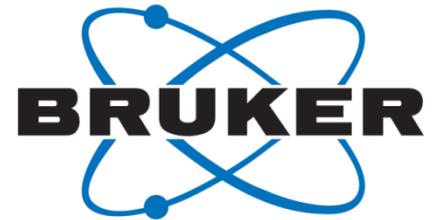


Bruker microCT



**An overview of NRecon:
reconstructing the best images
from your microCT scan**

Method note

MCT-62

1. Introduction

NRecon is the Bruker-microCT program that reconstructs your microCT scan projection images into cross-section images. You can reconstruct your scan cross-sections as soon as the scan is complete.

This document describes how to go through the steps of reconstruction in NRecon to get the best image results. NRecon is the program used for reconstruction in all Bruker-microCT scanners.

The images obtained by your scanner during the “scan” or acquisition are called “projection” images. After the scan, cross-section images are made by NRecon by reconstruction of the full set of projections, and these images are called “reconstructed cross-section” images (see figure 1).

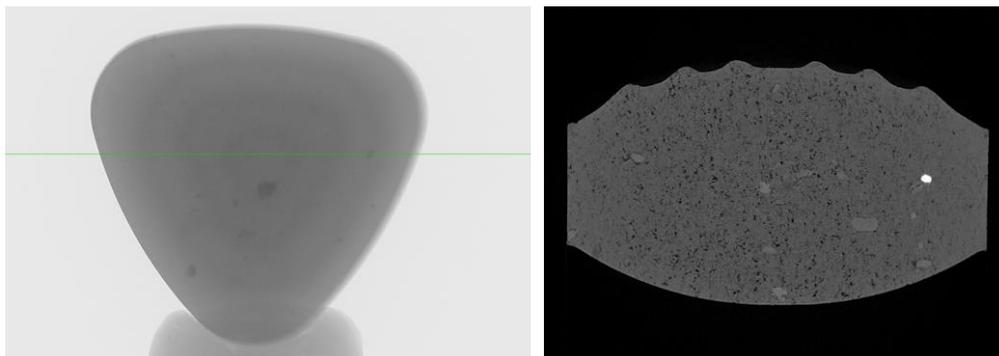


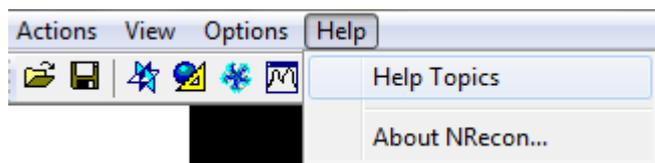
Figure 1. A scan projection image (left) and reconstructed cross-section (right).

It is important to understand the difference between projection and cross-section images, and the geometric relationship between these two types of image. The projection images are the primary, “raw data” of the microCT scan. The reconstructed cross-sections are secondary images, synthesised from the set of projection images. Each reconstructed cross-section image can be considered to correspond with a single horizontal row of pixels located at the same height in all the

projection images¹. Thus in figure 1, the cross-section image on the right corresponds to the horizontal green line shown in the projection image to the left.

The NRecon help file

Please note that NRecon has extensive help documentation. This is available under the “Help” menu and “Help topics”:



This help material is reproduced in a separate NRecon manual. Not everything in these help files will be reproduced in this method note, which is intended as a concise overview, or “getting started” document.

As this method note is intended as an introduction to NRecon, it will not cover in detail all advanced techniques that are available², but will help a user to learn quickly how to do effective reconstruction with correct parameters to get the best image results from microCT scans from all Bruker-microCT scanners.

¹ Although a cross-section image corresponds to just one horizontal row of pixels in all the projection images, the reconstruction of that cross-section usually requires more than that row of pixels only, depending on the position in the projection. Except at the optical axis near the center of the projection image, where only the one row of pixels is needed for “fan beam” type reconstruction, all other projection image pixel rows, above and below the optical axis (where there is a finite “cone angle”) require a number of pixel rows both above and below the reconstructed row, in order to interpolate a cross-section in the plane parallel to that of the optical axis reconstruction.

² Separate method notes exist for advanced functions such as the post-scan XY movement correction, which is implemented in NRecon.

The two programs – why is there NRecon plus a “Server” program?

The program NRecon provides the user interface for controlling and launching reconstruction. But NRecon itself does not perform the reconstruction calculations. These are done by a second “server” program (or reconstruction engine). Three server programs are available: NReconServer, GPUReconServer and InstaReconServer.

NReconServer is the standard server program which uses your computer’s CPU to do the reconstructions.

GPUReconServer uses your computer’s graphics card to perform reconstruction more quickly. The graphics card – at the time of writing – needs to be an NVidia™ card of the Quadro, GeForce GTX or Tesla ranges.

InstaReconServer™ is a third party product, supplied by the Illinois, USA based company InstaRecon Inc. It performs reconstruction using the CPU but with an innovative algorithm that speeds up reconstruction up to 100 times.

Thus, one reason for the separation of the NRecon from the server program is to allow different software solutions for reconstruction, aimed at increasing the speed of reconstruction.

Cluster reconstruction

A second reason for the separate server programs from NRecon, also connected with reconstruction speed, is to allow a “cluster” reconstruction. Cluster reconstruction is where one or more other PCs, running the server program and linked to the control PC by a fast network connection, carry out the reconstruction calculations together in parallel. Reconstruction is controlled by NRecon on the control PC, and the scan projection images are typically (but not necessarily) also located on the control PC.

A diagram of cluster reconstruction is shown below in figure 2. Cluster reconstruction can be done with one or more separate PCs network-linked to the control PC. Where both NRecon and the server program run on the same computer, and there is no other computer running the server program, this is called “local” reconstruction. Figure 2 shows three “scenarios” for local and cluster reconstruction.

Licensing of NRecon

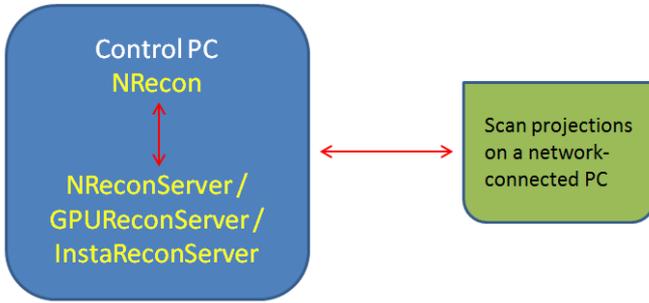
As explained in figure 2, there are “local” and “cluster” versions of NRecon and the server programs. The license requirement is as follows:

- NRecon (local) requires no license (figure 2, scenario 1)
- NRecon (cluster) can run in a 1-1 single linked PC cluster without a license (figure 2, scenario 2)
- NRecon (cluster) needs the cluster license in order to run in a multiple linked PC cluster (figure 2, scenario 3). A different license is required for running NReconServer and GPUReconServer.

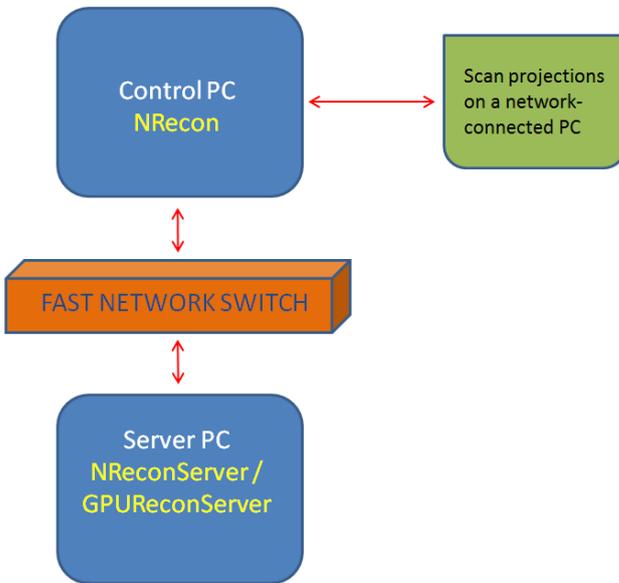
InstaRecon is no longer available in a cluster. Increases in computer power causes network speed to become a more limiting factor for speed gains from a cluster.

Figure 2 also shows that the projection images being reconstructed can be located either on the control PC, or on one of the server PCs (not a common situation but possible) or on a separate network connected PC outside of the reconstruction cluster. This gives you the maximum possible flexibility for reconstruction; you can run it from your office for instance without even visiting the scanner room.

Scenario 1: Local reconstruction



Scenario 2: Cluster, single connected PC ("1-1 cluster")



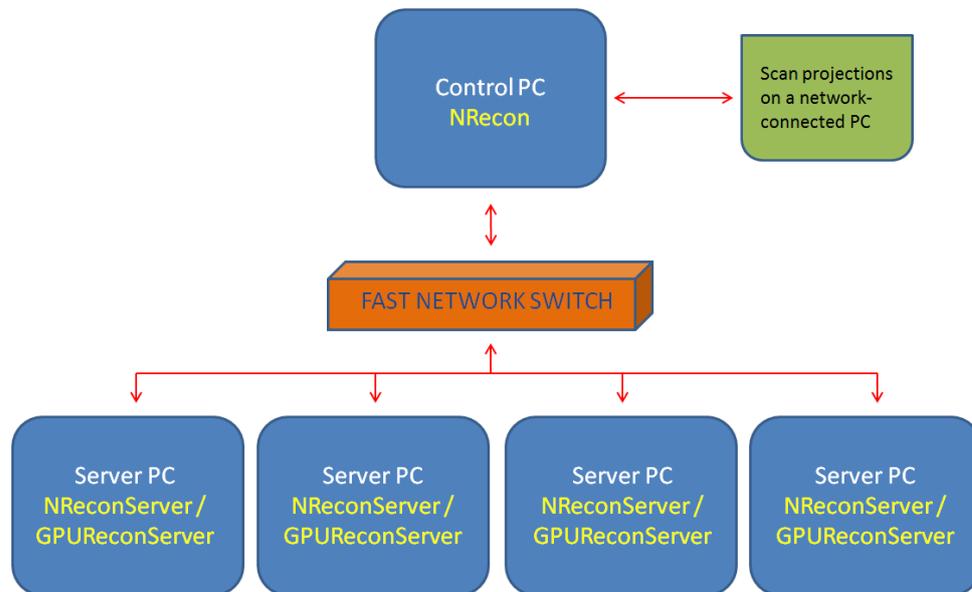
Scenario 3: Cluster, multiple connected PCs

Figure 2. The three local or cluster “scenarios” for using NRecon. The first is “local”, it can be run with the local versions of NRecon and any of the three server programs (NReconServer, GPUReconServer, InstaReconServer). The second is a cluster using a single offline PC connected by a fast network switch (“1-1” cluster). The third is a cluster using more than one offline PC, up to a maximum of 5-6 (typically 4). InstaReconServer runs only with the local version of NRecon. It is not possible to run InstaReconServer on a cluster, either a 1-1 or multi-PC cluster. Note also that in all three scenarios, it is possible to reconstruct projection image datasets that are located either on the control PC or remotely in another, network-connected PC (green boxes), as well as on any of the server PCs.

GPU technology

Graphics card or GPU (graphics processing unit) technology has advanced rapidly in recent years making possible a big acceleration of reconstruction at quite low cost. In fact such is the gain in reconstruction speed from even a single powerful GPU, run either locally or in a single offline PC, that currently the scenario 3 of multiple offline PCs is being used less often than in former years. Or if it is used, it is done in a different way: scenario 3 can be implemented in the form of a GPU server unit containing up to eight single or multiple NVidia Tesla GPUs. This represents the fastest possible solution for microCT reconstruction in SkyScan systems, suitable for

the very large data formats produced by high-performance scanners such as the SkyScan1272.

To reiterate, the graphics card – at the time of writing – needs to be an NVidia™ card of the Quadro, GeForce GTX or Tesla ranges, on which the NVidia “CUDA” programming language is available.

InstaRecon

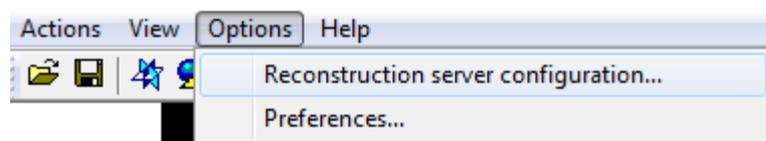
InstaRecon™ is a program provided by InstaRecon Inc., Illinois, USA, which uses algorithmic acceleration to increase the speed of reconstruction from 20-100 times depending on scan resolution and pixel format. InstaRecon has been optimised to run on all Bruker-microCT systems as the InstaReconServer. More information on this innovative and unique reconstruction solution can be found at:

<http://instarecon.com/cbr-product/>

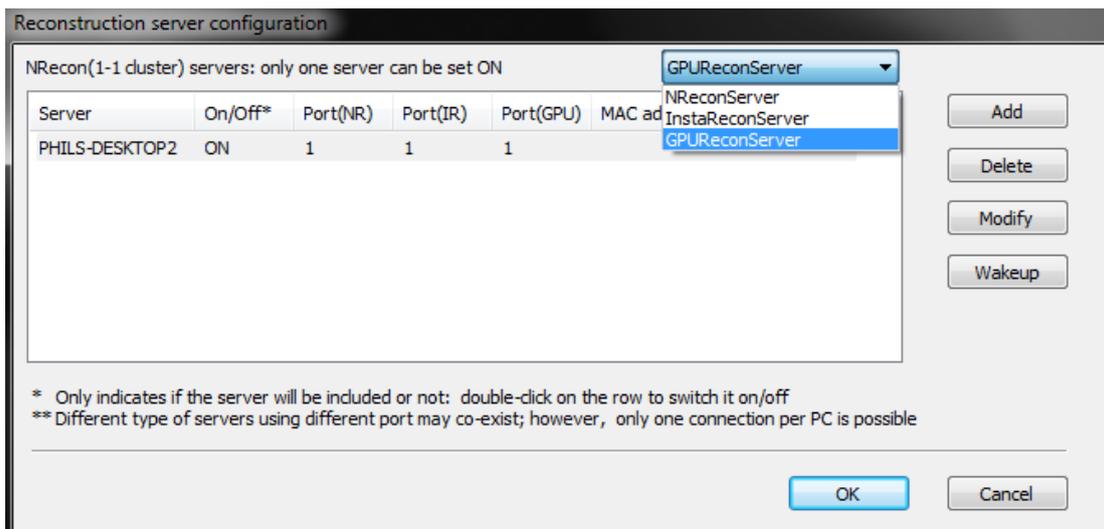
A license for this software can be obtained from your Bruker-microCT distributor.

Selecting your reconstruction option in NRecon

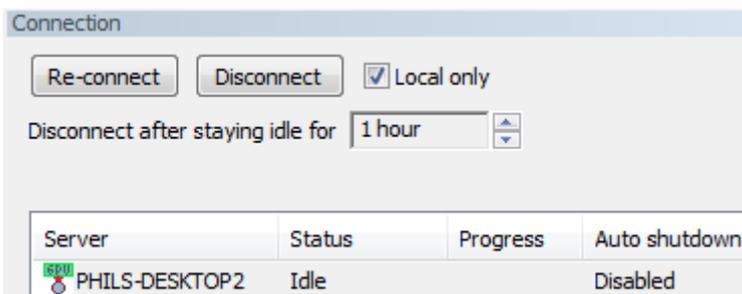
In the options menu, choose the first item, “Reconstruction server configuration”:



This window will open:



The drop menu at the top right allows you to choose between NReconServer, InstaReconServer and GPUReconServer. Note that InstaReconServer will only be available if you have a license for it. Select a reconstruction option and click on OK³. If you are running the local version of NRecon, then the selected server program will open automatically when you click on OK (providing the server program executable file is in the same folder as the NRecon exe file). However if you are running the cluster version of NRecon, then you should open the server program first, before selecting it in the NRecon “Reconstruction server



³ At the right of this server configuration window, the buttons “add” and “remove” allow you to add and remove networked computers to join a reconstruction cluster. This will have been done by the service engineer on installation of your scanner. The button “modify” gives you options to assign different COM ports e.g. 1, 2 and 3 to the three server programs to allow them to be opened simultaneously (if they share a port e.g. 1 then only one can be opened at one time). However since NRecon can only use one server program at a time, assigning different ports is usually unnecessary. The “wakeup” button can be used to wake up a connected server which is in sleep mode.

configuration” window and clicking “OK”.

When connection to a server program is successfully established, then in the “Connection window” at the lower left of the NRecon program window, you will see the symbol of the server program under the heading “Server”, followed by “Idle” under the “status” heading (see the above image for the case that GPUReconServer is selected).

Note that if NRecon and the server program are both running but not

Server	Status	Progress
PHILS-DESKTOP2	Offline	

connected, you will see the status of the server shown as “Offline”. In this case you should click on the “Re-connect” button at the top left of the Connection window and the connection will be restored (the status becomes “Idle”).

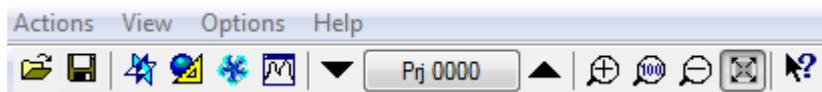
The layout of NRecon

The layout of the NRecon interface is shown in in figure 3. Four parts of the interface are shown in coloured boxes. These parts and their role in reconstruction are briefly described below. A fuller description of each part of the interface will follow in the reconstruction steps section below.

