Image registration in DataViewer

A brief user guide

Last modified: May 2014, version 1.5.1.2

Starting from version 1.5.0.0, DataViewer provides tools for image registration for both 2D and 3D.

In both cases, only rigid transformation is considered. Two sets of images are required as input: a reference image which stays stationary, and a target image which needs to be transformed to match the reference image.

1. In a nutshell

The aim of the registration option is to provide a basic tool for volume registration. The emphasis is more on visualization, speed and interactivity, rather than a fully automatic registration which can be extremely challenging.

There are three important components in a registration algorithm: transformation (rigid or non-rigid), similarity/dissimilarity measure (or matching criteria) and the searching algorithm. Currently, only rigid transformation is used to transform the target image. More specifically, the rigid transform contains x/y/z translations and 3D rotation which is expressed internally using Euler angles (z-x-z convention). There are many different similarity/dissimilarity measures for registration. For the time being, we use sum of square difference as criteria, which is suitable for images coming from the same modality (e.g., CT) and when the 2 images are not too different in geometry and intensity from each other. For the searching engine, we use Powell's method (see “Numerical recipes in C++”) to find the transformation parameters.

A true 3D registration algorithm searches for the set of parameters (in our case, 3 translation and 3 rotation parameters) which gives the minimum difference between target image and reference image. A basic operation in the searching is the calculation of image difference. The searching engine determines how the parameters should be adjusted according to the image difference. The number of this basic operation needed increases strongly with the number of parameters to search.

To speed up registration, we use a pseudo-3D registration method as default, while the true 3D method is provided as a reference. Instead of searching all 6 parameters simultaneously using the 3D volumes, we break the searching task into smaller and easier steps. The viewing option in DataViewer displays three orthogonal views at any given location: the x-y view (transaxial view), the x-z view (coronal view) and the z-y view (rotated sagittal view). We perform a 2D registration (two shifts and one rotation) on each of the views iteratively to achieve volume registration. One loop, or one iteration, consists of three 2D-registration steps: matching on x-y view, then on x-z view, and then on z-y view. This procedure can be repeated to refine the registration results. The number of iterations/loops is default to 3, which works fine in most cases. One can change this number in the options. Internally, the progress is monitored: if the improvement becomes very small, the registration will stop automatically before the number of loops is done. Due to the decomposition of a 3D task to many 2D tasks, the registration can be done almost instantly in a semi-automatic fashion.

In the options dialog, one can switch between pseudo-3D method and true 3D method easily. In general, the pseudo-3D method is very flexible and is much faster. The result is often fairly good. The true 3D method does refine the registration to a certain degree by reducing the difference slightly more. However, the difference is often barely visible.

For registration to succeed, it is important to choose a proper region for comparison. The Volume-Of Interest (VOI) can be defined interactively. The VOI can be shift/resized with mouse action. It is also editable with right-mouse click on the VOI itself. The smaller the VOI is, the faster the calculation will be. The VOI is chosen based on intuition: for the automatic registration to succeed, it should contain similar and clear structures in both volumes.

The 3D registration module in DataViewer accepts both 16-bit (grey value ranging [0, 65535]) and 8-bit ([0, 255]) images, although images displayed on the screen are always 8-bit. Color images containing RGB channels can be accepted in limited cases. However, they are always converted to gray values by loading. The method of RGB-to-gray converting can be defined in the preferences: if original is chosen there, the method is default to “luminance”.

To maximize flexibility, a 3-layer data structure is used internally: bottom, middle and top layer. The top layer is the images displayed on the screen; the middle layer is what the registration algorithm will use effectively; and the bottom layer is what the program keeps in the memory. The bottom layer is the raw data of both target and reference images, which are loaded into memory either in 16-bit or 8-bit as it is in the original images. Resampling and undersampling may have occurred during loading, as explained in the paragraph below. At this layer, no image transformation takes place. The middle layer is always 8-bit and is extracted from the bottom layer with image transformation applied to target image. In case of pseudo-3D method, only the 3 orthogonal views are extracted and kept. In case of true 3D, the full volume needs to be extracted. User defined dynamic range, with the guidance of a corresponding histogram, is applied to convert 16-bit/8-bit raw data to 8-bit data. This allows a user to choose image contrast properly. Images at middle layer are the basis for image registration and profiling. The top layer is the image displayed on the screen. Artificial colors (color look-up-table), contrast stretching, and zooming are applied to the 8-bit middle-layer-images.

