

April 2017

SmartAlign *for DigitalMicrograph*

Image Registration and Scan-distortion
Compensation Plug-in

SmartAlign Manual 1.2 (Provisional)

HREM Research Inc.

Preface

This user manual describes how to install and operate the *SmartAlign* image registration and scan-distortion compensation plug-in (version 1.x) for Digital Micrograph. The software is based on the full-length manuscript by Dr Lewys Jones and Prof Peter D. Nellst:

Jones, L., et al.. "Smart Align—a new tool for robust non-rigid registration of scanning microscope data". *Advanced Structural and Chemical Imaging*, **1 8** (2015). [Link](#)

As this tool applies post-processing steps to experimental data, users should state explicitly when such steps have been applied by citing the above paper such as:

“image alignment and scan-distortion compensation was performed using the SmartAlign software described in [1]”.

Most of this manual describes the processing of atomic resolution image series acquired in the scanning transmission electron microscope (STEM), however the same steps apply to data from the scanning tunneling microscope (STM) and the atomic force microscope (AFM). A separate section at the end of this manual describes settings specific to these physical probe techniques.

Contact

Enquiries regarding sales, licensing, installation and operation of the plug-in should be directed to:

support@hremresearch.com

For questions about the academic manuscript referenced above, please contact Dr Lewys Jones at:

lewys.jones@materials.ox.ac.uk

Copyright Statements

© Copyright 2015-2016 HREM Research Inc.

All rights reserved. This manual is protected by international copyright laws and treaties. Unauthorized reproduction and distribution of this manual, or any portion of it, will be prosecuted to the maximum extent possible and may result in severe civil and criminal penalties.

DigitalMicrograph is a trademark of Gatan Inc. _____

Contents

When to use this Plug-in	5
Installation	6
Software Requirements	6
Software Installation	6
Typical Workflow	7
Plug-in Menu	8
Supported File Types.....	8
Data Acquisition	9
'Control' Tab.....	9
Auto Rigid Alignment on Finish	11
'Signal' Tab	12
Data Preparation.....	13
Data Shaping Options.....	13
Data Scaling Operations	14
Scan Metadata Options.....	14
Rigid Registration	16
Selecting the Area for Rigid Registration	16
Selecting the Rigid Registration Options.....	18
Data pre-processing Options.....	18
Registration Method	19
Output Options	22
Non-rigid Registration	24
The Types of Iterative Registration Cycles	24
The Number of Iterative Cycles.....	25
Types of Non-linear Cycle.....	25
Smoothing Kernel Size.....	26

Registration Convergence - the Target, Acceleration & Number of Cycles	26
Storing the Transformation Data	27
Choice of Reference Image	27
Output Options	28
The Progress Bar	28
The NR Outputs and Image Cropping	29
Aligning Extra Image Signals	29
Aligning Spectrum Image Series	31
Additional Menu Items	32
Output Display Options	32
Advanced Options	32
Special Case Examples	34
Digital Super-Resolution Example	34
Processing STM & AFM Data	36
Advanced Filtering for Rigid Registration	38
Troubleshooting	39
Auto-save function not working with some detector channels.....	39
References	40

When to use this Plug-in

The *SmartAlign* plug-in for Digital Micrograph is designed to correct the most common image distortions encountered in atomic-resolution image series recorded in the scanning transmission electron microscope (STEM), including image drift between frames and scan-distortion within each frame. These are shown schematically below:

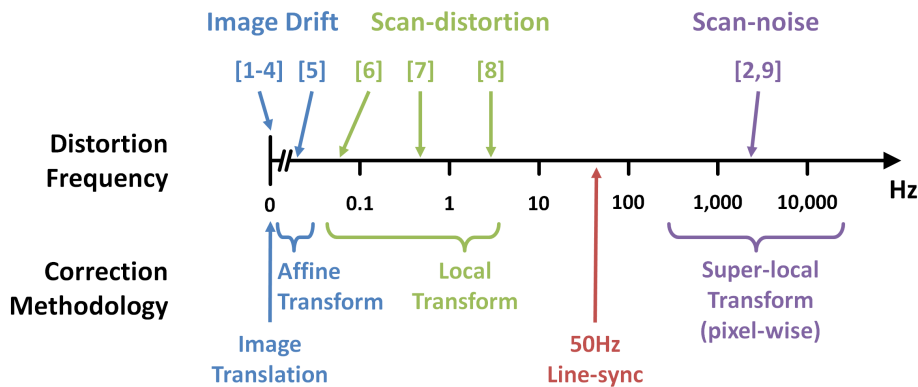


Image distortion frequency domains. For a description of the references above, please see the associated academic manuscript (Jones et al., 2015).

SmartAlign is used to correct image-drift and scan-distortion. Scan-noise however, caused by instabilities up in the kHz regime, can be corrected using the *Jitterbug* plug-in also available from HREM Research (Jones & Nellist, 2013).

Image offsets between frames in a series are readily identifiable. These can be easily corrected to simply the analysis of time-series or focal-series data. Scan-distortion can appear as additional image transformations, including affine transformations (stretch, skew / shear, small rotations) and non-linear distortions, causing otherwise straight crystal lattice planes to become ‘wavy’.

If your data includes these types of shifts and/or distortions, then it may be possible to recover some or all of the lost performance caused by scan-distortion. Of course if the image is very heavily degraded it will not be possible to compensate this entirely. Equally this plug-in cannot create resolution that did not exist prior to the incorporation of scan-noise. For further information and some related reading please see the reference given on page 2.

Installation

Software Requirements

The following is a list of the software requirements necessary to run the *SmartAlign* plug-in:

- DigitalMicrograph (GATAN™) with DigiScan and EMControl
- USB Key Driver
- HREM Mouse Tool Plug-in (Free-ware downloadable from www.hremresearch.com)
- IPU Plug-in (Free-ware downloadable from www.hremresearch.com)
- Acquire Image Series (Free-ware downloadable from www.hremresearch.com)
- Matlab MCR (Matlab Compiler Runtime) v8.2.

Software Installation

Installing USB Key Driver

The user key driver should be installed by following the instructions given by the key driver installer. The key driver installer can be found on our web site.

Installing SmartAlign Plug-in

The plug-in can be installed by drag-and-drop copy to the folder “PlugIns” (The PlugIns folder should exist under a normal installation of the DigitalMicrograph.)

When the DigitalMicrograph is launched after placing the plug-ins into the PlugIns folder, SmartAlign menu commands will appear under “SmartAlign” menu.