The main image

Occupying the main part of the interface toward the top left is the main image. When NRecon is first opened, the image displayed is one of the scan projection images. After a cross-section is reconstructed, then this cross-section can be displayed in the same main image position by selecting the “output” reconstruction step tab (see below). Either a projection or a cross-section image is shown in the main window (see figure 1), depending on which step tab is selected.

The top menus and buttons



The Actions menu contains a number of advanced options such as a movement correction method and resaving with optional down-binning of the projections. It also has options to load reconstruction parameters from files such as protocol or log files, to help with standardising reconstruction. This menu also contains links to start other Bruker-microCT programs for viewing and analysing reconstructed datasets.

The View menu allows you to show or hide different windows in the interface. It also contains zoom functions.

The Options menu has two items. One is a configuration window where you choose which of the three server options you will use (NReconServer, GPUReconServer or InstaReconServer). The other is a preferences window with many reconstruction options.

The Help menu links you to extensive help documentation, as well as telling you the version of NRecon.

The Button row below the menu headings provide options also covered in the menu items, such as links to start other Bruker-microCT programs for viewing and analysing reconstructed datasets. There is also a profile line function, a bar to step through projections, the zoom controls, and a help button.

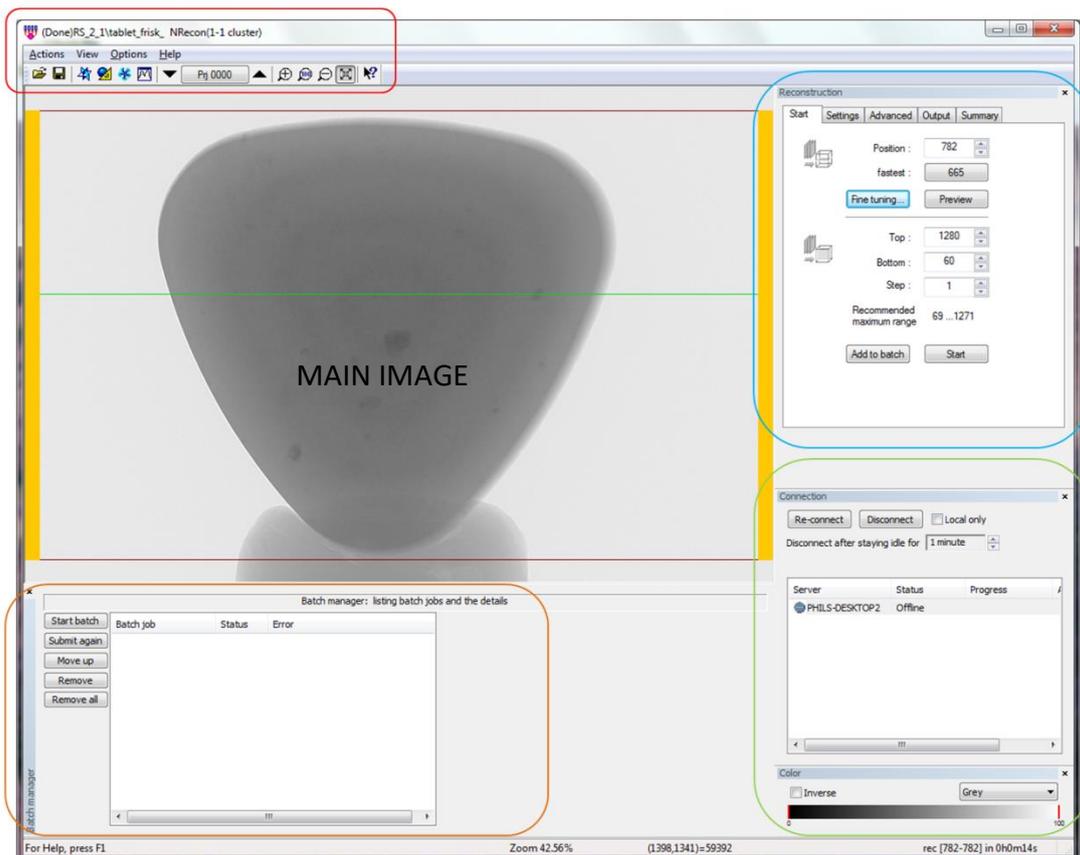


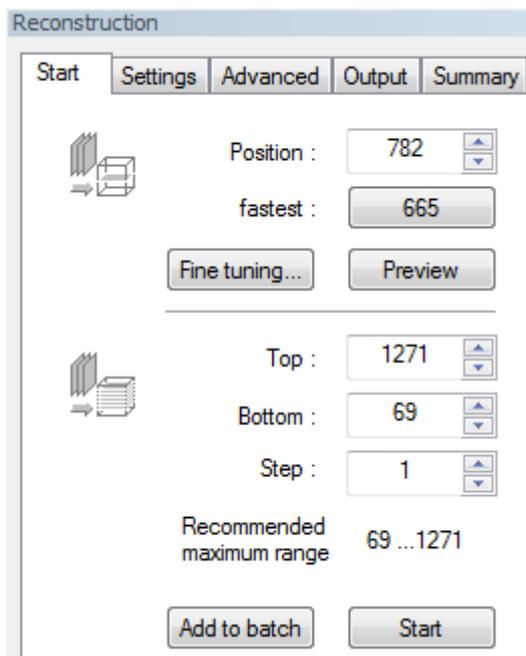
Figure 3. The NRecon user interface. In addition to the main image, four sections of the interface are shown in coloured boxes. These are: the top menus and buttons (red box), the reconstruction step tabs (blue box), the connection window and colour palette (green box) and the batch manager (orange box).

The reconstruction step tabs

The top right part under the heading “Reconstruction” contains the five step tabs which you step through to set up and run a reconstruction:

Start: here you set the level for a one-slice reconstruction preview and the range for the full reconstruction.

Settings: here you select image



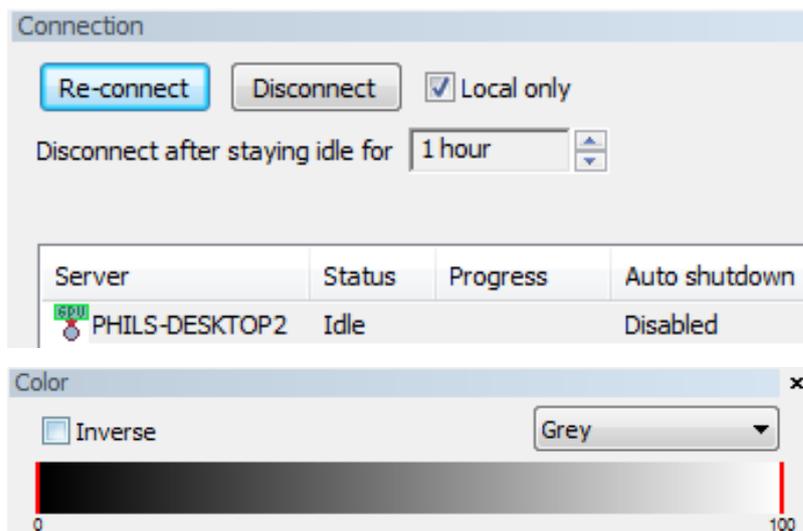
reconstruction parameters such as smoothing, ring artefact and beam hardening correction.

Advanced: some advanced options.

Output: here you set the density contrast range (a kind of brightness and contrast) and also choose options such as a restricted region of interest (ROI), a folder location and image file type for reconstructed results.

Summary: this tab does not always need to be visited for reconstruction but displays summary scan data and allows saving a protocol file to help standardise reconstructions and load previously used parameters.

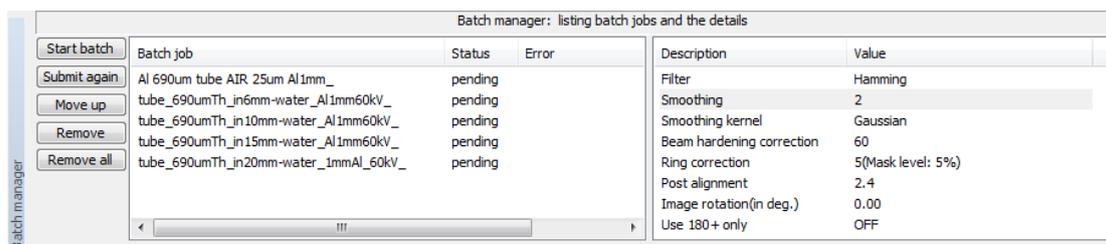
The connection window and colour palette



The Connection window shows if a server program such as GPUReconServer is actively connected to NRecon. It also shows if the connected server is idle or busy with reconstruction (and if so, the progress).

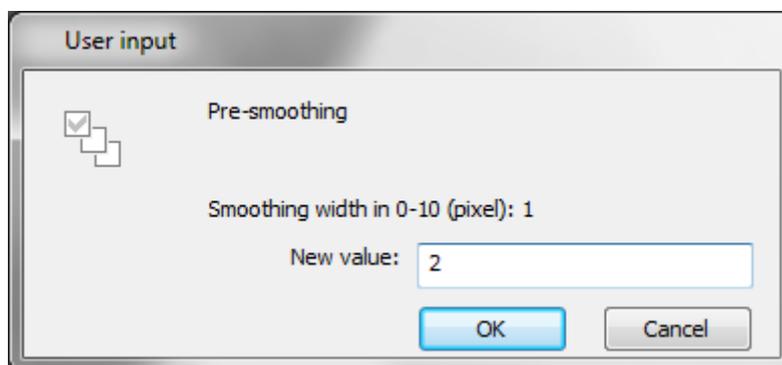
The **Color palette** allows you to invert the displayed image or apply a false colour palette, with a dual slider bar for range adjustment.

The batch manager



The batch manager displays reconstruction jobs that are queued for automatic batch reconstruction. Each job is a scan dataset to be reconstructed⁴. The left window shows the list of jobs waiting, and the right window shows all reconstruction parameters for each job.

These parameters can also be edited from the batch window by double-clicking on the parameter line in the right pane of the batch manager. The buttons



on the left allow removal of jobs, change of job order and start of the reconstruction batch run.

⁴ You can if needed make several reconstructions from a single scan, for instance if it contains different objects or parts. In that case each separate reconstruction should be placed in a separate sub-folder ("output" tab) to prevent over-writing.

If a reconstruction job is cancelled, and you then wish to run it again, select the job and click the button “Submit again”; the job will return to the “pending” status.

Stepping through the reconstruction process

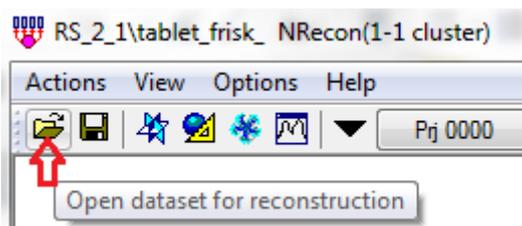
As a worked example, the reconstruction will be shown of a “Frisk” mint tablet (figure 4).



Figure 4. A “Frisk™” mint tablet is used as an example scan reconstruction.

Load the projection dataset

At the top left of the NRecon window click on the button at the far left, the yellow open file icon:



In the windows open file dialog that opens, navigate to the tiff format projection images from your scan. Select any one of the scan projection tif images which is followed by the sequence number (with either 4 or 8 digits). Do not click on a file other than a projection image. To open the whole dataset, either double click on one of the projections as shown, or single-click to highlight the projection image then click on the “Open” button at the bottom right of the “Open dataset” window:

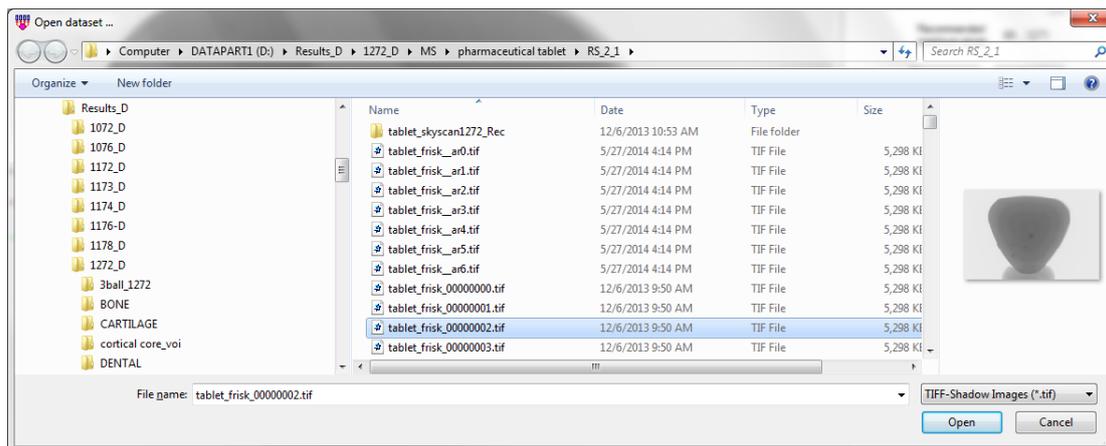


Figure 5. The open dataset dialog box.

Wait while the projection dataset is processed and loaded, after which the first projection image (numbered zero) is displayed in the main image window (figure 6 below).

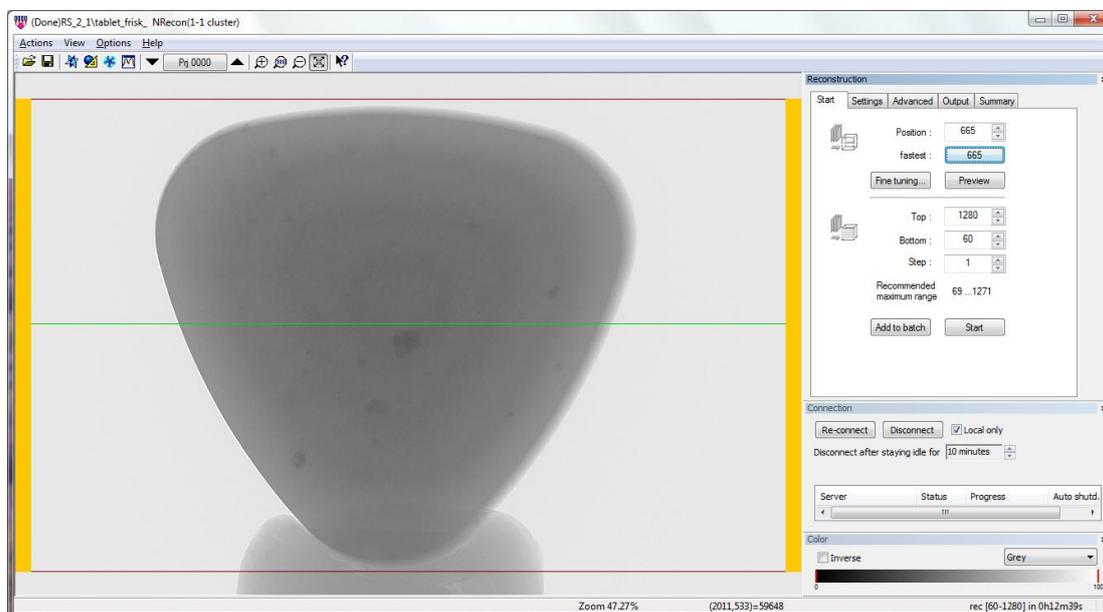


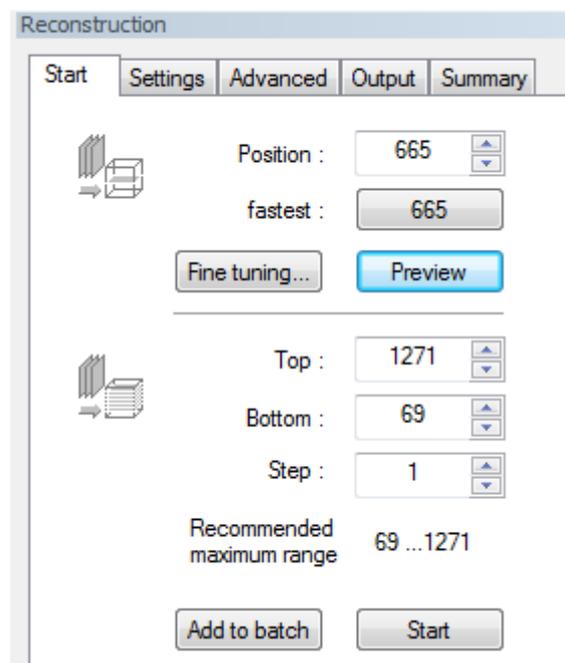
Figure 6. The display of the projection in the Start tab after opening the dataset.

Please note that on the displayed projection image are one green and two red horizontal lines. The green line shows the default level for a single-cross-section

“preview” reconstruction. The red lines show the default range for full dataset reconstruction (see below).

The start tab and preview

With the dataset open, look at the window on the top right under the heading “Reconstruction”. There are five tabs: Start, Settings, Advanced, Output and Summary. At first the tab which will automatically be opened is the first tab, “Start”⁵.



Connected to the Start tab display are the red and green lines shown on the projection image.

The green line is the preview line. By default this is set to the “fastest” midline cross-section, but can be moved to any level. The current level set for preview is shown at the top of the Start tab after “Position”.