For registration, the target image and reference image are not required to have the same matrix size or pixel size. This implies that re-sampling is applied if necessary . At loading time, a common matrix is calculated based on one simple rule: just large enough to contain both image sets, with the chosen pixel size (by default, the pixel size in the reference image). In 2D case, no down-sampling is applied. In 3D case, both image volumes are loaded into computer memory. Therefore, down-sampling is often necessary. The down-sampling factor is kept as an integer number and is determined by the "max. preview dimension" as defined in options (if N = max. preview dimension, the maximum number of voxels is then N^3), and the available physical memory on the PC. Noting that 3D registration is very memory-demanding, it is therefore very important to install a 64-bit DataViewer on a 64-bit computer: a 32-bit application has the intrinsic limitation of maximum usable memory (2GB).

The registered images can be visualized immediately. There are several ways to examine it:
- Show the reference and target images alternatively on the screen by clicking on the line showing file names, or with keyboard keys ↑ , ↓, Page Up or Page Down.
- Show fusion images. Fusion is done with alpha blending. Simply explained: each grey-value pixel is mapped into a so-called RGB-triplet (R: red; G: green; B: blue) to be displayed on the screen; if we denote a RGB-triplet pixel in target image as (R_t, G_t, B_t) , and the corresponding pixel in reference image as (R_r, G_r, B_r), the pixel will be displayed as $(\alpha R_t+(1-\alpha)R_r , \alpha G_t+(1-\alpha)G_r, \alpha B_t+(1-\alpha)B_r)$, with $\alpha \in [0., 1.]$. The mapping between a grey value and a RGB-triplet is done using a look-up-table (there are a few of them in DataViewer, grey, color 1, color 2, etc., w/o inverting, w/o contrast stretching). The colors are chosen in different ways. In fusion mode *a*, if the user has selected different color for each image, these colors are used; otherwise a set of pre-defined contrast colors are used instead. In fusion mode *i*, inverted grey colors are chosen to enhance the differences between the 2 set of images.
- Draw a profile on the fusion image. This draws a profile on each of the two images and shows the difference at pixel level: not only in geometry, but also in image intensity. As mentioned above, the profiles are extracted from the "middle layer".

2. 2D registration.

This function is activated via menu Actions->2D registration, as shown in fig. 2.1.

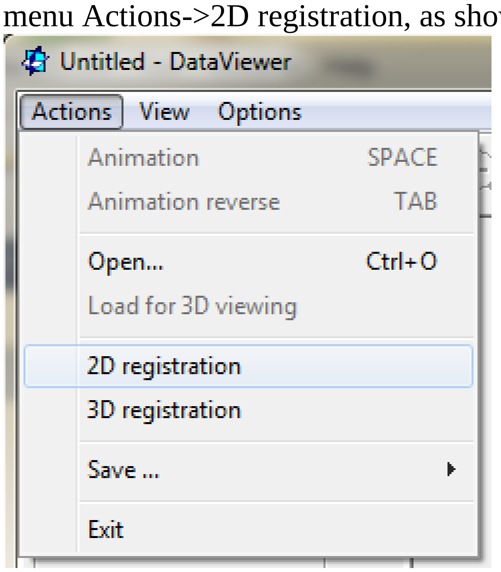

Fig. 2.1 How to activate 2D image registration function.

The 2D registration window, as shown in fig. 2.2, will appear and the normal image control window will be toggled off. For the first time use of this function, you might need to re-adjust the window size so that all controls are shown properly. This window is a "floating window", meaning that it cannot be attached to the main window, but it can be toggled on and off with "Actions->2D registration" at any time. Even if the registration window is toggled off, the registration procedure maintains its status until another task (e.g., open a new dataset, start 3D registration) is started explicitly.

For registration, two images are needed: a reference image and a target image. During registration, the target image is transformed into the coordinate system of the reference image which stays stationary. By double-clicking the corresponding line in the list with left mouse button, one can browse to select the corresponding image. Once both images are selected, one can load by clicking on the Load button.