Installing Mouse Tools Plug-in, IPU Plug-in and Acquire Image Series Plug-in

All the files relating these tool plug-ins, can be installed by drag-and-drop copy to the folder “PlugIns.” (The PlugIns folder should exist under a normal installation of the DigitalMicrograph.)

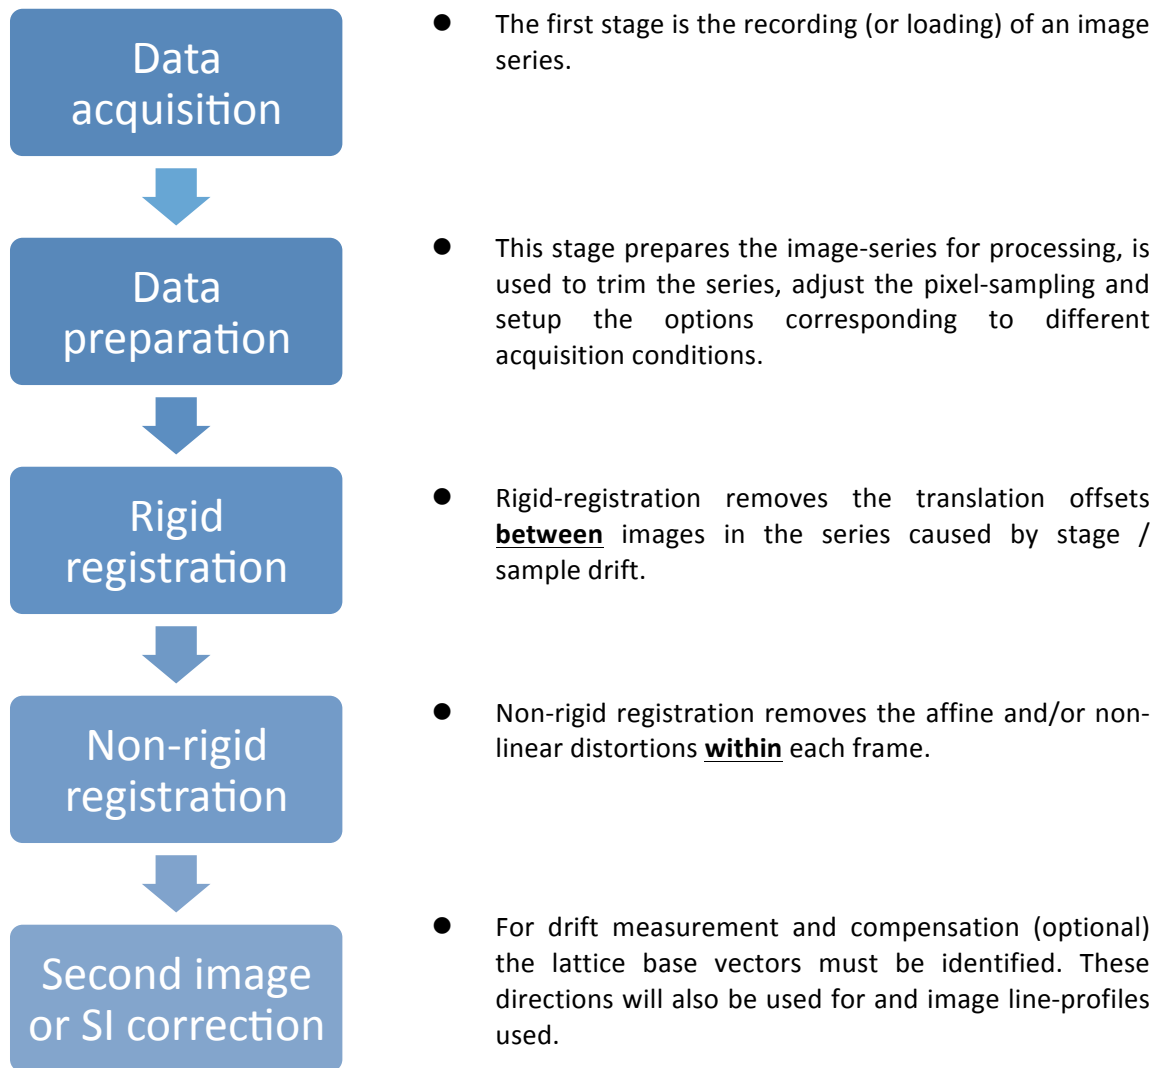
When the DigitalMicrograph is launched after placing the plug-ins into the PlugIns folder, the Mouse tool will appear as an addition to the standard tools.

Installing Matlab MCR

The Matlab MCR should be installed by following the instructions given by the installer. The Matlab MCR that is compatible with the SmartAlign can be found on our web site.

Typical Workflow

The SmartAlign software is very flexible in how it is used, however, most tasks will follow a similar workflow. These are summarized below and then discussed in more detail in the following sections.

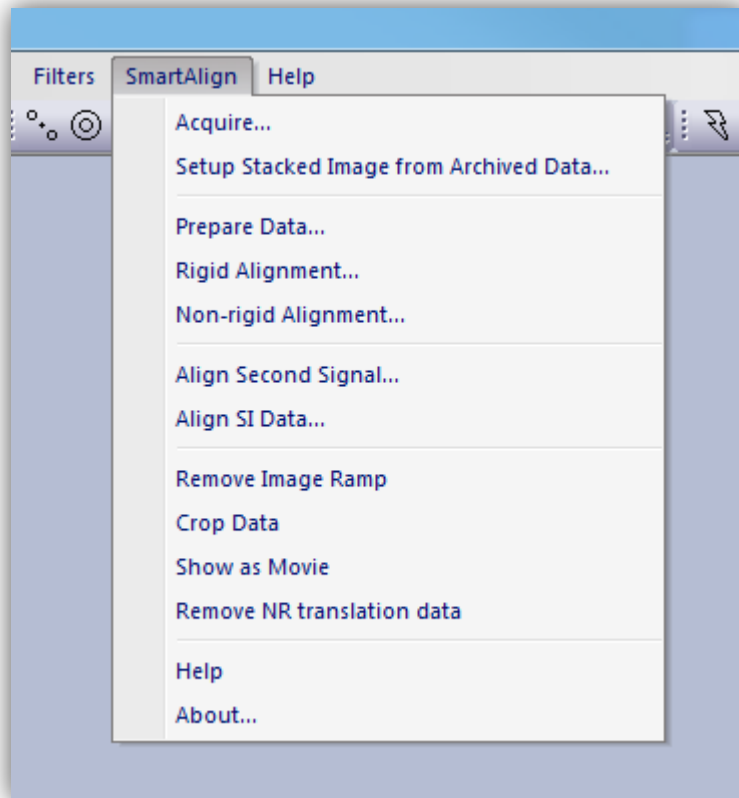


The plug-in is designed so that each stage in the processing creates a new data window for its results. If for some reason the calculation fails, or the user wants to go back one stage and choose a different setting, the new window can simply be closed and the result from the previous stage will still be there.