⁵ Unless in preferences you choose the “show alignment image first” option in which case, the Settings tab containing this image will open automatically, not the start tab.

The two red lines show the range between which the full dataset reconstruction will be performed. By default these red lines are set at the widest possible reconstruction range. This means that the small part above the top red line and below the bottom one will not be reconstructed⁶. This maximum range is shown at the bottom of the Start tab as the “Recommended maximum range”.

First step: click “preview” to see the scan result (unadjusted)

In the start tab near the top right there is the “preview” button. Preview means to make a reconstruction of a single cross-section image.

Just above the preview button is a button with a cross-section number – here 665 – and to the left of it the label “fastest:”. This is the midline cross-section for which reconstruction is fastest – it gets slower above and below the midline⁷. That is why

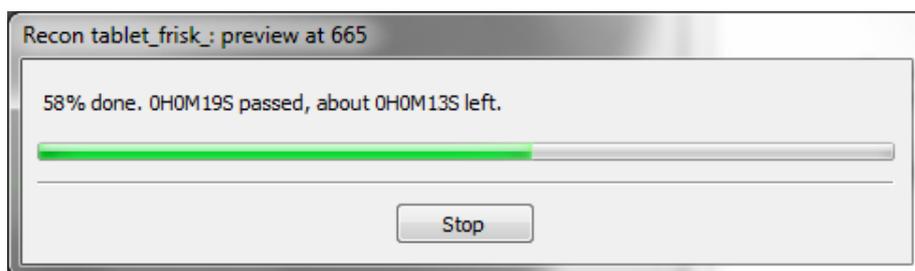
⁶ This is connected to the Feldkamp reconstruction method which at non-zero cone angles always requires some cross-sections above and below the reconstructed level for interpolation of a parallel cross-section from an inclined source-camera geometry. Thus the very top camera pixel row for instance could not be reconstructed since there are no slices above it for interpolation.

⁷ This means that the cross-section number 665 is the “optical axis” of the scanner. This is the cross-section representing the horizontal row of camera pixels which is the closest to the x-ray source emission spot – the only line at which x-rays from the source strike the camera at exactly a 90 degree angle. Above and below this line, the x-ray path relative to the camera plane is inclined upward or downward. This angle of inclination is called the “cone angle”. For all cross-sections with a non-zero cone angle, the Feldkamp reconstruction algorithm always interpolates over several horizontal rows of camera pixels to reconstruct a single parallel cross-section. As cone angle increases, so does the number of pixel rows that have to be interpolated. However, uniquely at the optical axis, no interpolation is needed and reconstruction uses that single pixel row only. That’s why it’s the fastest.

the default position of the green preview line corresponds to the fastest midline – here you get the fastest preview.

You can however obtain a preview at any height in the image, by either drag-and-drop of the green line up or down with the mouse left button, or by just a left mouse double-click anywhere in the image.

Once you have pressed “preview”, a progress bar will run...



After which you will see the reconstructed image. Looking at the tabs under the “Reconstruction” heading, you will notice that NRecon has automatically opened the “Output” tab, which displays the reconstructed cross-section at the selected preview level (figure 7, below).

Thus NRecon has jumped from the first tab (“Start”) to the fourth (“Output”) to display the result of the preview reconstruction.

However it may be necessary to step back to the second tab, “Settings”, to adjust reconstruction parameters and then repeat the preview one or more times.

The “start” and “add to batch” buttons

Below the settings for range top and bottom and the recommended range, are the “start” and “add to batch” buttons.



These are the buttons to press at the end, once you are confident that all settings for the reconstruction of the dataset are correctly set. “Add to batch” adds the reconstruction to the batch manager.

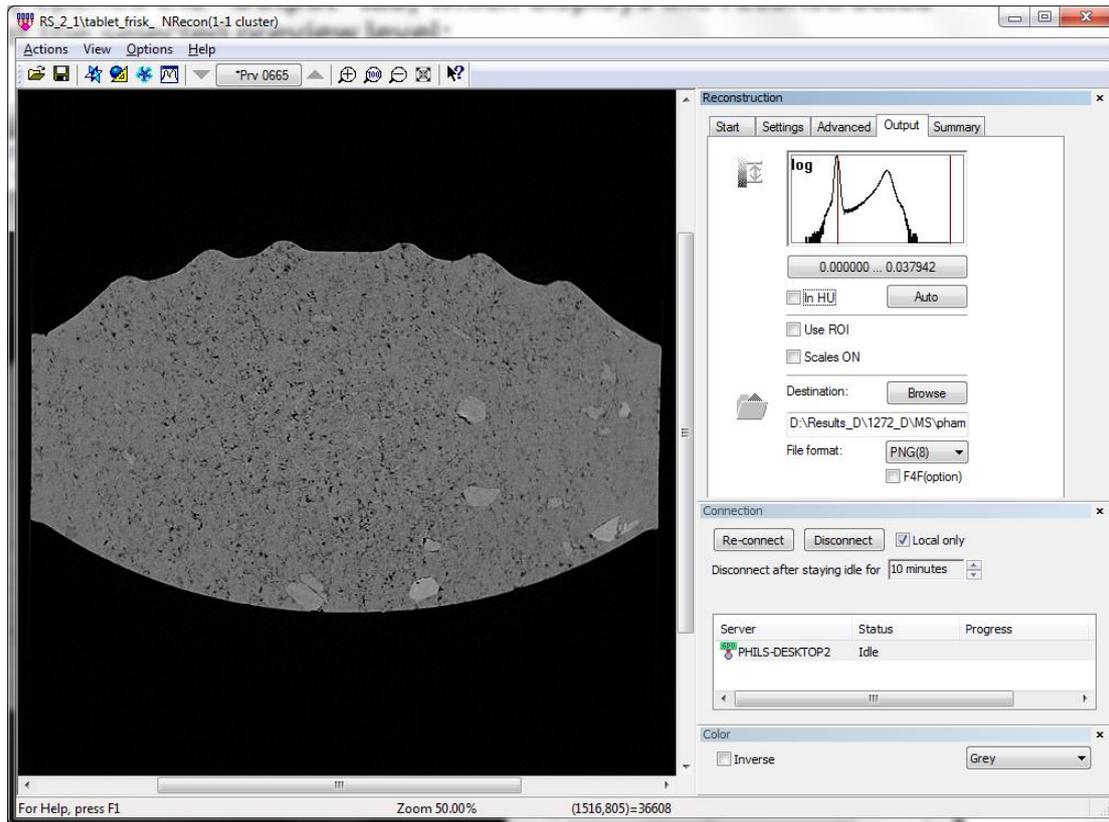
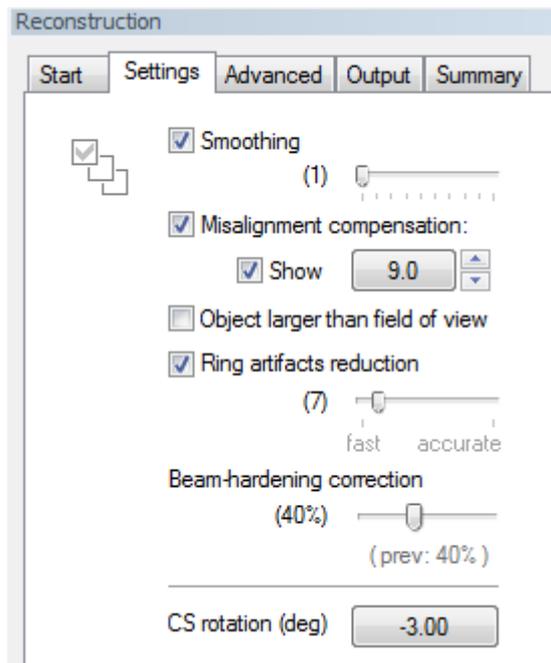


Figure 7. The display of the reconstructed cross-section that opens automatically in the Output tab after doing a preview

The Settings tab: choosing parameter values



The Settings tab is where most of your choices of numerical values of reconstruction parameters have to be made. The 6 parameters to set are:

- Smoothing
- Misalignment compensation
- Object larger than field of view
- Ring artefact reduction
- Beam hardening compensation
- CS rotation

Smoothing

Smoothing or “blurring” is sometimes needed to reduce noise in the reconstructed image. Please see the images in figure 8 below. You can judge the appropriate amount of smoothing visually⁸. Smoothing is set as a value on a scale from 0-10. (Note – in the advanced tab you can choose the smoothing algorithm, between a box kernel and Gaussian smoothing.)

⁸ Note that if the noisiness of the image is severe, such that a smoothing value of more than 5 is needed, this indicates that the scan acquired too little image data. Four factors affect the amount of image data: (a) the setup of the scan mode should be correct, such that the raw image (empty air, no flat field) has a mean intensity of around 60%; (b) a smaller rotation step will increase the number of projections, decreasing noise in the reconstruction; (c) increasing frame averaging will also increase the volume of image data, reducing noise; (d) if smooth images are hard to get from dense objects, then don't scan at maximum resolution but reduce the camera resolution level (i.e. increase the camera binning to 2x2 or 4x4). Then higher frame averaging values become practical.

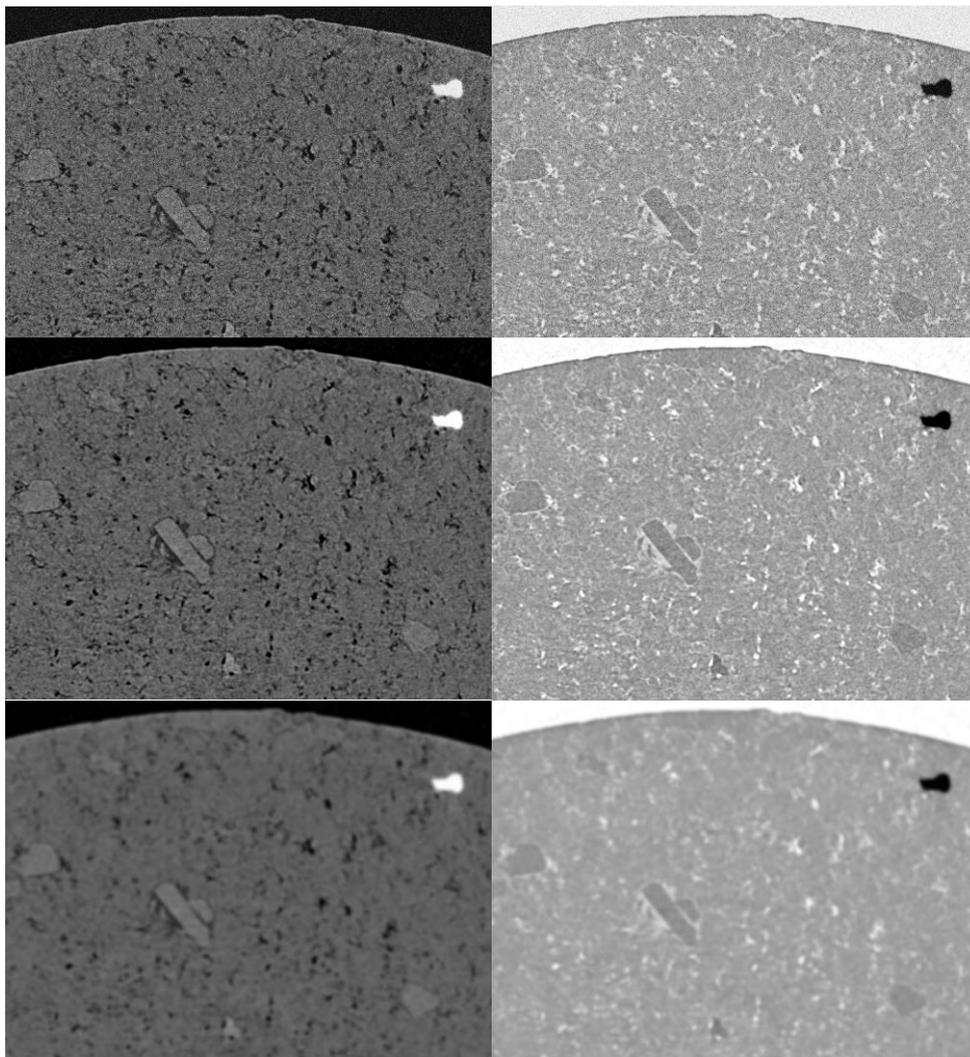


Figure 8. Top – no smoothing, too noisy; middle – appropriate smoothing, improved image; bottom – too much smoothing, loss of image detail.

Misalignment compensation

Alignment is a complex issue in tomography⁹. Every scan requires a test of alignment, and in most cases a misalignment or “post-alignment” compensation

⁹ Consider the axis of rotation, the virtual line around which the object rotates in 3D space during the scan. If this line could be projected by the x-rays onto the camera it would make a vertical line image. In a perfectly aligned scan this projected line would correspond exactly to the midline of the camera. In this case the misalignment is zero. A positive or negative misalignment value means that the

needs to be added. The value of misalignment refers to a number of pixels. Failure to make this compensation, or applying the wrong value, leads to artefacts shown in figure 9.

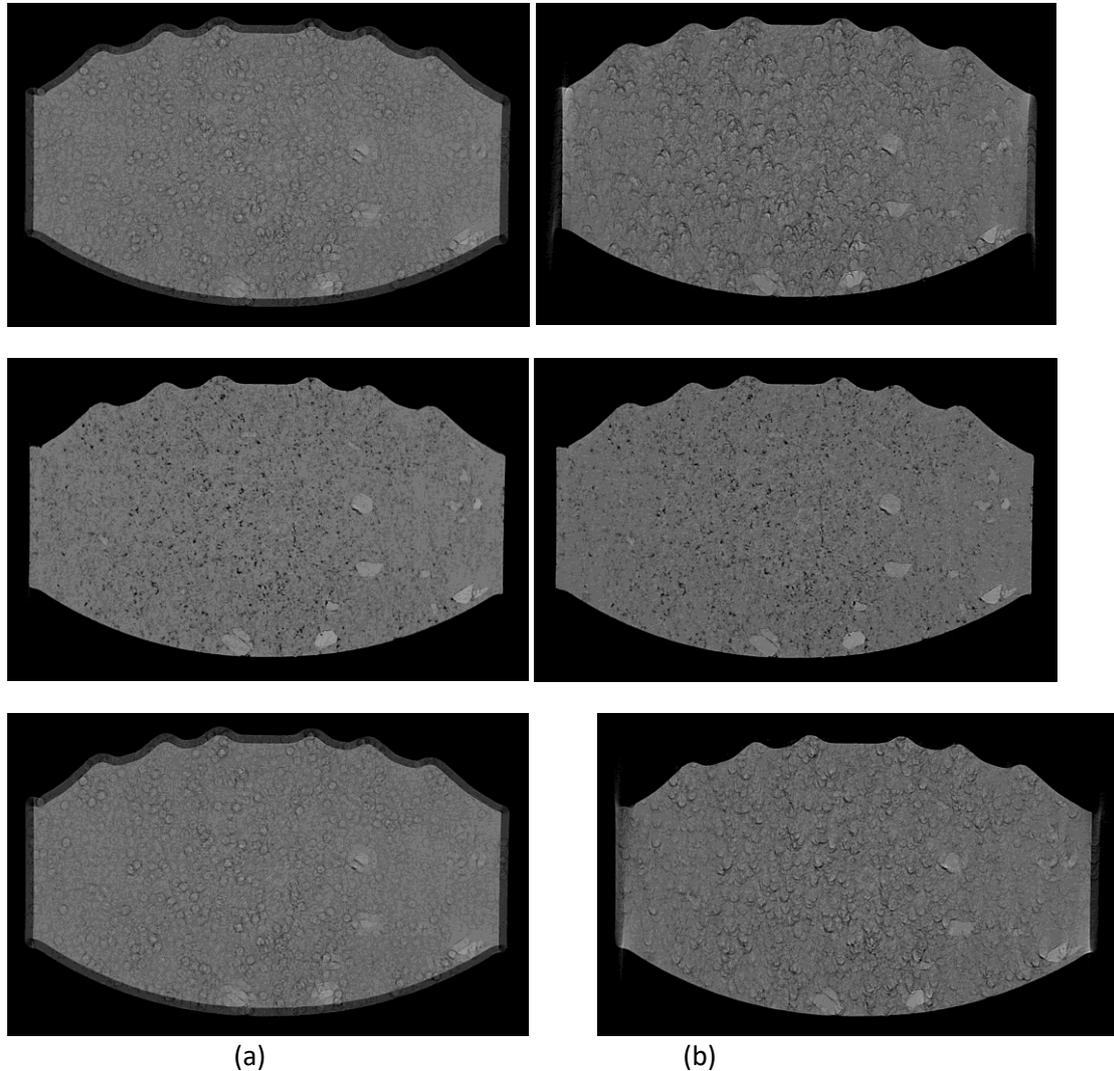
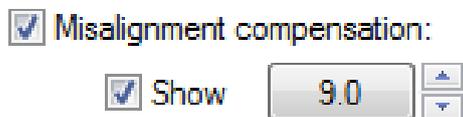


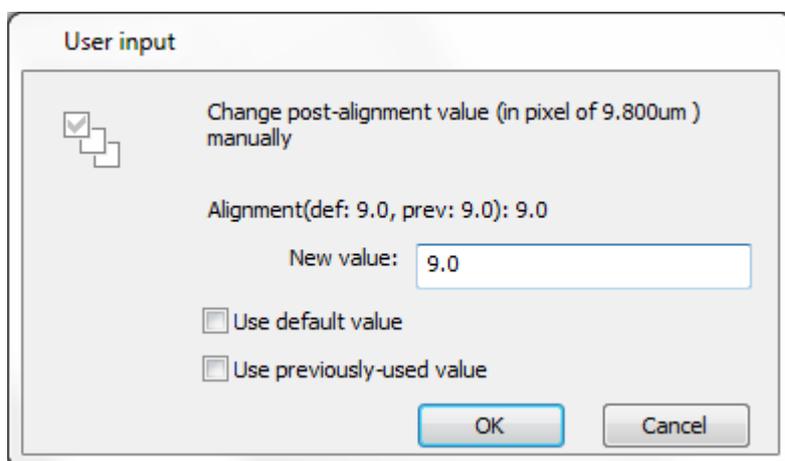
Figure 9. Reconstructed cross-sections with a misalignment compensation value that is (top row) 10 pixels too low, (middle row) correct, (bottom row) 10 pixels too high. Note the artefacts of 360 and 180 degree scans are quite different. The images in column (a) are from a full 360 degree scan, where the artefact of misalignment is an image doubling. The images from column (b) are from a half 180 degree scan, where the artefact of misalignment is curved streaks at the ends of objects facing downward or upward.

projected rotation axis is to the right or left of the camera midline by the given number of pixels. Note – you should run the hardware alignment test periodically on your scanner to keep the misalignment values low.