The screen image shown above is when the reference and target images are loaded successfully. Once loaded, the target image and reference image can be displayed either individually or as fused image, by clicking on the corresponding line in the list. Once highlighted, you can use Page Up and Page Down button to switch. In the fusion mode, a slider is provided to vary the weights (blending). The images are shown in the main window. When target/reference image is shown alone, one can change the color options and the dynamic range using the histogram shown.

To visually examine how good the registration is, one can alternate between target and reference images, or view the fusion images with various settings. Alternatively, you can also check profile along any line (right-mouse click and drag).

The target image can be shifted and rotated manually using the arrow buttons, or keyboard shortcuts (see below for detailed key combinations). You can alter the shift/rotation steps using button Options. More conveniently, you can use CTRL + left-mouse-button to rotate, or use CTRL+SHIFT+left-mouse-button to shift. Do make sure you are on the fusion image so that you see the effect immediately.

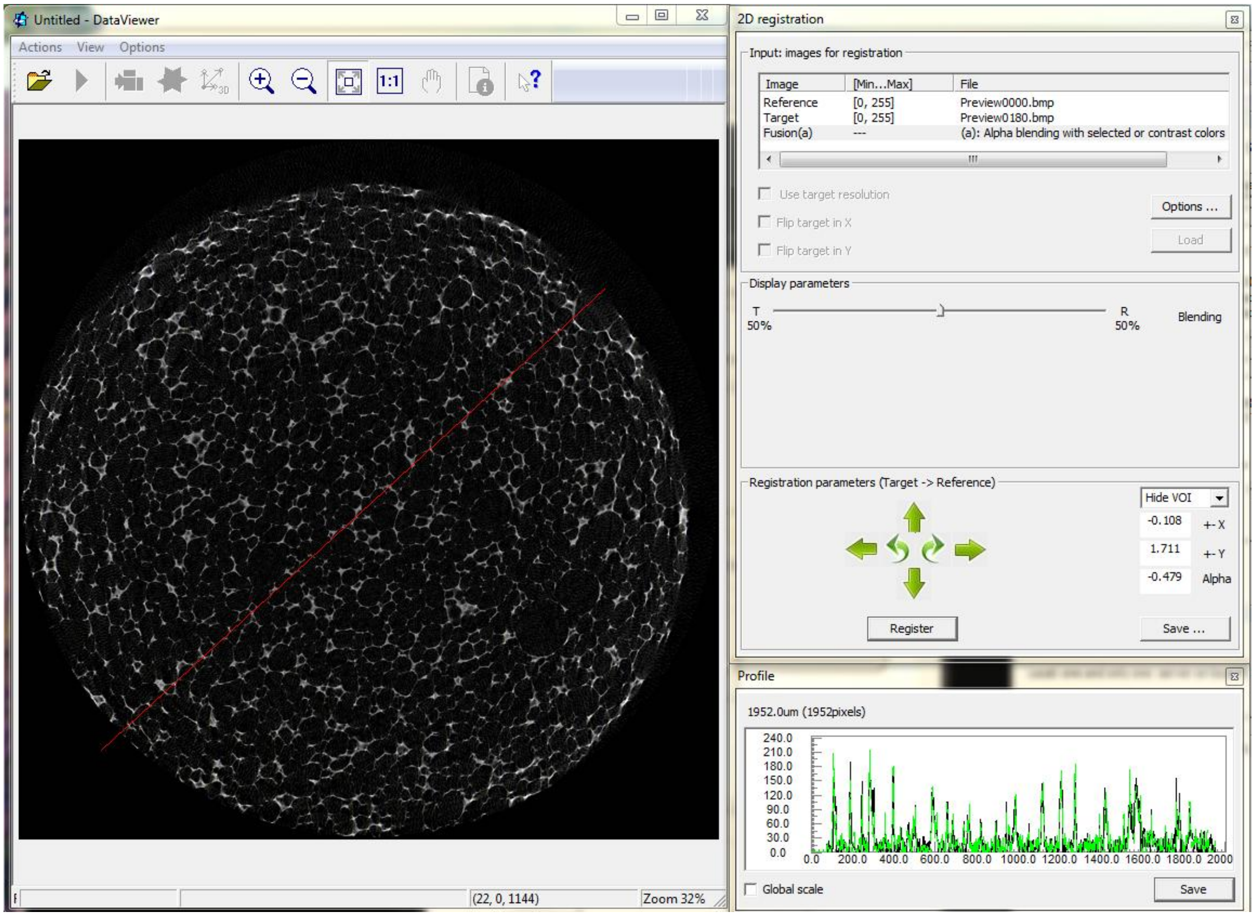

Fig. 2.2 Main window for 2D image registration.