The following examples will use the supplied example image-series of a SrTiO₃ nano-cube edge to demonstrate the various steps.

Plug-in Menu

The figure below shows the default menu for the SmartAlign plug-in:



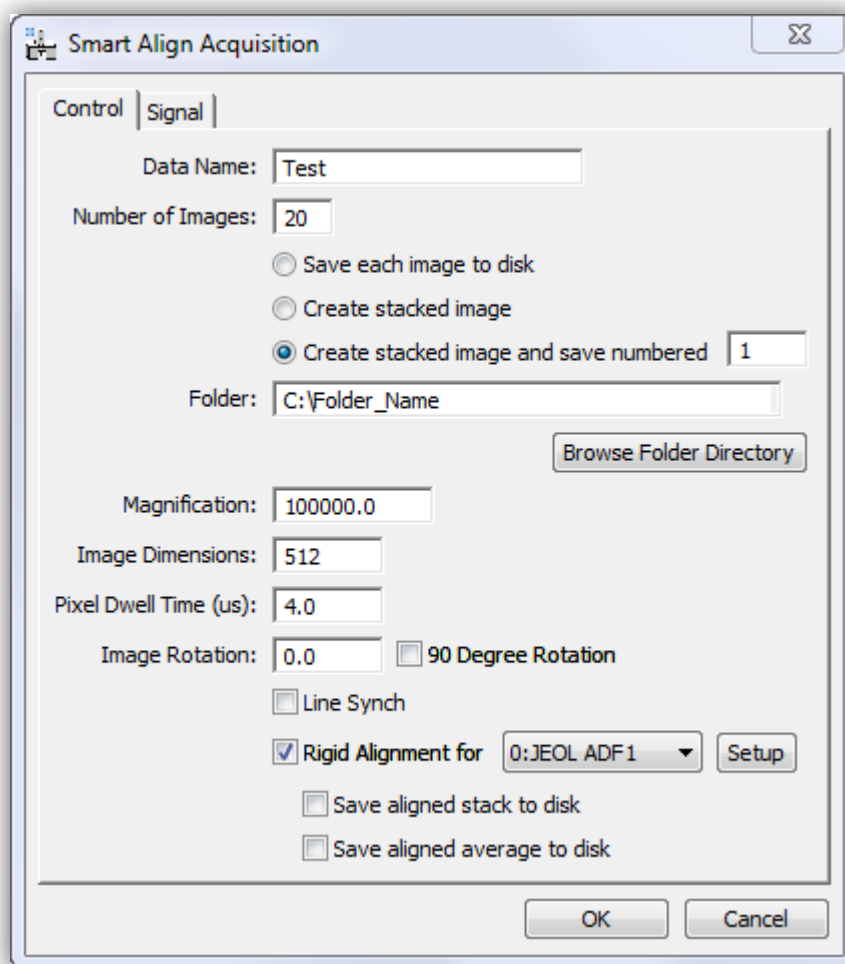
Supported File Types

To begin, the user should have the file open within the Digital Micrograph workspace. To open a file either open it from your computer or from within DM by going to File > Open. DM supports images-series being saved as a single file (as slice data), however if the images were recorded as separate files, they can be imported using the 'File' > 'Open Series...' option to produce the stacked data on import.

The plug-in supports all image formats supported by DM, however, it is strongly advised that the images to be processed have received no image compression (such as JPEGs).

Data Acquisition

A tool for the recording of series is included with *SmartAlign*.



'Control' Tab

Data Name

Here you can enter a filename that will be assigned to the recorded image series. If the user chooses to create a stacked image (see next section), *and* use the auto-numbering and save feature, each filename will be created using underscores to concatenate the fields with the following format:

“Data Name” + “_” + “Sequential Number” + “_” + “Detector Name”

For example:

My Sample Name_001_JEOL ADF1

If you change to a new sample / region / experiment you can choose a new 'data name' and reset the numbering by simply typing the next desired number. Numbers are given in a three digit format.

Number of Images

This specifies the number of images in the series. After recording, the series can always be trimmed. If the stability or damage tolerance of the sample is unknown, it may be helpful to overestimate the desired number of frames. If damage or contamination is observed, then the user can later decide to reject the later part of the series.

In addition to specifying the number of images, the user should choose whether to write the images individually to disk or to display the images as a single stack (as slice data). In general, a single stacked image is more convenient. However, if recording very large images, or an extremely large number of frames it may be preferable to write directly to disk.

In the case of creating a stacked image, the user can further select to auto-save this file.

Folder

If an auto-saving option was selected, then images will be saved into the specified folder. Either type to location directly or select the folder using the 'Browse' button.

Magnification

This will be automatically read from the microscope and added to the image series meta-data. This does not need to be edited unless the user changed the magnification after opening the acquisition menu.

Image Dimensions

This will be the height and width of the image measured in pixels. The recorded images are square.

Pixel Dwell Time

This sets the dwell-time for each pixel in units of microseconds. The time for one scan-line will be the image width multiplied by this number. To adjust other time parameters such as the flyback waiting time, use the advanced settings in the normal STEM Digiscan dialogue window.

Image Rotation

This sets the overall rotation of the scan, performing the same function as 'scan-rotation' in the Digiscan settings. The scan-rotation of acquired data has an effect on the ease of processing the series.

For the rigid registration, the image alignment will make use of correlations calculated in Fourier-space (more details later). To maximise the precision of this stage it is advisable to ensure that when recording image series, the Fourier spots (including their cross shaped streaks from the finite image border) do not impinge upon one another.

For the non-rigid registration, if a low-order crystal lattice plane is aligned to the fast-scan direction, then often some scan-rows will contain only minimal contrast. As it is image contrast which is used to identify and correct for the scan-distortion, it is preferable for scan-rows to not be devoid of contrast.

Both of these conditions can be fulfilled by adding a small scan-rotation off of the crystal plane (say 10° or so).