Note that NRecon automatically calculates the misalignment compensation. Thus when you open the Settings window you will already see a value given for misalignment:



This can be adjusted by the up and down arrows. Alternatively, if you click on the bar with the number – here 9.0 – another window will open allowing you to enter a chosen number by text entry:

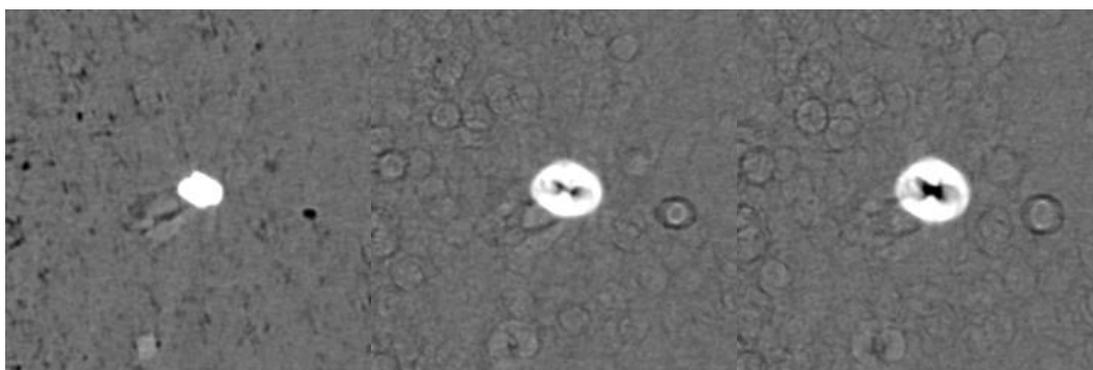


The value referred to as the “default” value is the one calculated by NRecon. The colored bar that you see in the image pane at the Settings page is the result of the image co-registration within a central band of the projection image, which NRecon uses to calculate the misalignment:

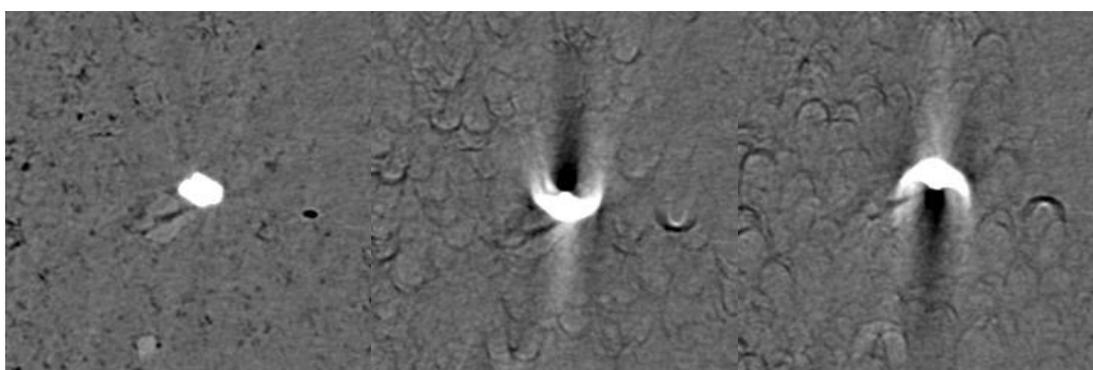


NRecon does not obtain the correct misalignment for 100% of scans. Therefore it is important to recognise the appearance of misalignment, so that where necessary you can adjust misalignment yourself. (for this reason you should always perform at least one preview reconstruction for all scans even where many experimental scans are being reconstructed with the same parameters, for comparison. This represents an important quality checking step. You should not “blindly” accept the results of reconstruction without viewing at least one preview cross-section.)

As shown in figure 9, the appearance of misalignment is quite different in 360 full scans and 180 degree scans half-scans. Misalignment in 360 degree scan gives an image doubling, blurring artefact, while from a 180 degree half-scan you get instead “u” or “n” shapes with upward or downward tails at the sides of features such as pores or particles:

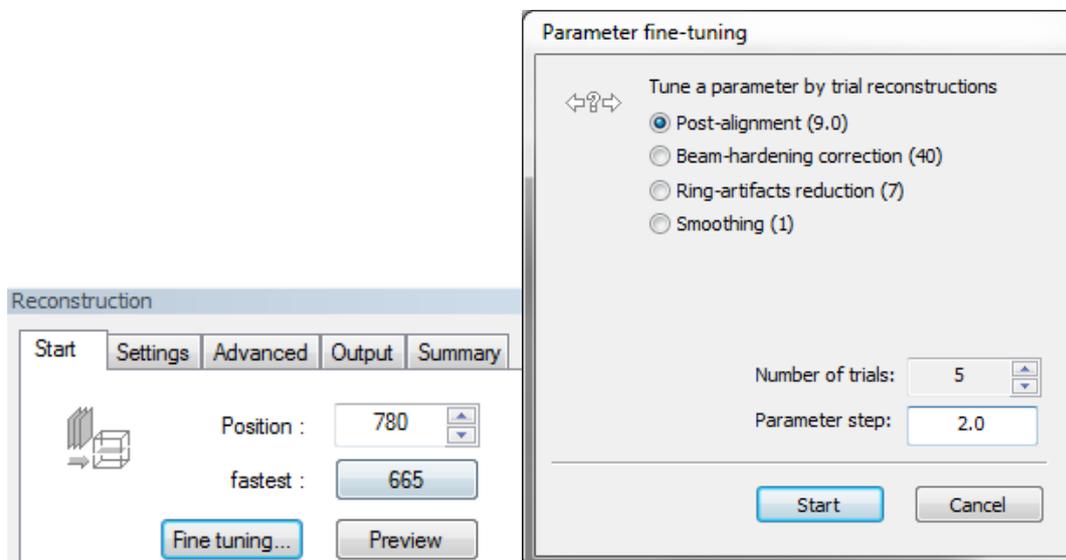


360 degree scan: correct (left) and misaligned (mid and right)



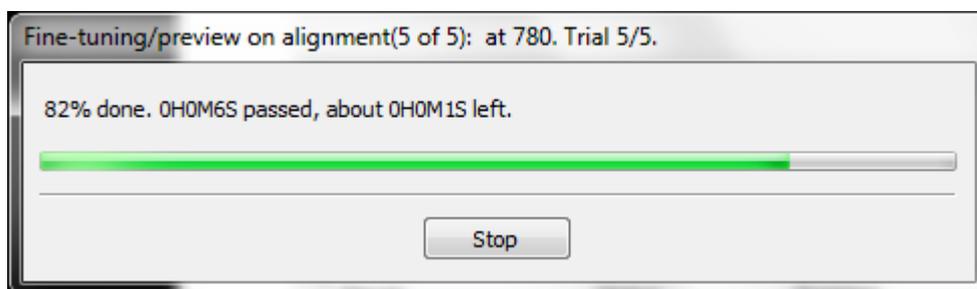
180 degree scan: correct (left) and misaligned (mid and right)

Finally – with misalignment as with other parameters at the Settings page, you can speed up the testing of different values by using the “Fine tuning” button which you will find at the Start tab, to the left of the “Preview” button:

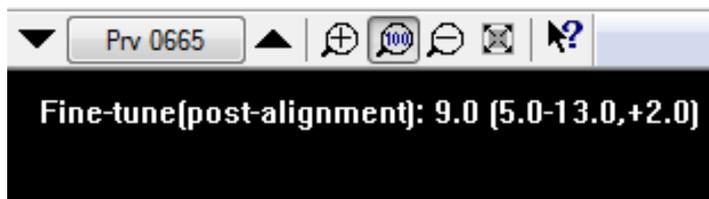


A selected “number of trials” will automatically preview images with values ranging below and above your initially set value – here a misalignment of 9.0. So for example, with 5 trials and a parameter step of 2, previews will be done at misalignment values of 5, 7, 9, 11 and 13 pixels.

Once you click on “Start” the trials will begin automatically:



Once they are complete, they are viewed at the Output page, where you can scroll through the previews with the different parameter values using the up and down arrows at the top of the displayed image:



Please note also: the misalignment compensation is a parameter that should be adjusted individually for every scan. Thus it is an exception to the rule that, when you are comparing microCT scans for experimental studies, all reconstruction parameters should be the same. This rule is valid for most parameters, but not misalignment.

Object larger than field of view

For an ideal tomography scan, the whole object should rotate within the left-right boundaries of the camera field of view (FOV). This means that a few pixels at both the left and right sides of the image should be occupied by air and not attenuated by any object or sample holder, at all levels of the image and for all projections. If this is the case then the object (actually its x-ray projection) is smaller than the FOV.

However if for part or all of the scan, any material (object or holder) moves beyond the left or right end of the FOV, then this is called “truncation” of the scan. It means that some attenuation information is lost from the reconstruction. It also means that the object is larger than the FOV. (Note that “object” includes any holder or surrounding material.)

In the case that some material crosses the right or left side during the scan, then you must check the option “Object larger than the field of view”.

Figure 10 shows examples of scans that are truncated or not truncated, i.e. where the object is either smaller or larger than the field of view.

Object larger than field of view

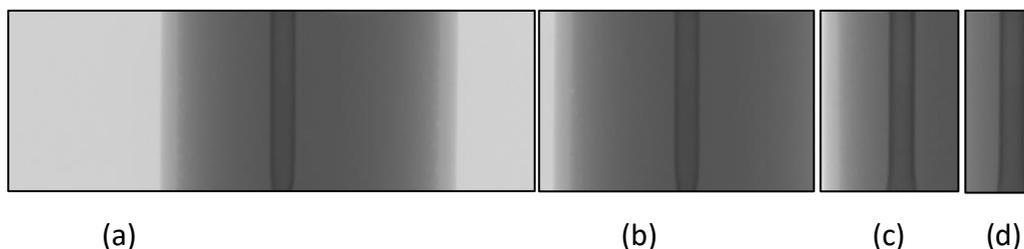


Figure 10. In scan (a) the whole object is contained in the field of view (left-right) so that there is empty space both sides. In all other scans (b-d) the object and/or holder are truncated, so that absorbing material rotates outside the FOV for part or all of the scan. In (b-d) the object is “larger than the FOV”.

It is important that if any part of the object, holder or mounting material crosses outside the left or right sides of the image at any time during the scan, then the “Object larger than FOV” option must be checked. Omitting this option when a scan is truncated can lead to serious errors of reconstructed intensity.

Ring artefact reduction

Ring artefacts are a ubiquitous feature of tomographic reconstruction, resulting from non-uniformity of the background image taken by the x-ray camera¹⁰. Note that

¹⁰ An object in the camera image which remains in the same place in all projection images of the rotating sample is interpreted as a ring by the tomography reconstruction. Examples of “objects” that can stay in the same image position include artefacts of small variation between camera pixels in their sensitivity to x-rays, or occasionally small particles in the beam path, such as on a filter or source collimator. These can be removed by physical cleaning of such locations (please ask your service engineer about this).

correct set-up of the scanner, such as having the correct camera exposure time, and choosing an appropriate filter and x-ray voltage (not too little or too much x-ray transmission through your scanned object) will minimise ring artefacts. It is also important to regularly update the flat field correction for the scan mode that you use – this will greatly minimise ring artefacts.

However even with the right scan settings, and an up-to-date flat field correction, some ring artefacts will often appear in the image.

Software ring artefact reduction is available in Ring artifacts reduction NRecon, with varying strength:



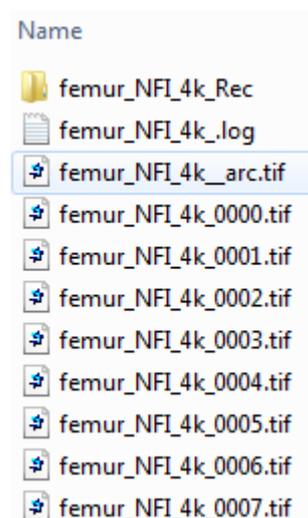
The number in the slider refers to the pixel diameter the operation. The maximum number available can be varied from 20-100, in the NRecon preferences:



Figure 11 below illustrates how to choose the best value for ring reduction by showing the results of increasing the ring reduction value from zero to higher values. Please note – in this cross-section to the right there is a high density particle. One needs to pay attention to higher density particles when setting ring reduction since a value of ring reduction that is too high will cause a new, “secondary” ring artefact around the particle.

In this example, when ring reduction (“RR”) is zero there are many rings visible. These are reduced but still visible at RR values of 2 and 4 pixels. An RR value of 7 removes almost all visible ring artefacts.

At RR values of 12 and 20, you can see around the high density (white) particle to the right of the image, a new

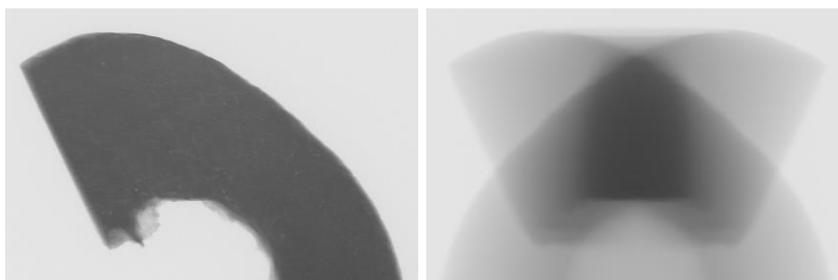


ring structure appearing. This is a secondary ring artefact, which is caused by the ring artefact reduction process itself, and results from using excessively RR high values.

The arc file: NRecon uses an image with the name “arc” to perform the ring artefact reduction. This is called the “arc file” and is saved in the same tif format as the scan projections:

The arc file – if you look at it – is a blurry image since it is simply the average of all the projection images. NRecon uses this averaged arc file to identify any invariant features that can lead to ring artefacts, and removes them from the reconstructed cross-section images¹¹.

Examples are shown below of a projection image and the arc file from the scan of a sample of femur bone:



Projection image (bone)

Corresponding “arc file”

¹¹ Note that there is an advanced ring reduction method, for problematic datasets, that involves making multiple arc files. By default a single arc file is made by averaging all projections over 180 or 360 degrees. However by adding a line to the log file it is possible to get NRecon to make “partial” arc files which average over multiple partial slices of the scan rotation. Making up to 10 partial arc files in this way can sometimes greatly improve the effectiveness of ring reduction. Another method note explains this technique.