The registration can also be done automatically using buttons Register after selecting VOI carefully. It is generally recommended to shift/rotate manually first so that the 2 images are sufficiently close to each other, before using the automatic function. You can alter the searching range using button Options.

Once the images are registered, a saving function with button Save... will allow you to save the transformed image. It saves optionally also the reference image. This makes sense if there is resampling during loading. A difference image can be saved too. The grey values in the difference image are calculated as such:

Diff = (target – reference) + 128

Note that difference smaller than -128 or larger than 127 is truncated to keep values within [0, 255]. The resulting images will be saved in the target folder with "_tar", "_ref" and "_diff" attached to the file name.

Keyboard shortcuts:

(For the keyboard shortcuts to work, click at an empty place on the main window so that it is on focus to respond to the shortcut signals.)

`Shift` + ←: shift target image to the left (x-axis) by one step.
`Shift` + →: shift target image to the right (x-axis) by one step.

`Shift` + ↑: shift target image further (y-axis) by one step.
`Shift` + ↓: shift target image closer (y-axis) by one step.
`Ctrl` + `Shift` + ←: rotate to the left (anti-clock-wise) by one rotation step.
`Ctrl` + `Shift` + →: rotate to the right (clock-wise) by one rotation step.
`Ctrl` + left-mouse-button: rotate to any orientation.
`Ctrl` + `Shift` + left-mouse-button: shift to any position at any direction.

3. 3D registration.

This function can be activated via menu Actions->3D registration, as shown in fig. 3.1.

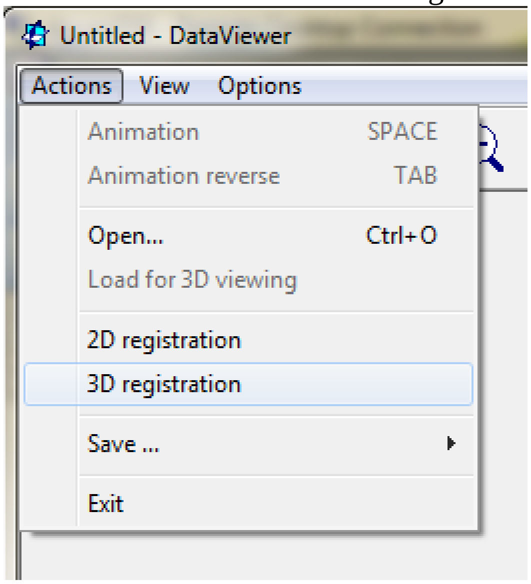

Fig. 3.1 How to activate 3D image registration function.

The 3D registration window, as shown in fig. 3.2, will appear and the normal image control window will be toggled off. For the first time use of this function, you might need to re-adjust the window size so that all controls are shown properly. This window is a "floating window", meaning that it cannot be attached to the main window, but it can be toggled on and off with "Actions->3D registration" at any time. Even if the registration window is toggled off, the registration procedure maintains its status until another task (e.g., open a new dataset, start 2D registration) is started explicitly.

For registration, two datasets are needed: a reference image volume and a target image volume. During registration, the target image is transformed into the coordinate system of the reference image which stays stationary. By double-clicking the corresponding line in the list with left mouse button, one can browse to select the corresponding image volume. Once both volumes are selected, one can load by clicking on the Load button.

The screen image shown in fig. 3.2 is when the reference and target image volumes are loaded successfully. Once loaded, the target image and reference image can be displayed either individually or as fused image, by clicking on the corresponding line in the list. Once highlighted, you can use Page Up and Page Down button to switch. In the fusion mode, a slider is provided to vary the weights (blending). The images are shown in the main window as 3 orthogonal views: the x-y view (transaxial view) at bottom-left, the x-z view (coronal view) at top and the z-y view (sagittal view) at right. When target/reference image is shown alone, one can change the color options and the dynamic range using the histogram shown.