A further robustness to sample / stage drift can be incorporated by adding an additional scan-rotation stepping of 90° between each image frame. In this case the images will be rotated back when added to the stack.

Auto Rigid Alignment on Finish

If selected, an automated rigid registration can be performed immediately after the acquisition has finished. This will only correct for image translation (whole frame offsets), however it allows the images in the series to be summed for inspection. The user can then check, for example, the sample tilt or image defocus.

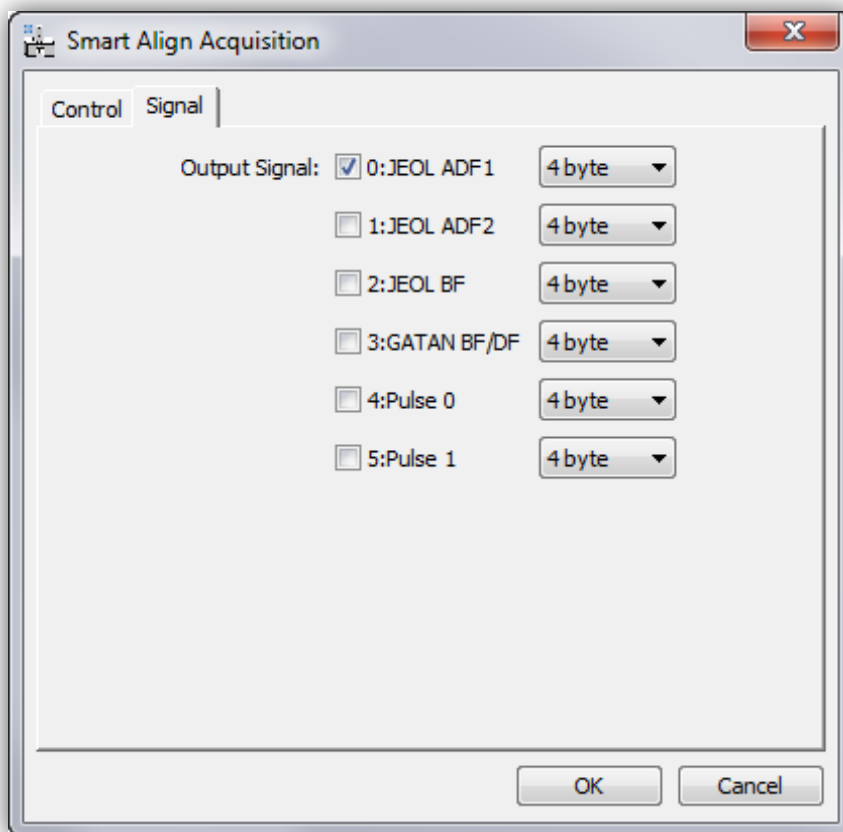
If this option is selected, then one signal can be selected as the inspection signal. Other signals can be registered later. The 'Setup' button allows for the rigid-registration settings to be adjusted.

To maximise workflow productivity, the outputs from this post-acquisition alignment can be immediately saved to disk, with the option to also save the average through the series as a single frame.

'Signal' Tab

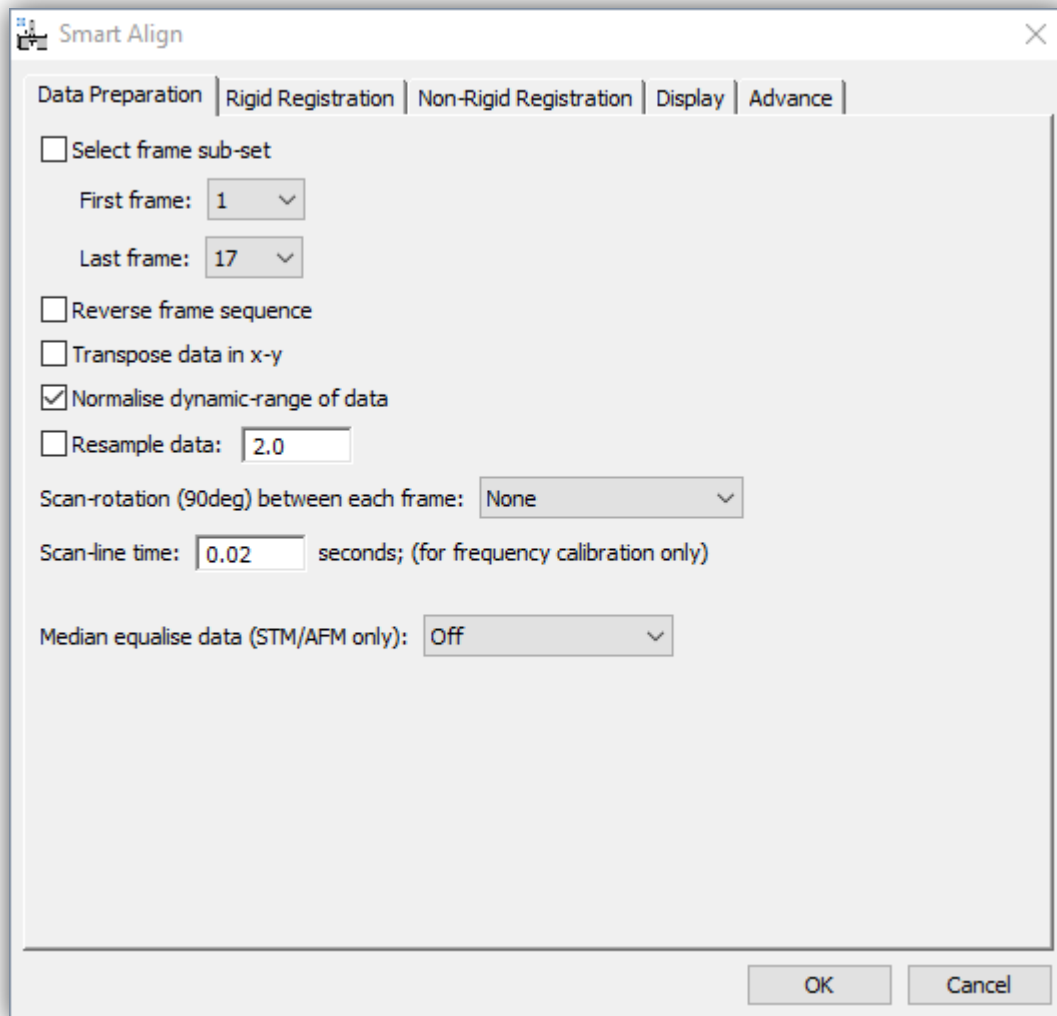
On this tab the user can select the signals to be recorded. If multiple signals are selected these will be recorded simultaneously. Importantly, this means that the same probe-offset vectors will apply to every signal.