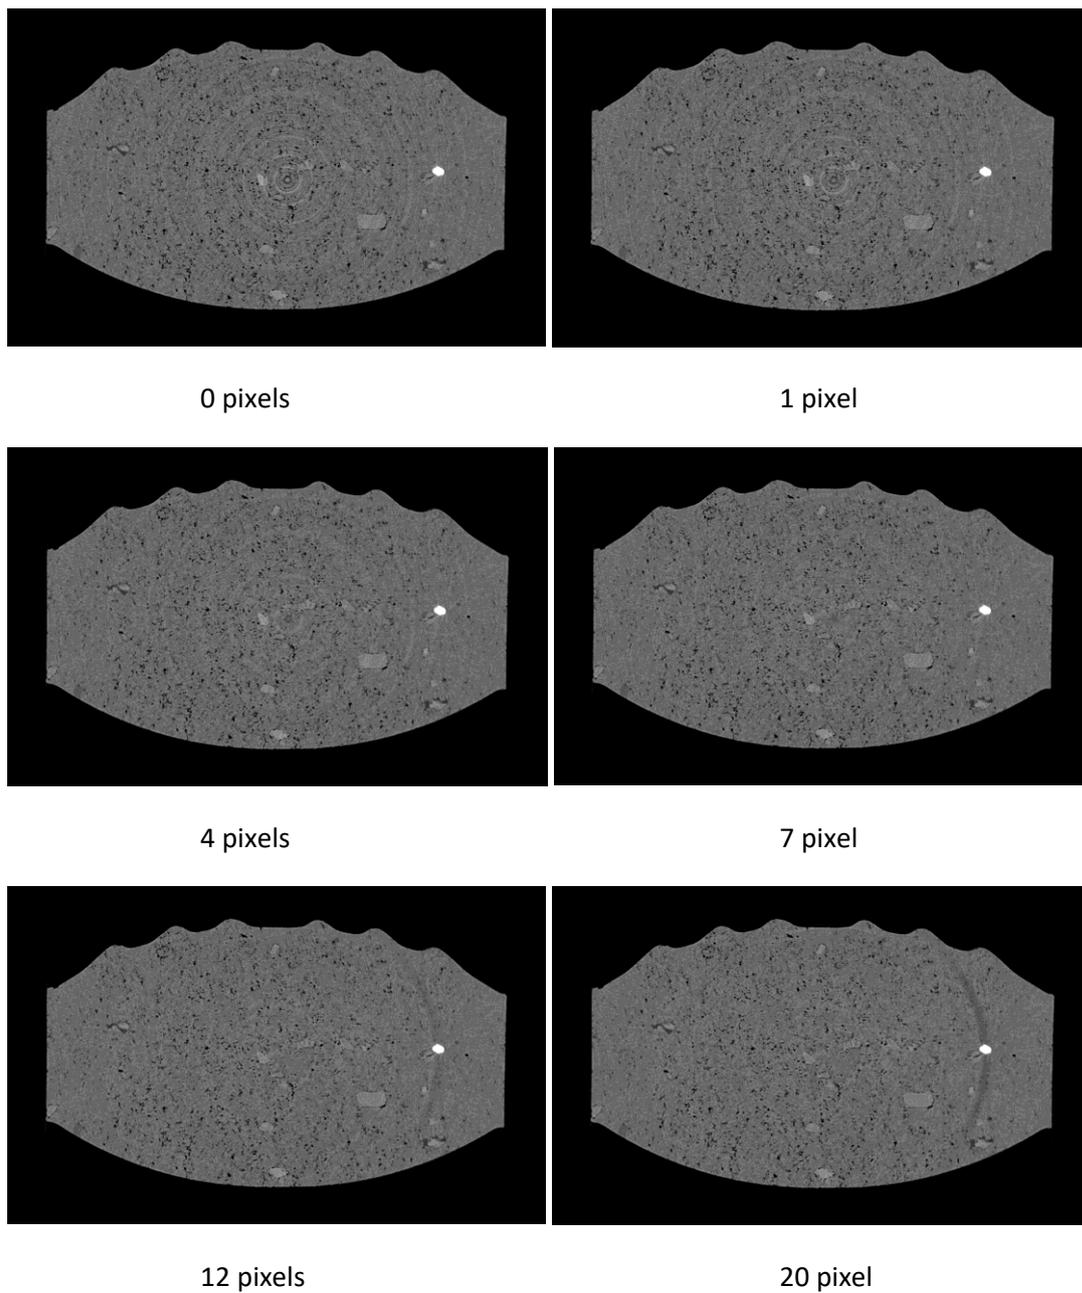


Figure 11. The effect of the ring artefact reduction in NRecon on the appearance of rings in cross-section images. Note the secondary ring around the dense (white) object at the right of the image with the highest RR values. In this example the best value of RR to use is 7.

Beam hardening correction

Beam hardening is an artefact of tomography¹² which makes the object appear artificially more dense at or near its surfaces, and less dense in its central parts.

Beam hardening is minimised by the correct choice of filter in the scan. Increasing filter will decrease beam hardening. Scans with no filter can be impossible to correct for beam hardening with software correction, so some filter should always be used if measurement of density is intended with the microCT scan results.



Beam hardening is selected as a percentage value from 0-100, with increasing percentage meaning a stronger correction¹³.

To help assess beam hardening in a cross-section image you can draw a profile line across the image using drag-and-drop with the *right* mouse button, over any distance and direction. (This right-button draw tool serves three functions: a ruler, an angle measurement – from horizontal – and display of an attenuation or “density”

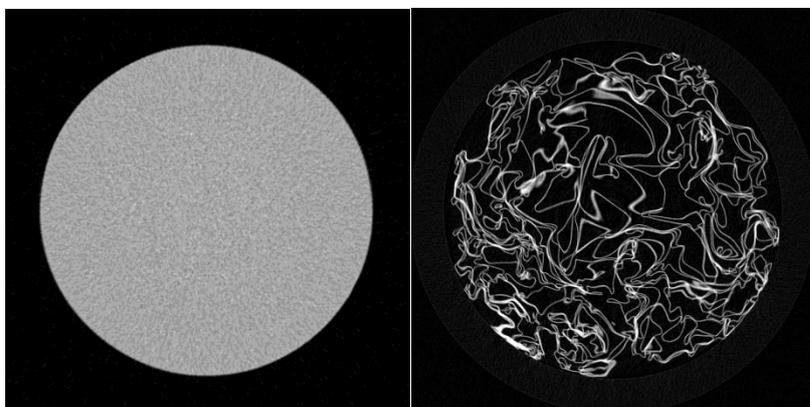
¹² Beam hardening is caused by the fact that x-rays produced in laboratory x-ray sources are “polychromatic”. That means that they have mixed x-ray photon energy. As this mixed beam of x-rays passes through the material of a sample, the lower energy (“soft”) x-rays will be attenuated more rapidly than the higher energy (“hard”) x-rays. Due to this soft depletion the average photon energy increases along the beam path – the beam “hardens” – causing absorption to decrease (for a given material). Since absorption translates into density, the result in the reconstructed cross-section is an artefact where object density is artificially increased near object surfaces and decreased deeper in the object.

¹³ Beam hardening correction is a second order polynomial transformation of the attenuation with thickness curve which tries to restore linearity – the percentage represents the fractional weighting of the quadratic term, i.e. 50 % would mean addition of just half of the squared term while 100% would be adding the full squared term.

profile histogram showing changing image attenuation along the drawn line). Figure 12 shows the profiles of attenuation horizontally across the reconstructed tablet, with values of beam hardening correction from 0 to 80 %. With values of 0 and 20 % beam hardening is clearly visible as the “cupping artefact” with elevated density peaks at the sides of the tablet. 40% correction gives an almost flat profile, while values of 60 and 80% show over-correction of beam hardening with inversed cupping. Thus 40% would be chosen here.

Therefore the profile of attenuation (“density”) across the sample can tell you if the beam hardening correction is too little, too much or just right.

Note however that this is only possible if your sample volume is occupied largely by material with more-or-less uniform absorption. If on the other hand you scan an object consisting of only thin structures, then there is no attenuation gradient of the type seen in figure 12. For example the two images below show cross-sections of a solid rod of aluminium where beam hardening profiles would be easily visible , and a scan of convoluted aluminium foil, where no such profile can be seen.



Left: solid aluminium;

Right: aluminium foil

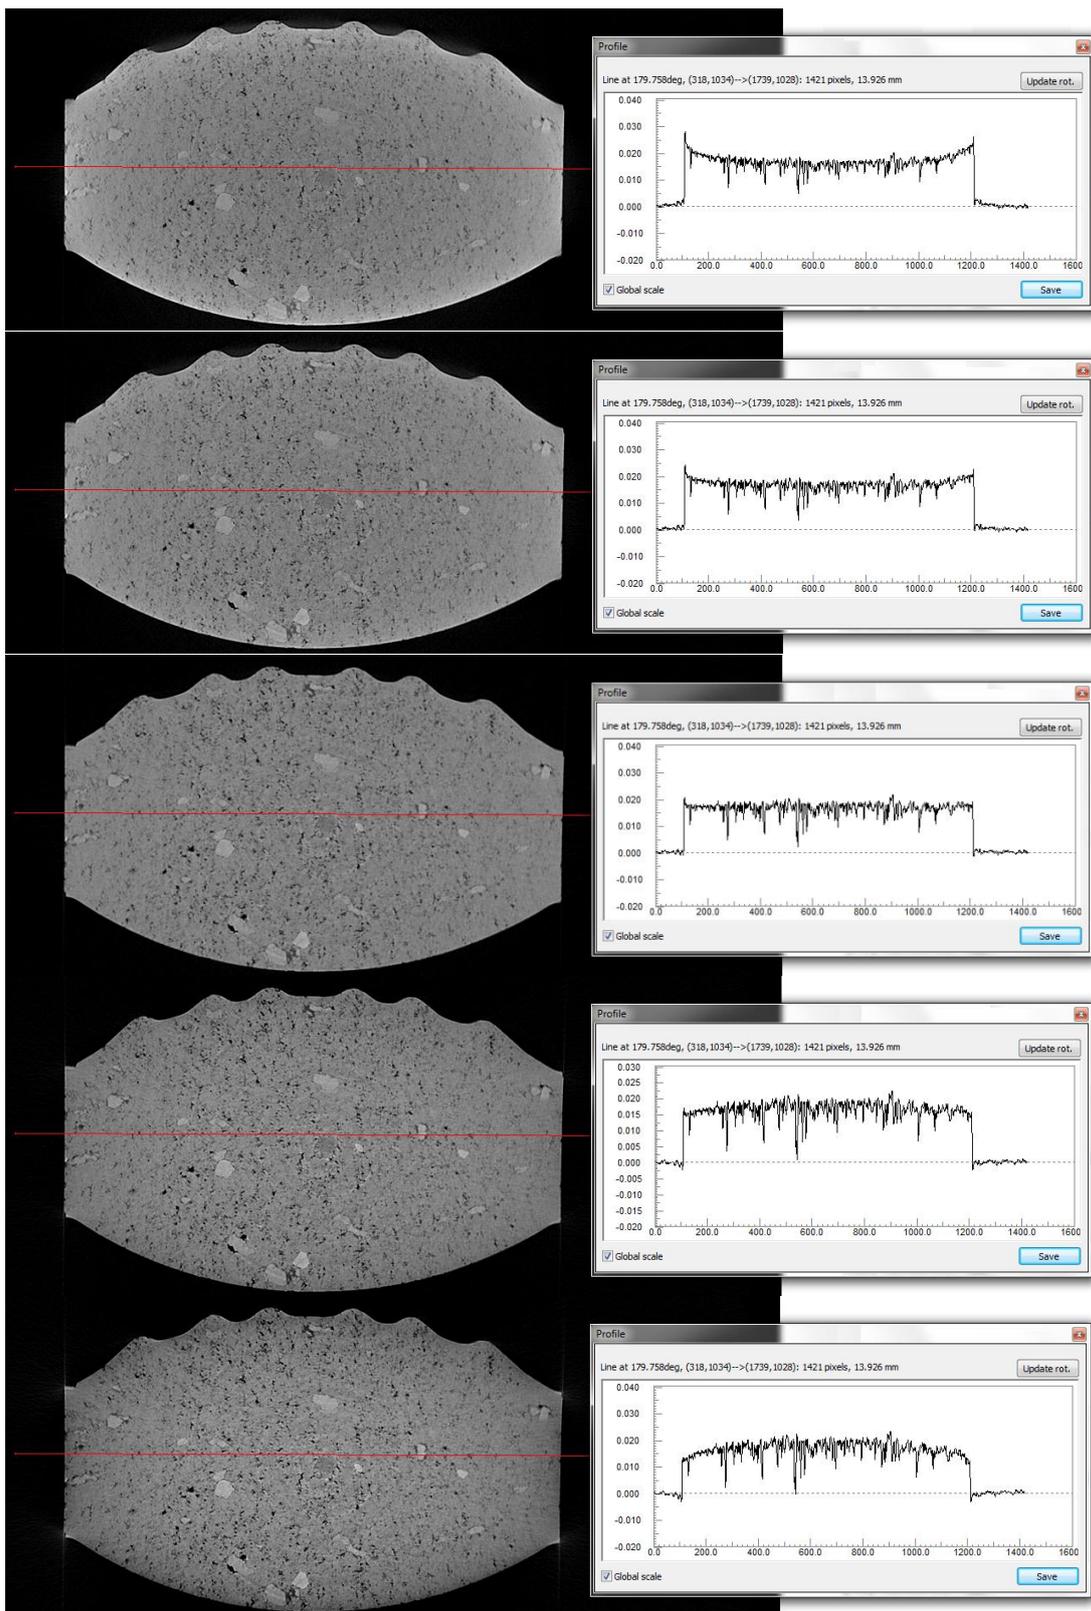


Figure 12. The Profile of attenuation across the reconstructed tablet from left to right (red line) for beam hardening correction values of 0, 20, 40, 60 and 80 percent respectively. It can be seen that 40% beam hardening correction results in an almost flat profile so is the best value to use.

In samples such as the aluminium foil where no profile can be viewed, one has to use one's judgement based on experience with the same or similar materials, in selecting the beam hardening correction value.

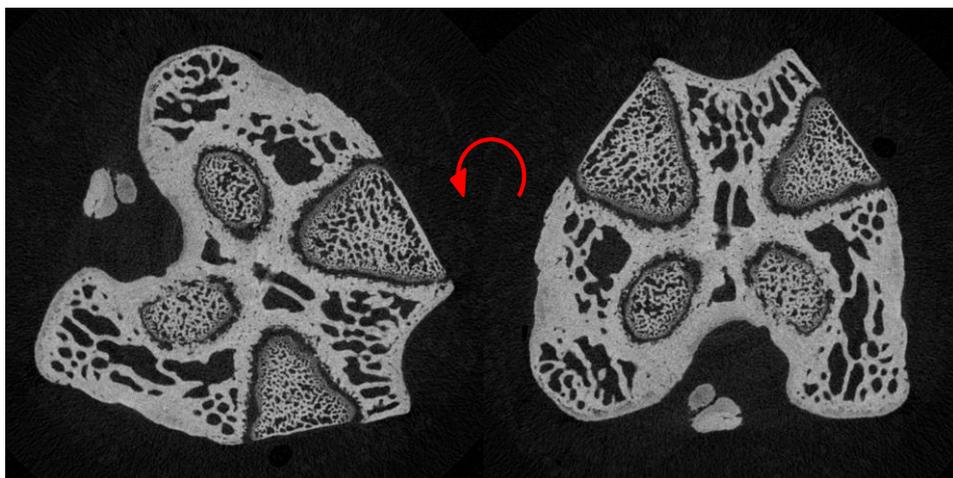
Cross-section rotation

There is the option to rotate the reconstructed cross-section by a selected angle. This rotation is within the XY plane of the cross-section only. Note that the rotation does not change the reconstructed image quality or its resolution in any way.

For objects with asymmetry in the cross-sectional plane it can be useful to standardise the orientation of reconstructed images. For instance in bone biology, standardising the bone cross-sections so that the femur or tibia condyles and growth plate are in a consistent orientation can give more consistency to the manual drawing of volumes of interest, for morphometric and densitometric analysis.

In the example below, rotation of a rat femur cross-section in the vicinity of the growth plate aligns the condyles horizontally:

CS rotation (deg)



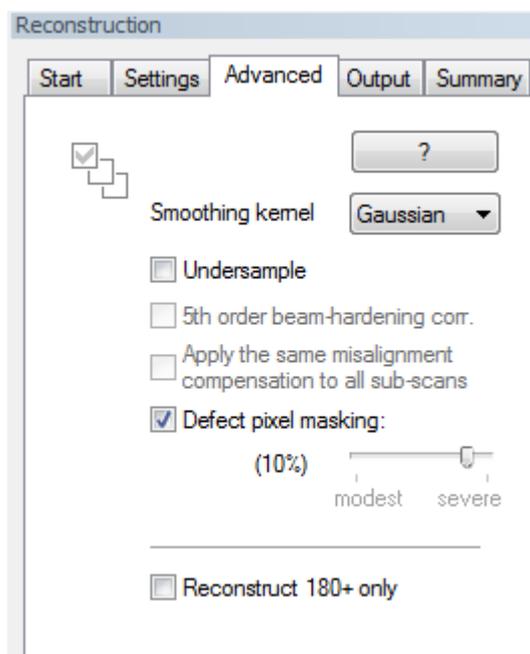
Left: cross-section with no rotation; Right: after -114 degree CS rotation.

Note that a positive value results in a clockwise rotation while negative values rotate the image anticlockwise.

The Advanced tab

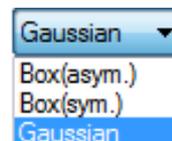
Some additional options are available in the advanced tab. These will not all be described in detail in this document which is intended to give the essential working overview of NRecon. More details on all functions are available in the help documentation within NRecon.

The advanced tab is shown below:



Smoothing kernel

You have a choice of three types of smoothing kernel: box asymmetric, box symmetric, and Gaussian. Generally we would recommend choosing the Gaussian kernel.



Defect pixel mask

The defect pixel mask is a second type of ring artefact reduction. It applies to what might be called “hard” ring artefacts, that is, ring artefacts associated with a specific dot on the camera image caused by a physical particle in the beam path.

An example is shown below in figure 13, from an *in-vivo* scan of a mouse in the SkyScan1176 scanner. Due to rotation of the gantry, occasionally a particle can settle on a surface such as a filter. This can cause a pair of dots – a white and a black dot – to appear on the projection image, as shown below. The black dot is the shadow of the particle itself. The white dot is the flat field correction of the particle, which “misses” the particle when rotation of the gantry causes the particle to shift slightly. Thus this white-black dot pair is characteristic of a particle in the beam path.

