lock forward
Ctrl + ,	Lock back
/	Move forward/back (depending on lock mode)
Ctrl + Alt + -	Hide image
Ctrl + Alt + =	Unhide image
Ctrl + Shift + =	Unhide all
Ctrl + Shift + .	Advance to hidden (also use spinbox & slider)
Ctrl + Shift + ,	Retreat to hidden (also use spinbox & slider)

Transformations:

Ctrl + shortcut key always smaller than default transformation	
Ctrl + shift + shortcut key always greater than default	
0	Rotate clockwise
1	Rotate anticlockwise
]	Enlarge
[Shrink
Up Arrow	Move image up
Down Arrow	Move image down
Left Arrow	Move image left
Right Arrow	Move image right
Ctrl + Shift + Z	Reset image
Ctrl + Alt + Z	Reset graphics window (i.e. scroll bars etc.)

Crop Area	Mouse preferable for all but very small adjustments
Markers	Mouse always preferable.
Ctrl + =	Zoom in
Ctrl + -	Zoom out
Ctrl + R	Zoom 100%
Ctrl + Shift + R	Fit image to window

Note: On some international keyboards some of these shortcuts may not work – for example on German laptop keyboards ‘=’ and ‘0’ are on the same key, disabling the zoom in and some rotate functions. If users experience problems relating to keyboard setup, please contact the authors, who may be able to supply a custom-version of the program with different keyboard shortcuts.