In this tab the number of signals available to the Digiscan, as well as their names, will be automatically detected.



Data Preparation

The data preparation stage should be executed before any either of the registration stages:



Data Shaping Options

In this menu the user can choose to select only a sub-set of the frames in the series. To do this select the checkbox and then choose the first and last frame number (note frame counting here is one based and not zero based).

The user can also select to reverse the frame order; this may be useful when processing or comparing focal series.

The user can select to transpose data; this may be useful if the data has been imported from an FEI *.ser file or from another package such as MatLab.

Data Scaling Operations

Normalising the Data Grey-level Range

For most applications the numerical values of the STEM image pixels are not needed. In this case the data may be scaled (subtracting some constant and dividing by another) with no loss of qualitative information or relative contrast.

There are two reasons why the user may want to do this; for data with a large constant offset (for example a large amplifier black-level) the data will have a large mean grey-level. This will give a large DC offset in the Fourier based cross-correlations used for rigid registration. Subtracting a constant reduces this offset term and may help with the rigid registration. Additionally, the later non-rigid registration calculation is based on image intensity gradient matching and dividing the range of the data by a constant may accelerate the calculation speed.

Note, if performing a quantitative analysis, where the pixel intensities do have real meaning, this option should not be used.

Data Resampling

This option allows the user to resample the data in x-y before any onward processing. Resamplings smaller than one will reduce the number of pixels in the image (more coarse sampling), and values larger than one will increase the number of pixels in the data (more fine sampling). Reducing the sampling may be useful for condensing data from large CCD cameras or to do initial quick calculations on less powerful computers. Increasing the sampling may be useful to achieve digital super-resolution. In general, increasing the sampling by more than, say, three or four will not lead to any improvement but will slow the subsequent calculations significantly due to increased memory usage.

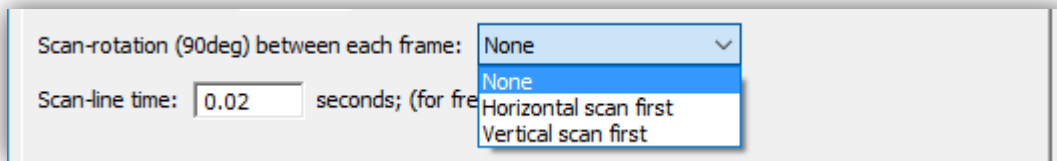
Digital super-resolution is an advanced topic and a separate example is given at the end of this manual.

Scan Metadata Options

Scan Rotation Between Frames

This option toggles different modes in *SmartAlign* depending on how the image series was recorded. Specifically, it changes the 'row-locking' and 'row-fitting' behavior used in the non-

rigid registration options. Users should see the associated manuscript for a further description of the 'row-locking' and 'row-fitting' behaviours.



The default setting is 'None', this means that no change in scan-rotation was used during the experimental acquisition. Alternatively, *SmartAlign* can accept data with multiples of 90° scan-rotation change between frames. For this data can be recorded as 'orthogonal pairs', that is scanning at 0°, 90°, 0°, 90°... etc., or as an 'orthogonal set, i.e. 0°, 90°, 180°, 270°, 0°, 90°... and so on.

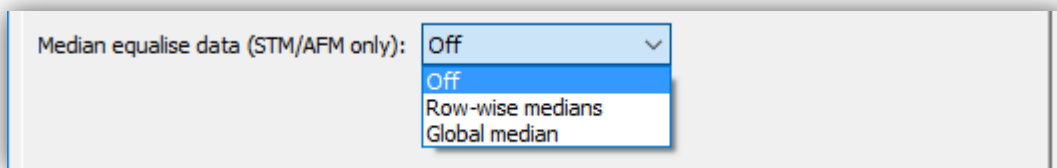
If either orthogonal type option is used, then the user should select whether the fast-scan direction is horizontal or vertical in the first image of the series *after* any trimming of the series and/or transposing has been applied from the earlier options.

Scan-line Time

Here the user can specify the time per fast-scan line. This is only used to calibrate the frequency axis in the instability frequency analysis.

Median Equalisation (Tip Height Correction)

This is an advanced option which attempts to correct for STM/AFM tip height errors.



For STM/AFM image series exhibiting tip-height errors, the data can be adjusted in height to bring each scan-row onto either a row-wise median, or global median height.

The default mode for STEM or TEM imaging is 'Off'. The only exception to this would be for STEM data recorded with a cold-FEG gun where the emission fluctuation compensation has failed (but this is very rare).

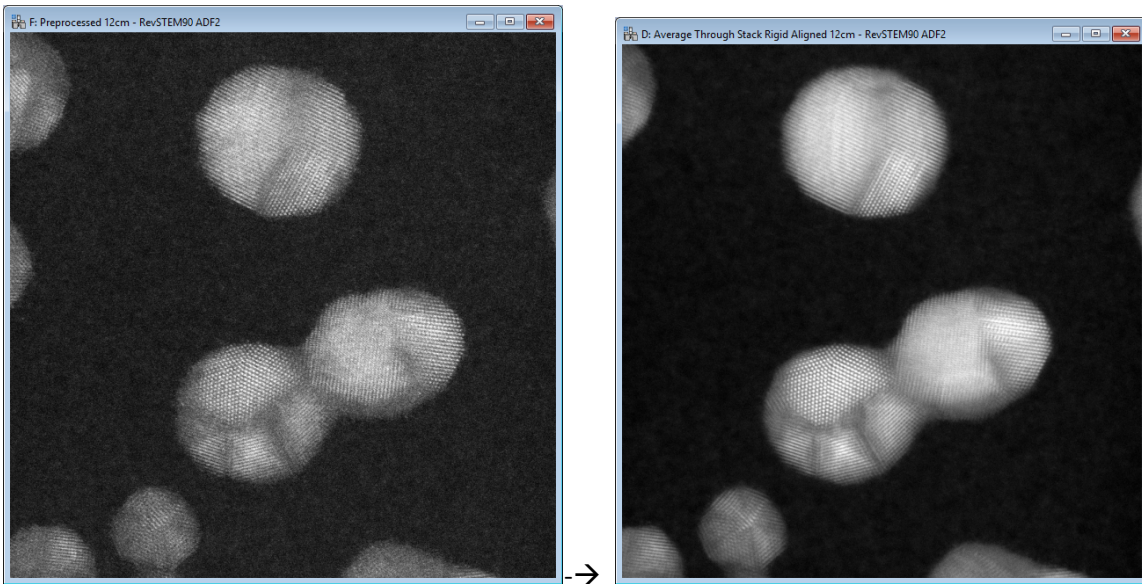
Rigid Registration

Selecting the Area for Rigid Registration

The *Smart Align* software can perform the rigid-registration in one of two ways; either on the entire field of view, or on a selected 'region of interest' (ROI).