The percent value for defect pixel masking is set by a slider, and is adjustable in the range 3-50%. The strength of the correction is the inverse of the selected number – thus a value of 50% is the weakest correction and 3% the strongest. This is shown in figure 13 – only the 3% value can correct this hard ring artefact in the mouse thorax scan.



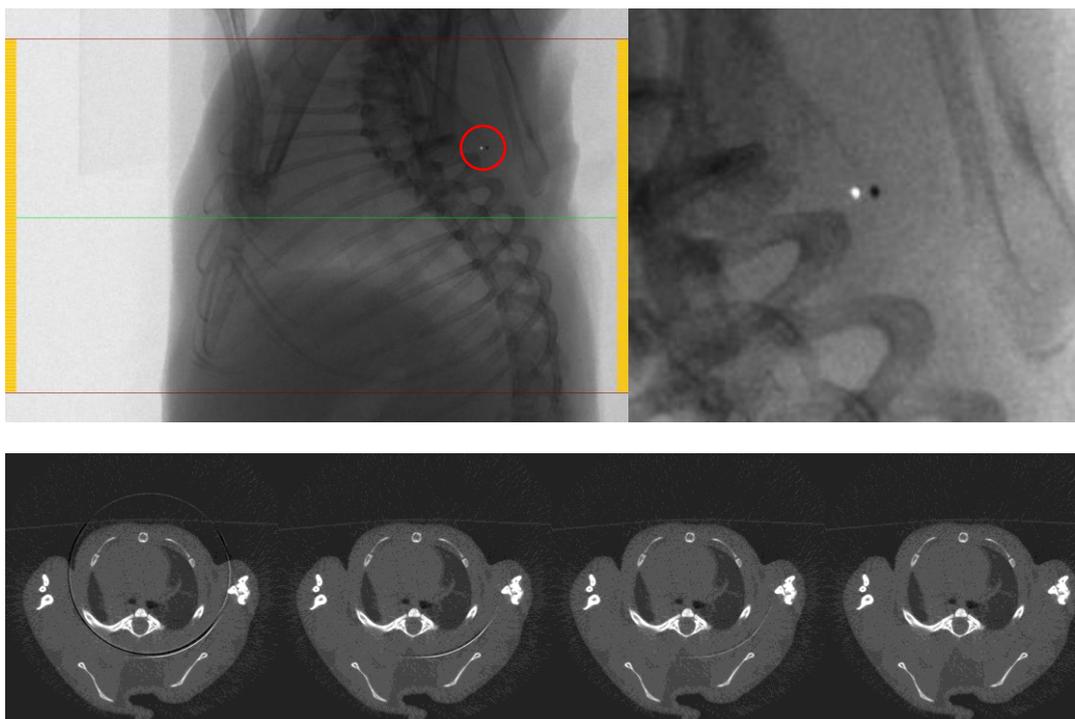


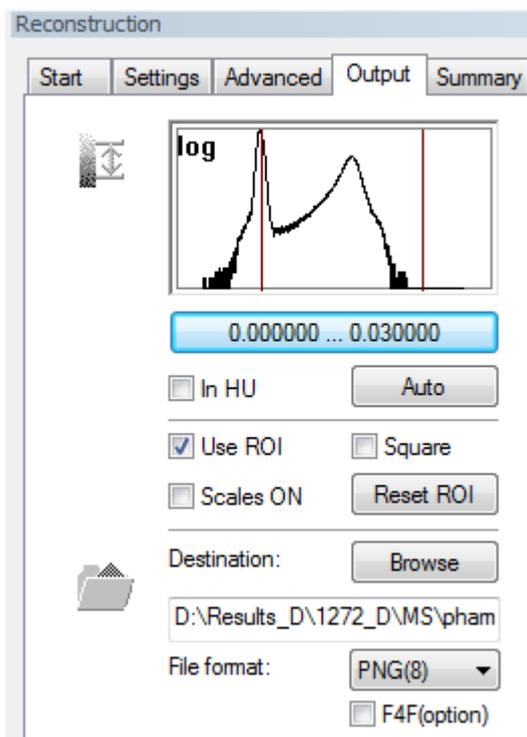
Figure 13. The use of the defect pixel mask in the advanced tab, to remove hard ring artefacts that arise from x-ray dense particles in the beam path. Top left – a mouse *in-vivo* projection image with the particle dot-pair shown by the red ring, and the same dot pair enlarged (top right). Lower images: the reconstructed cross-section at the level of the dots, using no defect pixel mask (left image) and – from left to right – defect pixel mask values of 10, 5 and 3 pixels.

The Output tab

The output tab opens automatically after every preview, with the reconstructed cross-section displayed in the main image pane.

At the top of the output tab is the histogram of attenuation coefficient or “density” of all the pixels in the reconstructed cross-section image.

Tip: double-click with left mouse on the



histogram itself to switch between linear and **log scale** in the vertical Y axis. Generally log is better, you can see the structure of the whole distribution more clearly. The word “log” appears when the log scale is chosen.

The contrast limits

The bar just below the histogram shows two numbers. These are the contrast limits, and they correspond to the two vertical red lines on the histogram. The contrast limits are a very important item to select in setting up reconstruction.

The histogram at the output page shows the reconstructed attenuation of all voxels in the cross-section image¹⁴. The cross-sections are reconstructed in either 8 bit (256 grey levels) or 16 bit (16536 grey levels). Say for example that 8 bit reconstruction is chosen (e.g. the BMP image format). The two contrast limits will represent the attenuation values – as shown by the histogram – corresponding to the lowest and highest grey scale, i.e. 0 and 255.

The lower contrast limit attenuation value will equal a grey scale of 0. The higher contrast limit attenuation value will equal a grey scale of 255.

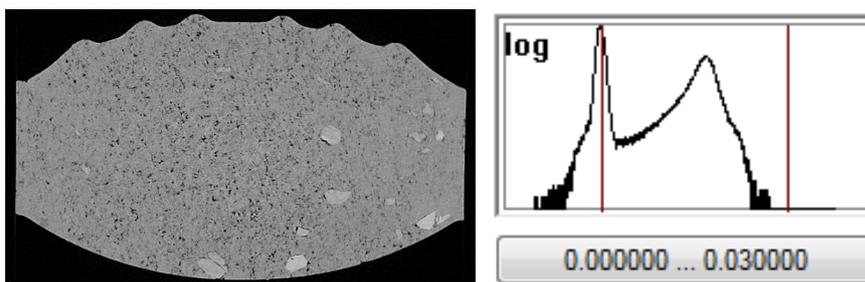
The general guidance for setting the contrast limits is:

¹⁴ Note that this includes voxels at or near the ambient density (i.e. that of air) which have negative attenuation. This is a result of the error spread in the calculation of attenuation for any voxel. For every “real” density phase in the reconstructed sample, attenuation is reconstructed as a bell-like curve centred on the true attenuation of that material, some voxels having higher and some lower attenuation than the true value. Thus for air, whose attenuation coefficient is set as zero, an equal number of voxels will be reconstructed with negative and positive attenuation, centred around zero.

Lower contrast limit: set to zero¹⁵ or a small negative number (5-10% of the high limit, but negative)

High contrast limit: set at an attenuation value about 10-20% higher than the maximum attenuation of the material of principal interest for reconstruction.

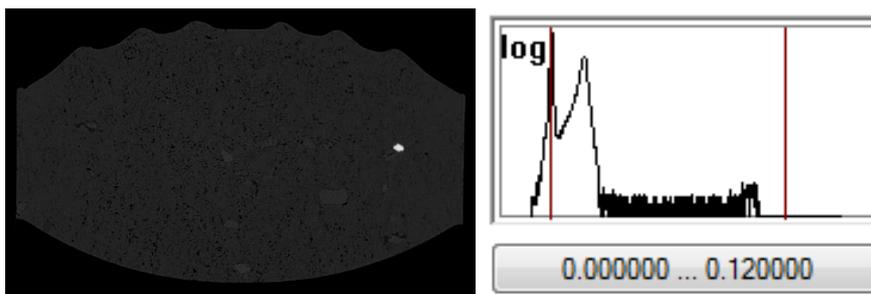
These rules are a good general starting point, although there are exceptions to them. Some of these will be discussed.



In the above example, the lower contrast limit is set at zero, and is close to the center of the low attenuation peak corresponding to air. The high contrast limit is set at 0.03, a little way above the maximum density in the histogram. Since in this cross-section there is not a very large range of densities of material, giving a narrow range of attenuation values, all of the material in the tablet is well visualised with these contrast limits.

However in the cross-section below, from the same scan of the Frisk tablet, there is a particle present with much higher density than the rest of the tablet.

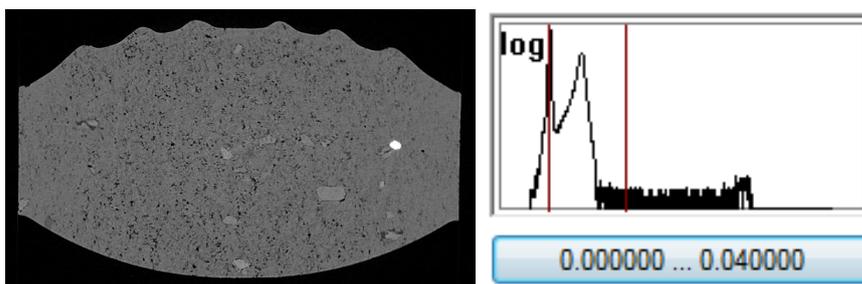
¹⁵ Setting the lower contrast limit to zero essentially links attenuation coefficient of ambient air to zero. This is useful in normalising the relationship between attenuation coefficient and material density, especially if one purpose of the microCT imaging is quantitative measurement of density – including calibration for instance to bone mineral density (BMD).



In this cross-section, the single dense particle results in a long “tail” at the right of the attenuation histogram (log view). If we follow the rule that the upper contrast limit is a little above the maximum, e.g. 0.12, then the particle is well contrasted, but the rest of the tablet is very poorly contrasted and appears pale. Note that in this case, the peak of attenuation values representing the majority of the material of the table, is limited to a small part of the contrast range.

For a material to be well contrasted and clearly visible, the density (attenuation) range of that material should fill most of the contrast range, between the contrast limits (red lines).

In this cross-section slice, we could reduce the contrast range to below the attenuation of the high density particle – to a value such as 0.04. This will cause the image of the particle itself to “saturate”, which means that the true attenuation of the particle is above the higher contrast limit.



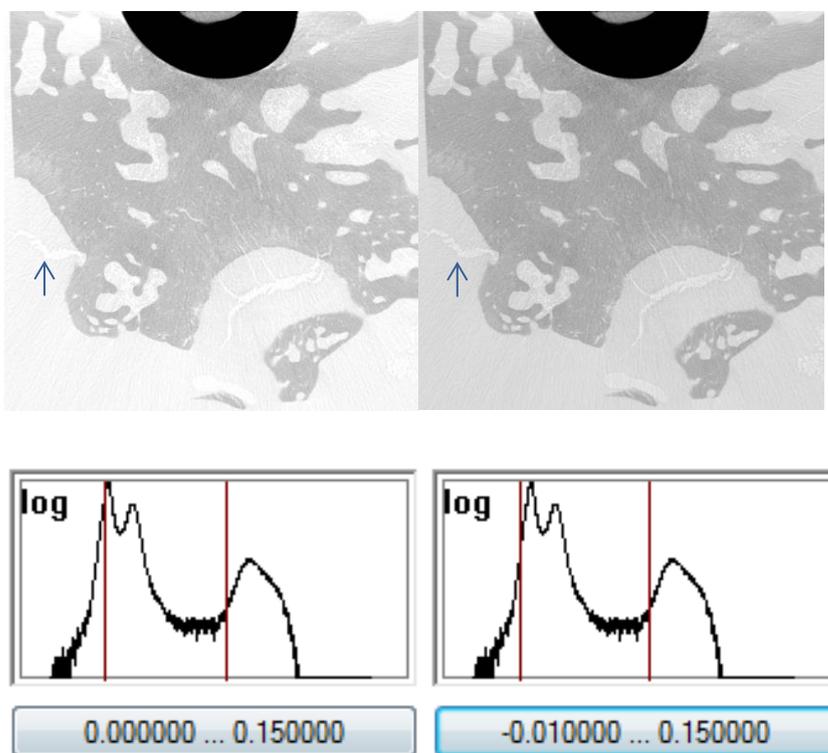
Now the main tablet material is well contrasted, and the high density particle is also clearly visible, although the object is “saturated”, meaning that its pixels all have the 255 white value not representative of the true attenuation of the particle which is much higher.

However, unless the density of the high density particles needs to be measured, it can be acceptable to allow saturation of a material representing a minor inclusion in a scanned object, so that the whole object is more clearly visually contrasted.

This represents an exception to the “rule” that the upper contrast limit should be above the maximum of the attenuation histogram.

Another exception to the rule involves the low end contrast limit. Below are two reconstructed images of a bone sample containing a titanium implanted screw, with the sample embedded in resin. Both the resin and even the bone in this image are relatively low density materials compared to the metal.

Remember that the “default” advice was that the lower contrast limit should be set at zero. In the left image below this is the case. However in the image to the right, the low contrast limit is reduced to a negative value below zero. The attenuation histograms and contrast limits for each image are shown.



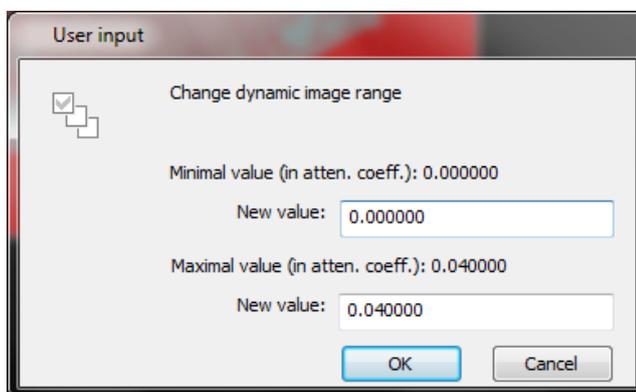
Pay attention to the cracks in the resin surrounding the bone – one such crack is marked with the blue arrows. The cracks can be seen much more clearly in the image where the lower contrast limit is reduced to below zero.

This again is the case where the reconstructed materials comprise a wide range of densities, and furthermore, the density of a material of interest is very near to the low end of the attenuation range. In this case, the low density material is better resolved when the lower contrast limit is decreased below zero.

Finally – you can open a text window to enter the value of the contrast limits directly.

To do this – click on the bar directly under the histogram window which displays the two contrast limit attenuation values.

This is much more convenient for instance for setting the minimum limit to zero, for instance, than trying to do this with the slider (probably not possible).



The “In HU” and “Auto” buttons

Just below the histogram pane and contrast limit settings there is a “In HU” and an “Auto” button.

In HU

Auto

The “In HU” button allows a Hounsfield unit calibration to be applied¹⁶. It is a limited HU calibration function since it requires the presence within the current cross-section of water or a material that can be considered equivalent to water. Therefore, while this button does provide the option for HU calibration of reconstructions, we recommend that a better alternative for HU calibration is to perform it in the Bruker-microCT CT-Analyzer (“CTAn”) program. This allows an HU calibration, based on an attenuation measurement of water in a separate calibration scan, to be transferred to other datasets. A separate Bruker-microCT method note exists for HU calibration with CTAn, which is applied retrospectively to datasets after reconstruction.

The “Auto” button applies an automatic method to set the contrast limits. This can be useful as a first step if the limits are grossly incorrect such that the image is either saturated (mostly white) or of very low intensity (mostly black). However it is best not to rely on the auto limits and determine them manually.

The ROI

The “Use ROI” tick box makes a red box appear in the reconstructed image. This is the “region of interest”. If it is selected when reconstruction is started, then only the image inside the ROI box is reconstructed, all of the cross-section outside the box is omitted.

¹⁶ Select a region of water or equivalent on the cross-section, then with the CTRL button held down, draw a square with right mouse button drag-and-drop. A text box will appear; within this box click (left-mouse) on “Hounsfield unit calibration”. Now a new “Hounsfield unit calibration” dialog box will open, showing at the top a field displaying the “HU value at the selected position”. This is the non-calibrated value. To apply the HU calibration i.e. to set the chosen region as a water reference, change the number in the field to zero, then click “OK”. Below this number field there is a display of the average attenuation coefficient for the same chosen region. The chosen region can only be a square / rectangle.

The ROI is important and useful, since it prevents you reconstructing large areas of image which are outside of the objects of interest and unnecessary for the further analysis or visualisation of the dataset.

Reducing the volume of the reconstructed dataset by applying an ROI, together with setting an appropriate vertical range for reconstruction, will reduce the data volume of the reconstructed dataset to the minimum actually needed (see figure 14 below).

Thus using an ROI where possible is strongly recommended.