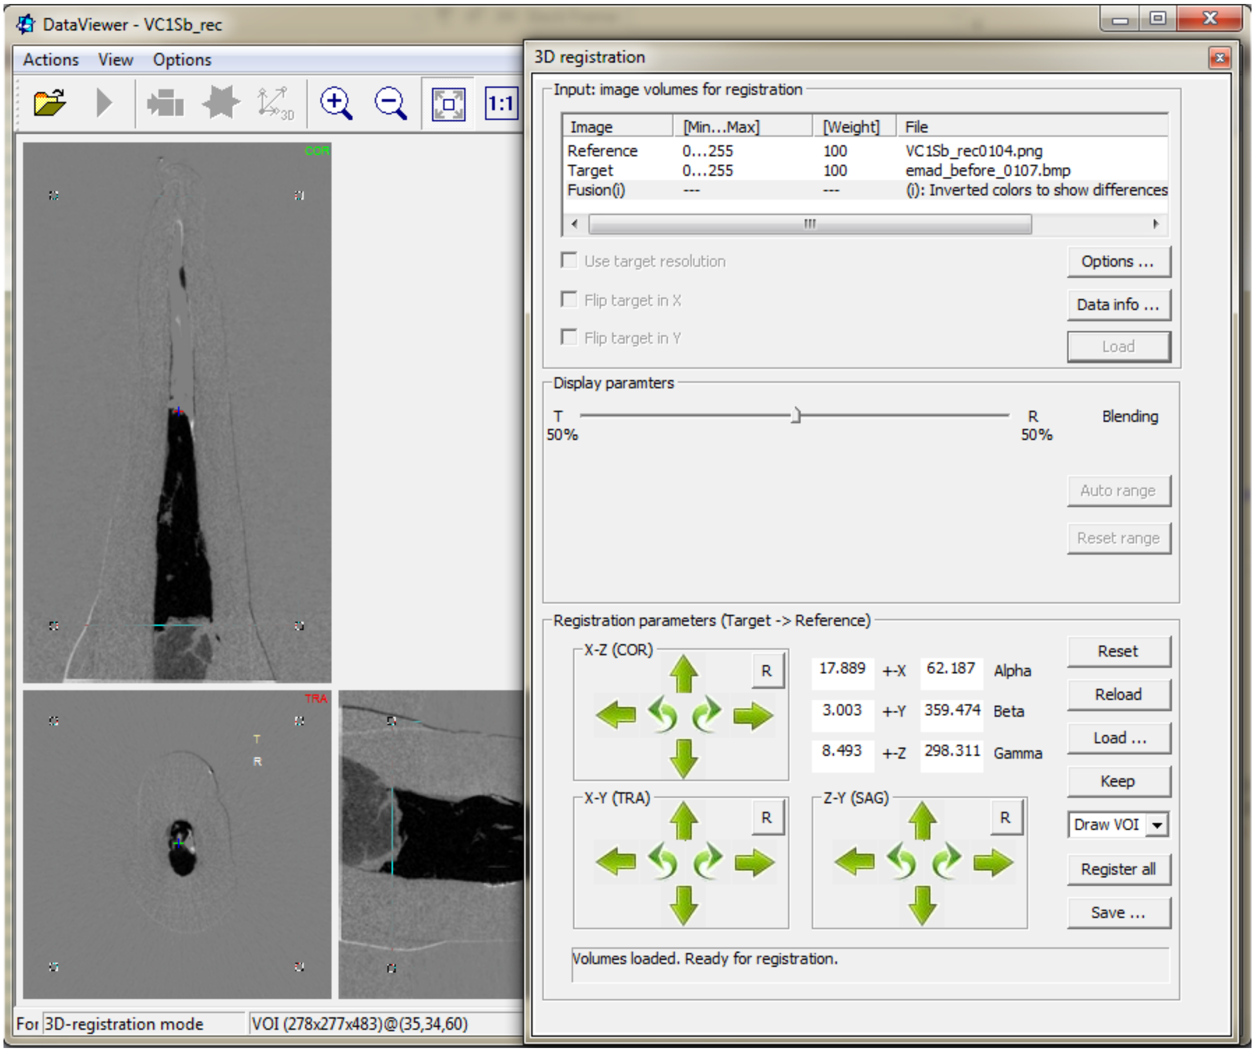

Fig. 3.2 Main window for 3D image registration.

To visually examine how good the registration is, one can alternate between target and reference images, or view the fusion images with various settings. Alternatively, you can also check by making profiles along any line (right-mouse click and drag).