Full field-of-view (without ROI):

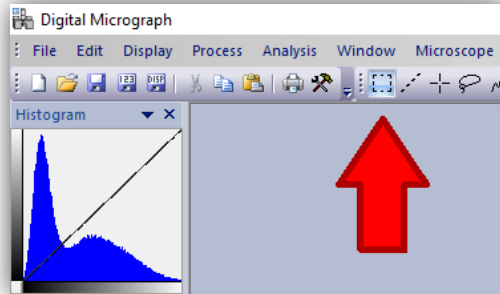
To use the full field of view, simply select the pre-processed data as the front window and then select "Rigid registration" from the *Smart Align* menu.



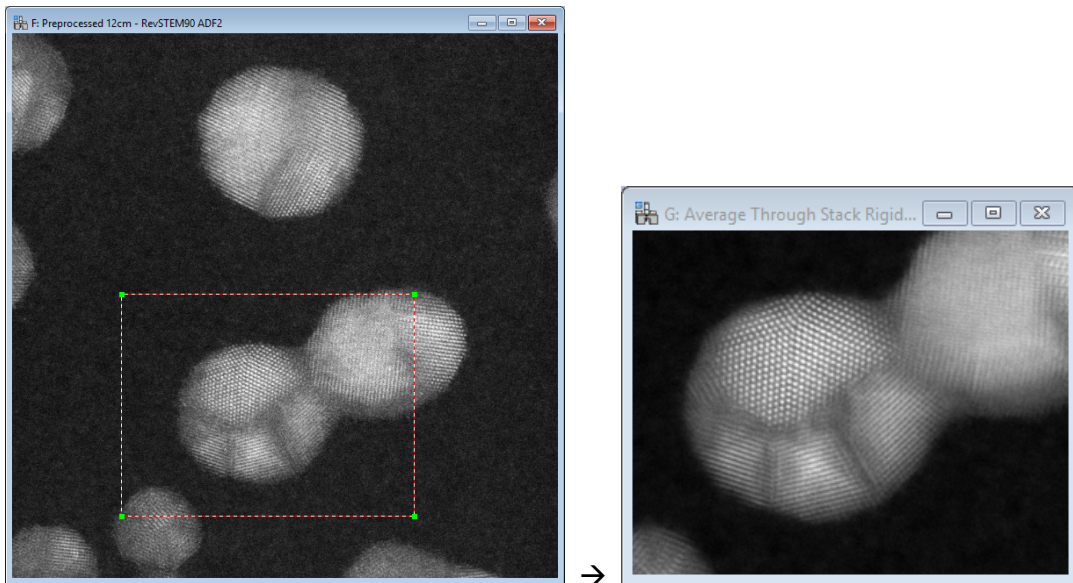
Using the entire image will yield an aligned image with the biggest possible field of view, but it may take longer to compute. For low-dose data, where crystal fills the field-of-view, then using the whole area is likely to increase alignment precision.

ROI Alignment

To use the ROI alignment mode, first draw a ROI on the pre-processed data using the DM ROI tool:



You should ensure that the ROI is large enough to fully encompass the feature of interest in ALL slices of the series data, not just in the first slice. You should inspect the slices before proceeding to ensure that the ROI is large enough to include the feature after sample/stage drift.