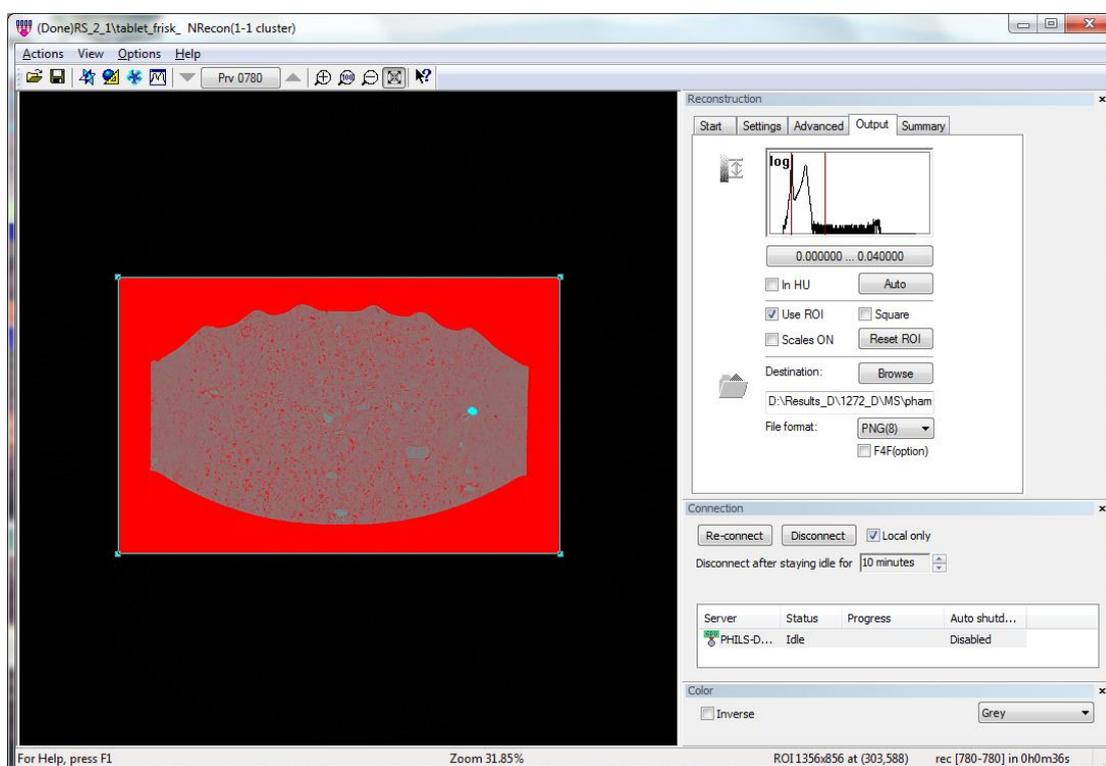


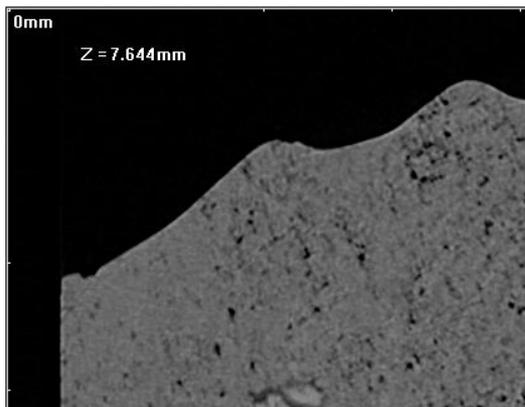
Figure 14. The region of interest (ROI) – indicated by the red rectangle, is an important step for restricting reconstruction to the object(s) of interest.

The scale bars

- Scales ON** Selecting “scales” adds scale marks to the borders of the reconstructed image.

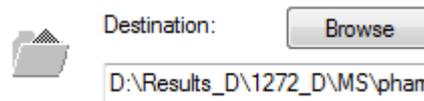
One should be aware that the binary scale marks are a permanent part of the images, and – unless steps are taken to remove them – the scale marks will appear as objects in 3D visualisations or in 3D analysis, which is not what you want!

In general scale marks are used where needed for display purposes or to satisfy image requirements of a journal referee for instance; but not for routine microCT analysis.



The destination folder

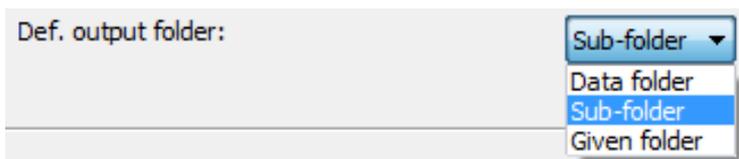
The destination folder into which reconstructed cross-section images are written, is displayed at the lower part of the Output tab.



To specify a folder manually, click on the “Browse” button.

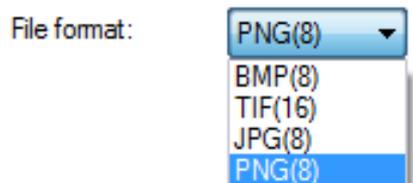
To see the full path of the destination folder, click in the displayed part of the path and hit the “end” button on your keyboard.

Please note – it is recommended that you specify in NRecon preferences that reconstructed cross-sections are always written in a subfolder, which is given the dataset prefix name followed by “Rec”. To do this, go to the options menu and preferences, and near the top, after “Def. output folder” choose “Sub-folder”:



The file format

After “File format” you have a drop menu of four image format choices for the reconstructed images. The number in brackets indicates whether the reconstructed images for each format are 8-bit (256 grey levels) or 16-bit (65536 grey levels).



Tif (16-bit) is recommended in situations where (a) there is a wide range of densities of material in the cross-sections, and (b) you might wish, after reconstruction, to analyse the images with different contrast limits for different density phases within the same images. With 16-bit tif images, you can change the contrast limits by reducing the range (but not enlarging the range) when loading into CT-Analyzer. So if this is your intention when choosing the 16-bit tif format, you should set wide contrast limits to include all materials present.

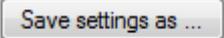
With the 8-bit formats the contrast limits are fixed and not adjustable.

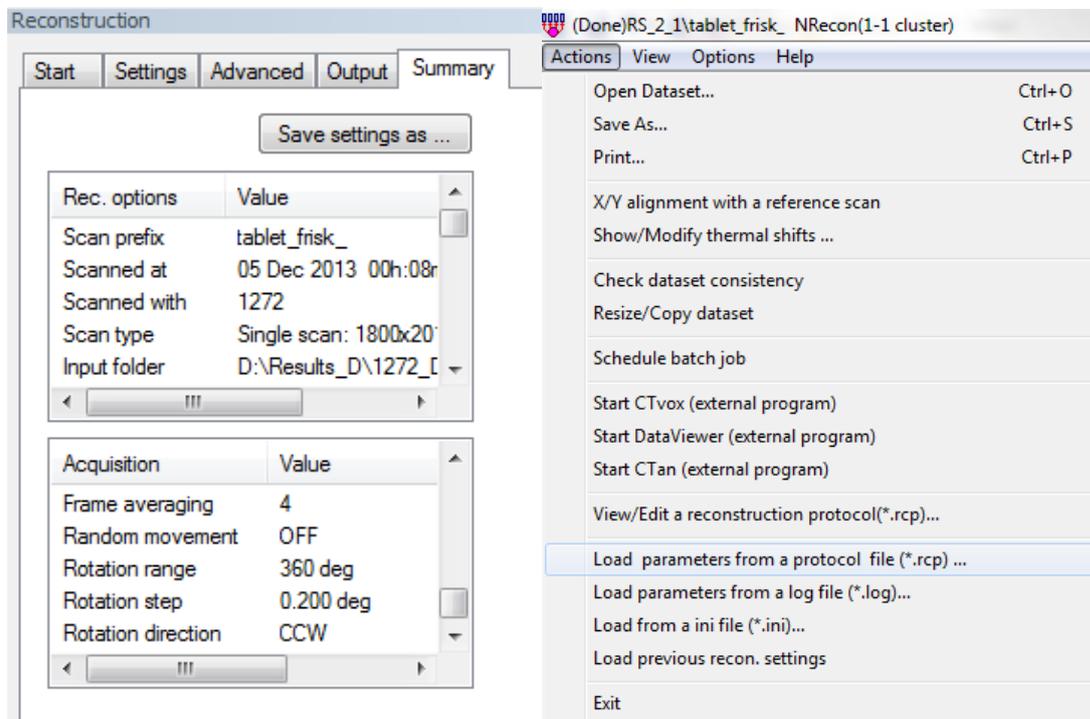
Refer to the section on contrast limits on pages 28-30. The tablet with high density inclusions, and the bone with titanium implant, are examples where you could consider 16-bit tif reconstruction.

The Summary tab

The summary tab displays technical information about the scan dataset that is currently loaded – data that can also be found in the scan log file (NRecon is reading this data from the log file).

The protocol file

An important function is located at the summary tab – the  “Save settings as” function that allows you to save a reconstruction protocol file.



The protocol file allows you to quickly load all the reconstruction settings from another scan¹⁷. Thus it is useful when you have many samples of the same type that you wish to analyse comparatively, so you need to scan them with consistent parameters.

The window that opens when you click on “save settings as” is shown below in figure 15. Let’s say I wanted to create a protocol file for scanning the “Frisk” tablet. I could give it a name such as “tablet”, as shown below, and save the file. What is saved is a file with the extension “.rcp” (reconstruction protocol file). If I later load another

¹⁷ You will see that under the Actions menu it is also possible to load reconstruction settings from the log file of another scan – where reconstruction has been done.

scan of a Frisk or similar tablet, and wish to apply the same reconstruction settings, then I can load these settings from the “tablet.rcp” protocol file by going to the “Actions” menu of NRecon, and choosing the item near the bottom “Load parameters from a protocol file (.rcp)” (see image above, right).

Some comments on the protocol file:

- Not all parameters should be taken from another reconstruction. The most important example is the **post alignment correction**, which can be different for every scan. In the “save settings as” window, always set the “Post alignment” value to the option “Use default if applicable”. This allows NRecon to calculate the misalignment compensation separately for each scan.
- The ROI is another setting that requires individual adjustment for each scan, since the position of each sample when scanned will be different.
- Likewise the “CS rotation” parameter for rotating the image, is also specific to an individual scan only. It is best to set CS rotation to zero when saving a protocol file.
- Protocol files are typically saved in the NRecon software folder.
- You can change the default settings in NRecon by saving these settings as “NreconDefault.rcp”.

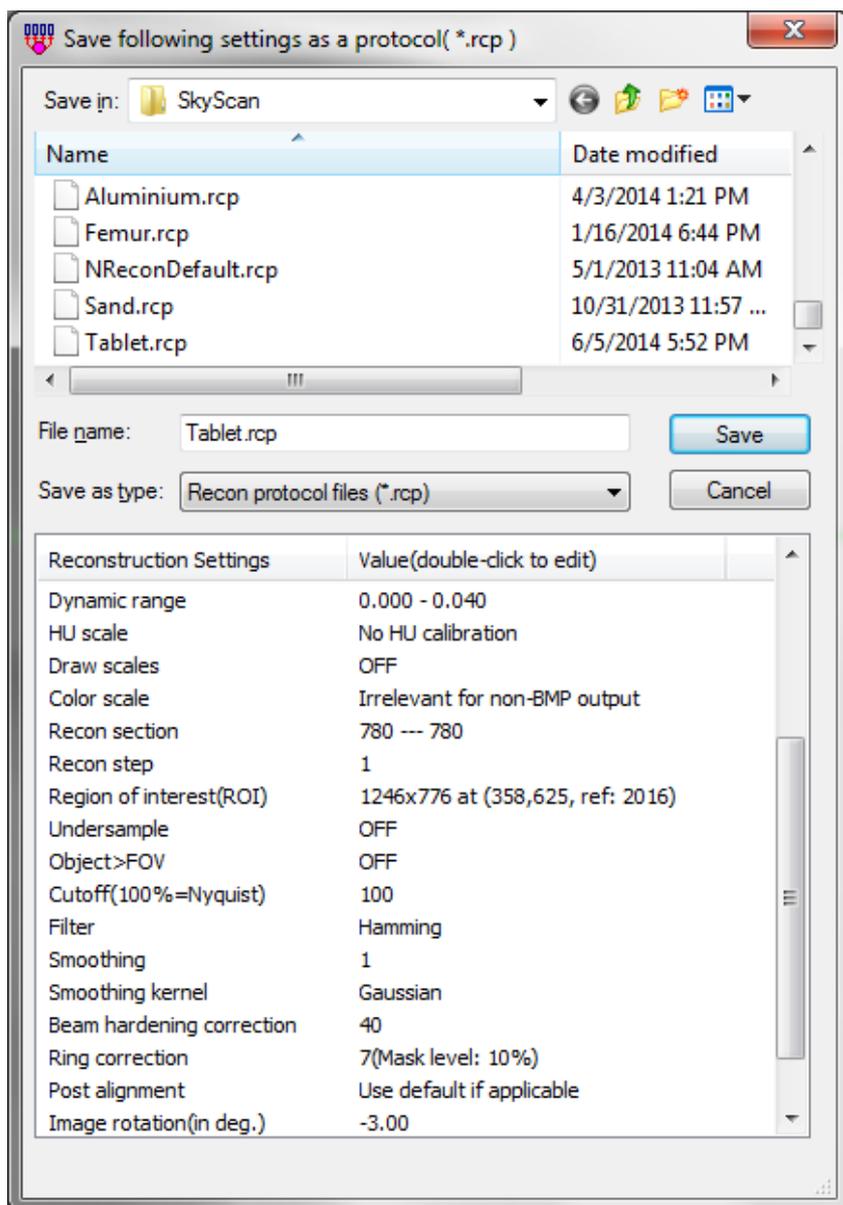
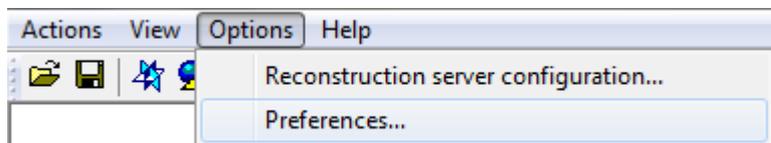


Figure 15. The window for setting up a protocol file. It takes the currently set parameters from the Settings, Advanced and Output tabs of NRecon, but the window is interactive, allowing editing of individual parameters.

The NRecon Preferences: recommendations

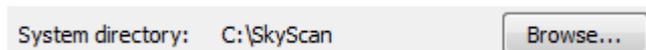
The preferences window is opened from the Options menu.



The contents of the preferences window are shown in figure 16 below.

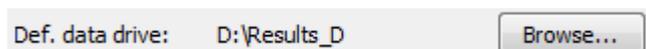
Recommendations for preference settings

The system directory



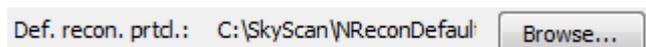
Set this to the folder containing *Nrecon.exe* and the other SkyScan software files. This is important – some NRecon problems can occur when this is not set correctly.

The default data drive



Set this to the drive and location where you most often store your scan files. (This is not critical, only for convenience).

The default reconstruction protocol file



This file contains default reconstruction settings. It should be located in the system directory (see above). If it is not, you can create this protocol file giving it the name

“NreconDefault.rcp” and saving it in the system directory, from the settings tab and the “Save settings as” button.

The default output folder



You are given a drop menu with three options for the default output folder. This means the folder where the reconstructed cross-section images will be saved. We recommend that you choose “Sub-folder”. In this case the cross-sections will automatically be saved in a sub-folder which is given the same name as the scan filename followed by “Rec”. So if your scan prefix was “sample_1_” then the subfolder would be named “sample_1_Rec”. This simplifies data management and copying of reconstruction image results.

Automatic shutdown of server PC(s) (NRecon cluster version only)

Automatic shutdown all server PC

This option makes the PC(s) running the server program shut down automatically after reconstruction is finished. This is not used so often now with the increased speed of reconstructions.

This preference item only appears in the cluster version of NRecon.