The target image can be shifted and rotated manually using the arrow buttons, or keyboard shortcuts (see below for detailed key combinations). More conveniently, you can use "CTRL + left-mouse-button" to rotate or "CTRL+SHIFT+left-mouse-button" on any of the 3 views to shift and rotate to any position and orientation. Make sure the fusion image is displayed to see the effect immediately. There are 3 sets of arrow buttons. Each of them corresponds to a given view: the controls and the views are arranged in the same way topologically. You can alter the shift/rotation steps using button Options.

The registration can also be done automatically using buttons R and Register all after selecting VOI carefully. It is generally recommended to shift/rotate manually first so that the 2 sets of images are sufficiently close to each other, before using the automatic function. You can select the 3D registration method (pseudo-3D or true 3D) and alter the searching range using button Options.

On each set of the arrow windows, there is a button R (standing for register). This action would do automatic registration on the corresponding view only. The Register all button does registration on the 3 views alternatively (in a few loops): after registration on one view, all 3 views are retrieved again using the newly obtained transformation parameters. All these actions can be applied repeatedly, and on different location of the volume. A Volume Of Interest (VOI) is defined for automatic match: only the area within the VOI is used for calculating matching criteria (by default, the square difference). This can be crucial for registration: the intuition is to use an area with clear and similar features for registration. The VOI can be hidden if desired, or can be enabled for drawing. To move around to visualize the images at any location, use left-mouse click as in the normal viewing mode. *Attention should be paid here:* if you want to move the cross-hair position into the VOI, you need to use double-click instead of single-click!

To keep the transformation parameters, three actions are provided: Reset, Reload and Keep. The parameters can be kept in the original log file (if absent, one log file will be created) in the target dataset. Reload would simply read the parameters back and apply the transformation. Reset is to discard all shifts and rotations in the current session. Once the parameters are saved in the log file, they will be loaded back automatically when you open the same set of data back again for 3D registration.

Once the images are registered, a saving function with button Save... will allow you to save the transformed images. This option has been extended quite a bit since the beta release. Fig 3.3 shows the window when you click on the button. By default, the images are saved from the loaded volumes in computer memory (thus, often with resampling); the current data dynamic range (as shown on the histogram) are applied before saving (therefore all saved images will be 8-bit); and the images are saved at full volume (the common matrix as used for loading). Depends on what you want to do with the saved data, you may want to change the default settings. However, each choice has its limitations.

You may want to save a VOI instead of the full volume. *Attention is required here*: the VOI used for registration is now used for saving; and more often than not, you may want to alter the VOI before saving, as a VOI good for registration is very probably much smaller than what you want to keep! There are two advantages of using VOI here. Firstly, you may save disk space by selecting VOI properly. Secondly, you may want to keep all features: by shift and rotation it is possible that some of the features have gone out of sight. By placing the VOI such that this side of the VOI goes outside of the display box, the VOI will be automatically extended to avoid missing features.

Quite some users prefer to save data with original data range (image intensity). If this option is chosen, the data range as shown on the histogram will not be applied before saving. In addition, if the original target/reference image is 16-bit, the corresponding image will be saved in 16-bit as well (in TIFF format, no matter which output file format is chosen). In addition, the difference image will be also in 16-bit TIIF format, if both target and reference images are in 16-bit. Each pixel in the different image is then calculated as the following:

Diff = (target – reference) + 128 (8-bit)

Diff = (target – reference) + 32768 (16-bit)

Where 128 and 32768 are the median gray value for 8-bit ([0, 255]) and 16-bit ([0, 65535]). Note that difference value smaller than -128 (or -32768) or larger than 127 (or 65535) is truncated to keep values within range.

Saving images at original resolution has been required by several users. As the original volumes are loaded at reduced resolution, files on hard drive have to be read back again. Selecting original resolution for saving implies that the saving would be much slower: the original dataset is often larger than the memory can handle, so that the saving is done piece by piece. If VOI is not chosen, the common matrix dimensions will be used, not the original target image dimensions. The target image will be saved with its original pixel size. *Attention is required here*: if the original pixel size in target image and reference image are different, both saving reference option and save difference option are disabled. This is to avoid unnecessary discrepancies. To calculate the difference image, the saved target and reference images are read back first. Therefore, the 3 sets of datasets are saved one after another.