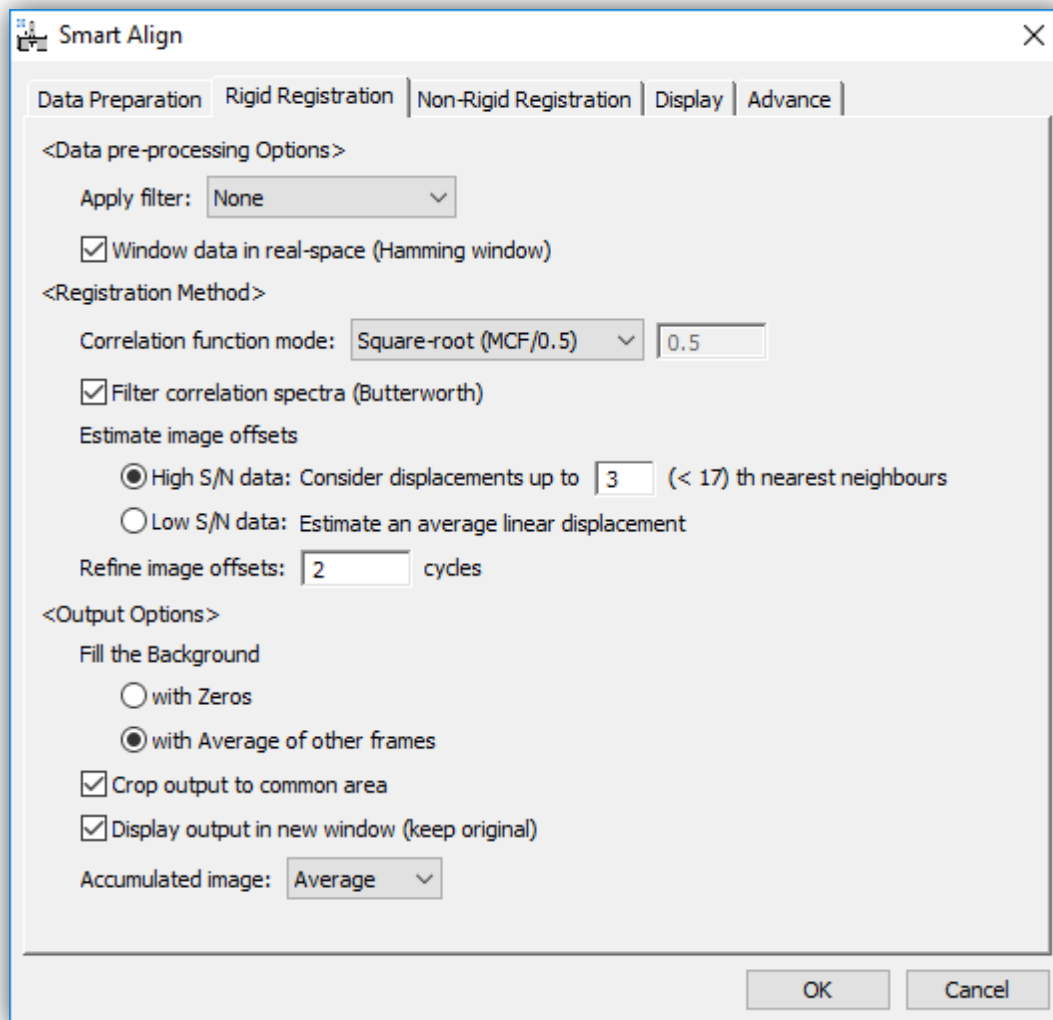


Where only one region of an image is of interest, then using a ROI rigid-align will generally be faster. This may also be useful if trying to align very large datasets where memory use would be an issue and a fiducial mark is present (see special FIB-SEM example at end of manual).

A hybrid of these approaches is to measure the image shifts on the smaller ROI data but then to correct the image shifts on the full field-of-view; see the later section on "Aligning Extra Image Signals".

Selecting the Rigid Registration Options

Each experimental dataset is unique; the rigid-registration menu contains options to adjust the settings to get the best possible image alignment for any given dataset.



The rigid-registration menu is divided into three sections:

Data pre-processing Options

Data Filtering

If the image-series is noisy then image filtering may improve the reliability of the registration. Two options are available; 3x3 median filtering and 5x5 Weiner filtering. Other advanced filtering may be used and this special case is discussed with a special example at the end of this manual.

Data Windowing

For images which have crystal filling the whole field of view, cross shapes often appear centered on the Fourier spots due to the abrupt edges of the image. These can affect the reliability of the image correlations and windowing the images in real-space may improve this. A Hamming window can be selected here.

Registration Method

Rigid registration is achieved in SmartAlign using correlation functions. Three specifically named correlation functions (between two images g and h) are shown below:

The cyclical cross-correlation:

$$\text{Corr}(g, h) = F^{-1} (F(g) \circ F(h)^*) \quad (1)$$

The mutual-correlation function (MCF) (Heel, Schatz, & Orlova, 1992):

$$\text{MCF}(g, h) = F^{-1} \left(\frac{F(g) \circ F(h)^*}{\sqrt{|F(g) \circ F(h)^*|}} \right) \quad (2)$$

The phase-correlation function (PCF) (Kuglin & Hines, 1975):

$$\text{PCF}(g, h) = F^{-1} \left(\frac{F(g) \circ F(h)^*}{|F(g) \circ F(h)^*|} \right) \quad (3)$$

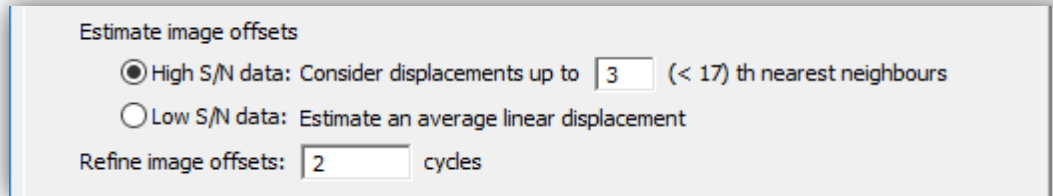
More generally these can be expressed as a continuum set where the exponent on the normalising denominator, α , is a variable between zero and one (Kuglin & Hines, 1975). Hereafter, for generality this will be referred to as the normalised correlation function (NCF):

$$\text{NCF}(g, h) = F^{-1} \left(\frac{F(g) \circ F(h)^*}{|F(g) \circ F(h)^*|^\alpha} \right) \quad (4)$$

The need to consider modifying the power on the normalising denominator lies in the range of specimens that may be needed to analysed. For example, specimens with very 'bland' crystallography, such as a constant thickness slab of Si[110], will require the registration to be achieved entirely using the non-periodic information so a larger value of α would be recommended (say 0.75). However, a free standing metallic nanoparticle free from contamination can be registered very accurately considering the silhouette of this periodic information, so a lower value of α would be most suitable (say 0.2).

Finally, if both the images being compared are noisy, the resultant correlation functions can themselves be noisy and an additional option to filter these with a Butterworth filter is provided.

Once the user has selected which filtering and correlation function options will be used, there are two further choices; how the offsets will be estimated and whether the offsets will be refined:



Estimating the Image Offsets

Here there are two modes, optimized for both high signal-to-noise (SNR) data and low-SNR data. For the majority of datasets, the high-SNR mode can be used, and only for the most noisy data is it necessary to use the low-SNR mode.

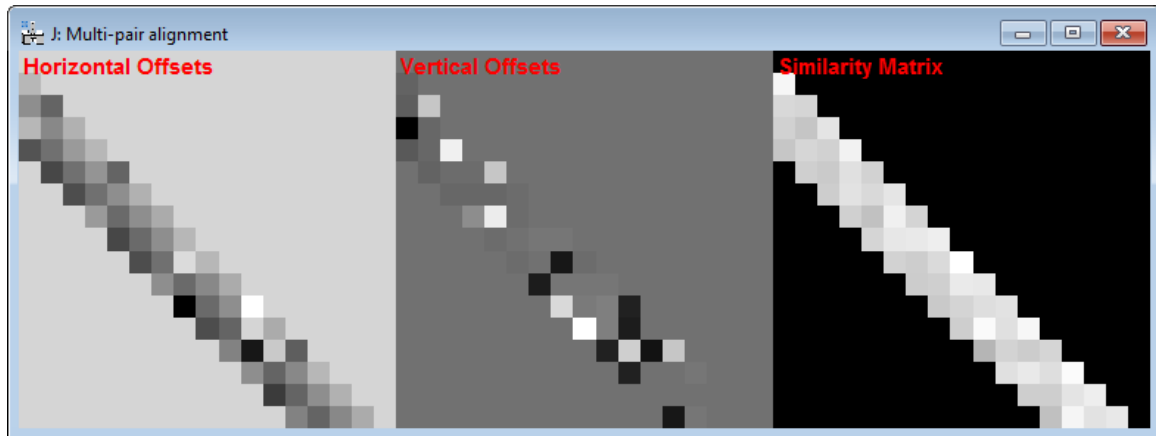
High SNR Mode:

In the high-SNR mode it is assumed that the correlation functions setup with the previous options will estimate the offset between images accurately but potentially imprecisely. To mitigate this the rigid-registration stage of the SmartAlign plug-in does not have rely on any one measurement to determine the offsets. Instead, when determining the image offsets, each image can be analysed relative to its direct neighbours in the image series, then its second nearest neighbours, third-nearest etc.