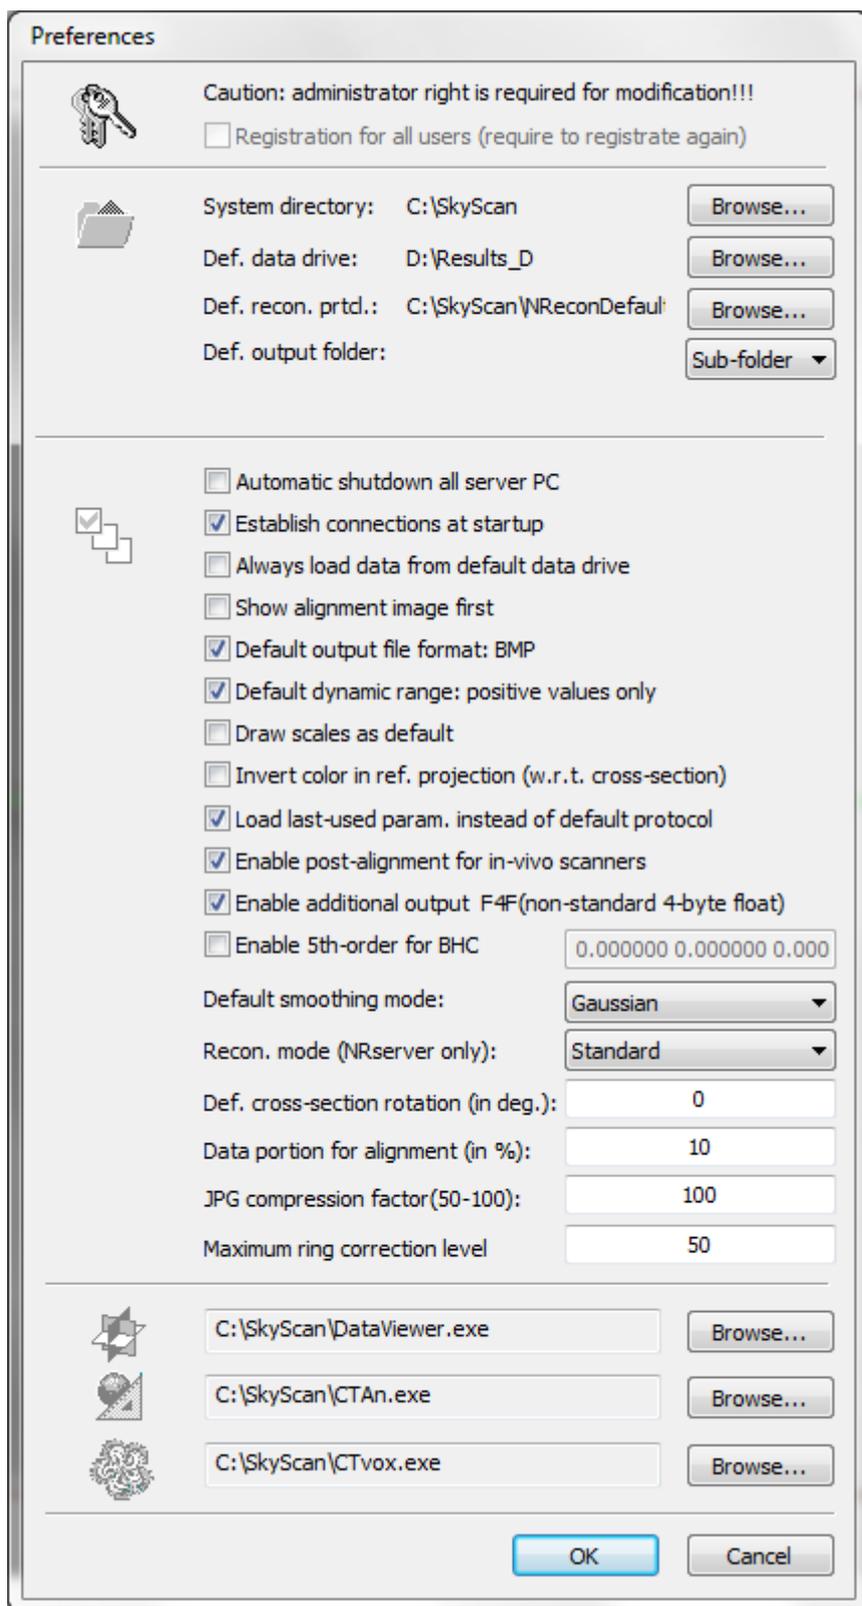


Figure 16. The contents of the preferences window. This is opened in the Options menu. (Ticked and un-ticked items are for illustration, not mandatory.)

Establish connections at start-up (NRecon cluster version only)

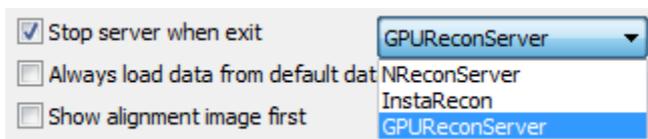
Establish connections at startup

With this selected, NRecon will make a connection with the server program that is selected in the “Options / Reconstruction server configuration” window – provided that program is open. In the local version of NRecon, that server program will open automatically with the opening of NRecon. We suggest that this item is ticked.

This preference item only appears in the cluster version of NRecon.

Stop server when exit (NRecon local version only)

With this selected, the server program will close automatically when NRecon is closed. To the right there is a drop-menu where you should specify the server program which you are using (as selected in the “Reconstruction server configuration” menu item – also in the Options menu just above Preferences).



Always load data from the default data drive

Always load data from default data drive

Not generally recommended, can be used where reconstruction is launched automatically after scanning.

Note that this preference item and all following preference items are present in both the cluster and local versions of NRecon.

Show alignment image first

Show alignment image first

When this is selected, then on opening a dataset NRecon will immediately open the “Settings” tab under the “Reconstruction” pane. Within this settings tab is a color-coded central strip of the projection image showing the super-position of the 0 and 180 degree projection images, used by NRecon for automatic calculation of the (default) Misalignment compensation value. Choose according to preference, not critical.

Default image file format

Default output file format: BMP

You can choose the default file format for reconstructed cross-sections, between BMP, PNG, JPG (all 8-bit) and TIF (16-bit). Note this is only the default, it can be manually changed for any reconstruction.

Default dynamic range: positive values only

Default dynamic range: positive values only

This makes the default contrast limit range have a minimum value of zero. Note that it can always be changed manually.

Draw scales as default

Draw scales as default

We normally recommend not selecting this option since the scale marks become a permanent part of the reconstructed images and require extra actions to remove them in order to make 3D visualisation or analysis of the dataset.

Invert color in ref. projection

Invert color in ref. projection (w.r.t. cross-section)

This inverts the “SPR” scan projection reference image that is created with each dataset reconstruction.

Load last-used parameters instead of default protocol

Load last-used param. instead of default protocol

This option can be useful for doing sets of scans for experimental purposes where most scan settings need to be kept the same (see FAQ discussion of this issue). Note it affects parameters only at the settings page, excluding the contrast limits at the “Output” tab and selections in the “Advanced” tab such as the defect pixel mask.

Enable post-alignment for in-vivo scanners

Enable post-alignment for in-vivo scanners

This option should normally always be selected especially when you are using an *in-vivo* scanner.

Enable additional output F4F (non-standard 4-byte floating point)

Enable additional output F4F(non-standard 4-byte float)

This option allows floating point cross-section output in the “raw” format.

Enable 5th order for beam hardening correction

Enable 5th-order for BHC

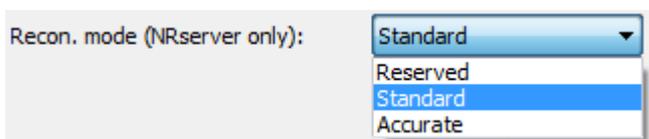
By default beam hardening correction involves the application of a second order polynomial transformation of the thickness-attenuation curve. Optionally however you can substitute this with your own (experimentally obtained) polynomial expression with up to five terms. Please consult the extensive on-line help documentation on this subject within NRecon.

Default smoothing mode



The mode or method of smoothing can be selected between three options in the advanced reconstruction tab. This preference setting allows you to choose the default smoothing type. Generally we recommend Gaussian smoothing.

Reconstruction mode



When using the NReconServer, there is an option to choose “Accurate” rather than “Standard” reconstruction which means slightly more background noise reduction¹⁸ but with longer reconstruction time. The “Reserved” option is for research purposes. Normally this should be left at Standard.

Default cross-section rotation



This should normally be set to zero, unless a fixed rotation is required for all datasets.

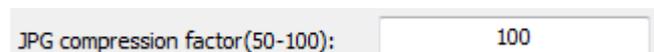
Data portion for alignment



¹⁸ Bilinear interpolation is used instead of nearest-neighbor interpolation in standard mode.

This determines the width of the central horizontal slice of the projection image which is subject to the “pseudo-parallelization” process and used to calculate the misalignment compensation in the setting tab. This is 10% by default. A wider angle can improve the results of calculation of the misalignment, but with increasing cone angle (i.e. distance above or below the midline) the pseudo-parallelization becomes less effective. Increase of this value up to 30% is acceptable.

JPG compression factor



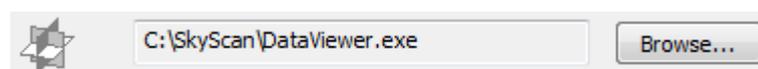
When JPG is chosen as the reconstructed cross-section image format, you can set in preferences whether this is lossless JPG (100%) or a bigger image file size compression value with a degree of image data loss. Generally we recommend setting this to 100% (no loss) unless a very large compression factor is required and some data loss is acceptable.

Maximum ring correction level



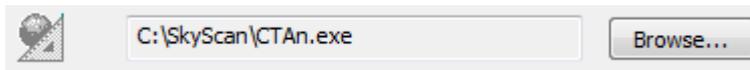
By default a maximum ring reduction value of 20 is available. However this maximum can be extended up to 100. Please note however – care should be taken with large ring reduction values that secondary ring artefacts are not being created (see figure 11, page 22).

Directory path for DataViewer



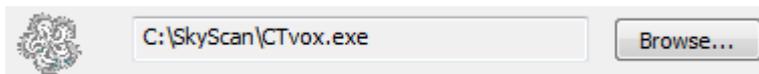
It is important to set the path to the *dataviewer.exe* program file to allow automatic opening of reconstruction results from NRecon by *DataViewer*.

Directory path for CT-Analyzer



It is important to set the path to the *ctan.exe* program file to allow automatic opening of reconstruction results from NRecon by *CT-Analyzer*.

Directory path for CT-Voxel



It is important to set the path to the *ctvox.exe* program file to allow automatic opening of reconstruction results from NRecon by *CT-Voxel*.

2. FAQ

What does “FAQ” stand for?

Frequently asked questions 😊

Section 2 says that I can reconstruct across a network, using projections on another computer. But I tried this and got an error message about permission, read-write access etc. How do I fix this?

What this means is that the folder on the network computer containing the projection images has not been shared, to give other PCs access to the files. Folders to be used for reconstruction across a network must be shared, giving the highest level of sharing permission including read-write access, since the reconstructed images will be written to the same folder – or a subfolder.

Is it OK to run reconstruction on the scanner workstation PC while a scan is in progress?

This can be done, but with certain restrictions and with caution. Several years ago the advice would have been “no”. Reconstruction uses a PC's resources (processor, RAM, disc read-write) very intensively, and with older PCs there was even a small risk of a complete computer crash (blue screen) if scan and reconstruction were run together. However current more powerful PCs with larger RAM, 64 bit architecture and multiple processors make a crash much less likely. However the InstaReconServer program uses the PC's resources even more intensively than the other server programs, so for InstaReconServer our advice would remain not to run reconstruction at the same time as scanning, if the InstaReconServer program is running on the scanner workstation PC. If however InstaReconServer is running on another PC in a “1-1” cluster, then it is (probably) OK to run reconstruction together

with scanning. This is also true for NReconServer and GPUReconServer on both a “1-1” or a larger cluster. Note that in the “Reconstruction server configuration” window under the options menu, you should make sure that the scanner workstation itself is not included in the cluster of PCs running the server program – so only the separate, connected PCs are doing the reconstruction server calculations.

I get error messages when I try to run GPUReconServer – what to do?

There are a few simple steps needed to make sure GPUReconServer works on your scanner workstation or other PC. First – find the file named “gpu_reg.reg” in your scanner software folder (copy it to another PC if necessary). Double click on it and then click on “Yes” twice. This makes a small change to your PC’s registry to enable smooth operation of GPUReconServer. Second – go to the NVidia website and driver downloads, and download and install the latest driver for your PC’s graphics card. To find what this card is, go to the Windows control panel, “System” and “Device manager” and look under “Display adapters”. Finally – right-click on your desktop and click “NVidia Control Panel”. Select “Manage 3D settings”. Find the option “Threaded optimisation” and turn it off.

Is it OK to use smoothing – will I lose image data?

Yes it is OK, sometimes excess noise compromises both viewing and analysis of images. Thus you can often improve image quality by smoothing. It should be judged visually – what looks best often is best.

Should I always use beam hardening correction?

Generally yes, always use a minimum of 10% beam hardening correction, usually more.

How important is it to get exactly the right value of each setting?

There is probably not a “perfect” value of most parameters, so use your judgement, generally what looks better is better. In many cases a little plus or minus will not greatly affect your image analysis results.

With beam hardening correction, is more always better? Should I not always use 100% to make sure beam hardening is corrected?

Definitely no! With beam hardening correction, too much is as bad or even worse than too little or none at all. Check the section of this document that shows that too much beam hardening correction causes an inverse distortion of the density profile, plus excess noise in the image.

With ring artefact reduction, is more always better? Should I not always use a high value to make sure ring artefacts is corrected?

Again definitely no. With ring reduction as with beam hardening correction, too much is as bad or even worse than too little or none at all. Check the section of this document that shows that too much ring reduction can cause secondary ring artefacts. Use only the minimum necessary value.

What is the defect pixel mask? How should I use it?

The defect pixel mask is a second method for ring artefact correction. It applies to ring artefacts caused by a dot feature on the camera image often caused by a

particle on a filter surface or elsewhere in the beam path. The range of values from 50% to 3% is inverse, so that the 50% value has the weakest effect while 3% is the strongest. If you have a problem with a specific ring artefact caused by a particle, the particle itself should of course be removed by cleaning the affected surface. This is a matter for your distributor's service engineer. However the defect pixel mask is a software method which can sometimes remove the ring artefact caused by the particle. Start with 20% and then try decreasing values of the defect pixel mask until the ring disappears. Note however – as with other correction methods, a strong (i.e. low) % value can sometimes cause new artefacts in the image, so you should carefully look out for any such secondary artefacts. Use a mask value at least 3-4 percent points above a value that causes artefacts. (Refer to the section above on ring artefact reduction.)

Which reconstruction settings need to be kept the same between scans, for quantitative measurements (comparing apples with apples, etc...)

The most important ones to keep the same are the contrast limits, beam hardening correction, smoothing, and the object larger than FOV option. These affect the fundamental grey distribution of the image. Ring reduction (and defect pixel mask) ideally should be the same but can be varied if necessary. Post-alignment correction is an exception – it should not be kept the same but optimised individually for each scan. Other settings such as ROI (size and position) image rotation and the range of slices can be varied between scans – indeed this is usually necessary.

Why is all this so complicated? Can't some automated magic just make all these decisions for me?

Sadly no. In microCT as in life in general, what you get out approximates to what you put in. It gets easier with practice – or at least it makes a bit more sense. Enjoy!

Appendix: Some examples of suggested parameters

The table below makes some recommendations for scan settings (filter, resolution level) and some reconstruction parameters for a range of different sample types. Please note that these suggestions are of a general nature and should be interpreted with flexibility.

Scan category	Scan resolution level ¹⁹	Smoothing	Beam hardening correction	Ring artefact reduction
No filter				
Paper etc., low density low volume	High	0	15	7-13
Insect, plant material, organic non-calcified, low-density tablet	Medium-high	0-2	25	7-13
Larger volume of biological tissue (brain, kidney, muscle, tissue biopsy etc.) or of other low density material	Medium	3-5	25	7-25
0.25mm Al filter				
Mixed density biomaterial e.g., porous calcium phosphate, tablet (medium density)	Medium-high	1-3	40-50	5-11
Soft biological tissue with partial calcification, e.g. late embryo	Medium-high	2-4	40-50	5-11
Larger volumes of low density material e.g. wood, plastic	Low-Medium	2-4	40-50	3-13
Very small medium density materials e.g. bone, ceramic, rock, polymer	Medium-high	0-2	40-50	5-13

¹⁹ Depending on your scanner type you usually have 2-3 resolution “levels” available. For example in the SkyScan1172/1272 desktop scanner, the SkyScan2211 nanoCT (using the CCD camera) and the in-vivo SkyScan1176, there are three resolution levels of 1k, 2k and 4k referring to the pixel width of the acquired projection image with 4x4 binning, 2x2 binning and no binning respectively.

Scan category	Scan resolution level	Smoothing	Beam hardening correction	Ring artefact reduction
0.5mm Al filter				
Mouse bone	Medium-high	0-2	45	4-9
Mouse hindlimb <i>in-vivo</i> (high resolution/high dose)	High	1-3	30	4-9
Mouse body <i>in-vivo</i> (high contrast/high dose)	Low-medium	3-5	30	2-5
Small rock, mineral, ceramic sample, very small tooth	Medium-high	2-4	40-70	4-9
Large volume – low density e.g. wood, plastic, water, bio-tissue	Low-medium	3-5	30	2-7
1mm Al filter				
Rat, rabbit bone	Medium	1-2	35	5-11
Small tooth	Medium	1-3	35-50	5-11
Mouse hindlimb <i>in-vivo</i>	High	0-2	30	4-9
Rat hindlimb <i>in-vivo</i>	Medium	0-2	30	2-5
Mouse body <i>in-vivo</i>	Low	3-5	30	2-5
Rat (small) body <i>in-vivo</i>	Low	4-6	30	3-7
Rock, ceramic, glass	Low-medium	1-3	30-50	4-9
Biomaterial, medium density	Medium-high	1-3	30-50	5-11
40 micron Cu + 0.5mm Al (equivalent to ~2mm Al)				
Large tooth, large bone	Medium	2-4	25-35	5-13
Metal containing sample	Low-medium	2-6	50-100	1-4
Rat (large) body <i>in-vivo</i>	Low	3-7	25-30	3-7
Metal implant in bone	Low-medium	2-4	30-50	1-5
Rock, ceramic, other material with high absorption	Low-medium	2-4	35-50	5-11

Scan category	Scan resolution level	Smoothing	Beam hardening correction	Ring artefact reduction
0.11mm Cu and higher filter				
Metal sample	Low-medium	2-4	30-60	2-7
Rock, ceramic, other material with high absorption	Low-medium	2-4	25-50	2-7
High-low density combination e.g. metal implant in bone	Low-medium	2-4	25-50	1-5
Very large rat <i>in-vivo</i>	Low	4-7	20	2-7