There is yet one more thing concerning difference images needs to be clarified. In the context of registration, the difference image is based on the grey values. By default, if the data range is selected to be "as displayed", the difference image is calculated using the middle-layer mentioned above, i.e., with the range given by the histograms applied. This gives an idea how good/bad the registration is. However, if the data range is selected to be "original", the range given by histograms are no longer applied. In addition, as mentioned above, if both target and reference images are in 16-bit, the difference image will be in 16-bit too. One particular question asked here is: can I still interpret the difference image in attenuation coefficient or Hounsfield Unit (HU)? The answer is "No, not exactly." As we mainly consider grey values in the context of registration, and the exact quantity of the gray values in different dataset are not necessarily the same. There is however one very narrow path which makes it possible: if both target and reference images have been reconstructed using the same dynamic range so that the gray values have the same meaning in both datasets; if original data range have been chosen during saving. If one wants to keep the HU unit, in addition to the above condition, the HU calibration should also be exactly the same!

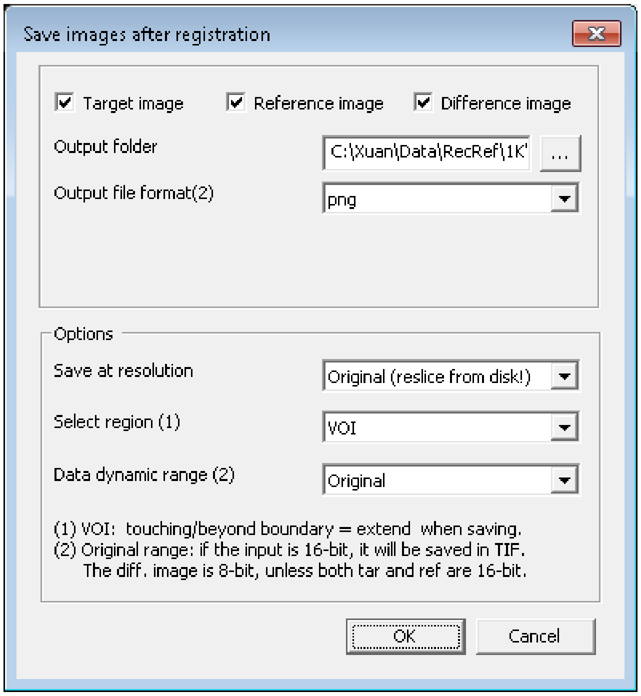

Fig. 3.3 Save after 3D registration.

Keyboard shortcuts:

(For the keyboard shortcuts to work, click at an empty place on the main window so that it is on focus to respond to the shortcut signals.)

Shift + ←: shift target image to the left (x-axis) by one step.
Shift + →: shift target image to the right (x-axis) by one step.
Shift + ↑: shift target image further (y-axis) by one step.
Shift + ↓: shift target image closer (y-axis) by one step.
Shift + **Page Up**: shift target image up (z-axis) by one step.
Shift + **Page Down**: shift target image down (z-axis) by one step.
Ctrl + **Shift** + ←: rotate x-y (tra) view to the left (around the z-axis) by one rotation step.
Ctrl + **Shift** + →: rotate x-y (tra) view to the right (around the z-axis) by one step.
Ctrl + **Shift** + ↑: rotate x-z (cor) view to the left (around the y-axis) by one rotation step.
Ctrl + **Shift** + →: rotate x-z (cor) view to the right (around the y-axis) by one step.
Ctrl + **Shift** + **Page Up**: rotate z-y (sag) view to the left (around the x-axis) by one rotation step.
Ctrl + **Shift** + **Page Down**: rotate z-y (sag) view to the right (around the x-axis) by one rotation step.
Ctrl + left-mouse-button: rotate to any orientation.
Ctrl + **Shift** + left-mouse-button: shift to any position at any direction.

4. Further development

New functions and features will be gradually added in the next releases. A brief list of these functions is given below.

- Alternative similarity/dissimilarity measures: e.g., co-variance, mutual information..., in addition to the sum of square difference.
- Alternative views for the pseudo-3D registration, e.g., MIP images in addition to the usual orthogonal views.
- Transformation: scaling factor in addition to translation and rotation.
- Pre-processing of the image views.
- Alternative visualization method for registered images, e.g., check-board viewing, difference image.

Suggestions and feedbacks are appreciated, via Kon_Applications@bruker.com.