Here the user can select the number of nearest neighbours to use; the minimum is one, in which case each image is only compared with the previous one, up to (n-1), where n is the number of images in the series. To help the user the menu will indicate the maximum number that can be selected alongside the option.

Increasing this number excessively will increase calculation time with diminishing gains in precision. Additionally, in the presence of gradual structural changes, such as beam-damage, including ever further afar neighbours may reduce registration precision. A typical selection would be three or four nearest neighbours.

The example below shows the result of computing the offsets for the first 4 nearest neighbours.



The figure shows three plots the horizontal offsets between the images, the vertical offsets, and the similarity between image pairs. Each value is plotted in the lower half (triangle) of the matrix representing the i^{th} - j^{th} images. The similarity matrix can be useful in determining the number of nearest neighbours to include, if the similarity begins to drop, this indicates that the images are no longer similar enough and would not improve (or possibly worsen) the precision of the offset estimation.

Once the matrix of neighbour data is compiled a mutual solution is determined.

Low SNR Mode:

In “Low SNR Mode” the image series data will first be condensed into four subsets to improve the SNR of low-dose images. The sample drift between image 1 and image 5, should be approximately the same as between image 2 and image 6, and so on. This gives redundant observations of the same drift-rate and allows for a more precise estimate to be achieved for noisy images.

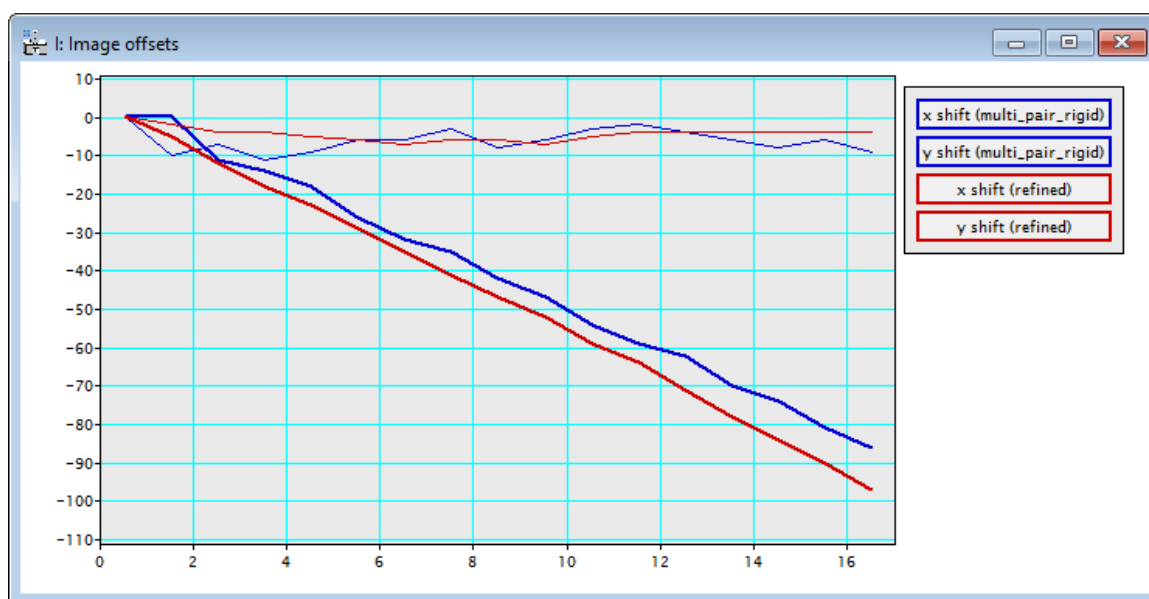
After the four subsets have been analysed to determine the overall drift-rate, the offsets between these corrected four sets of data are analysed to correct for any shifts induced by any 90° image rotations used.

Together these two information describe all the shifts, the images are realigned and can be refined as described in the following section.

Importantly, this mode is only applicable to image-series recorded in a single acquisition. If the series is generated from joining together two or more shorter series, then the prior knowledge of the monotonic drift behaviour may not be valid and this method will not be suitable.

Refining the Image Offsets

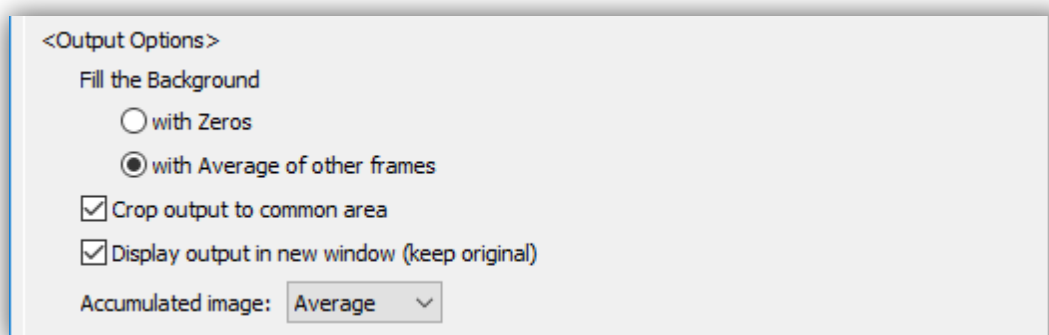
Using the image-offsets estimated from the previous stage the image series can be realigned. The average of this realigned series will have a higher SNR than any one image alone. Next, the user can reanalyse each individual image relative to this average image, and the offsets may be able to be improved. This procedure is somewhat iterative and the user can decide how many attempts should be made to refine the offset estimates.



A typical value might be two refinements; selecting zero will skip this stage.

Output Options

Once the image offsets are refined a new window with the results will be output.



Where shifts exist between images there are two choices to output the data; firstly, the canvas of the data can be expanded to incorporate the image drifts (results in a bigger size image series), or the data can be cropped to only keep the physical area common to all frames (results in a

smaller image series). The user can select which option they want with the “Crop output to common area” checkbox.

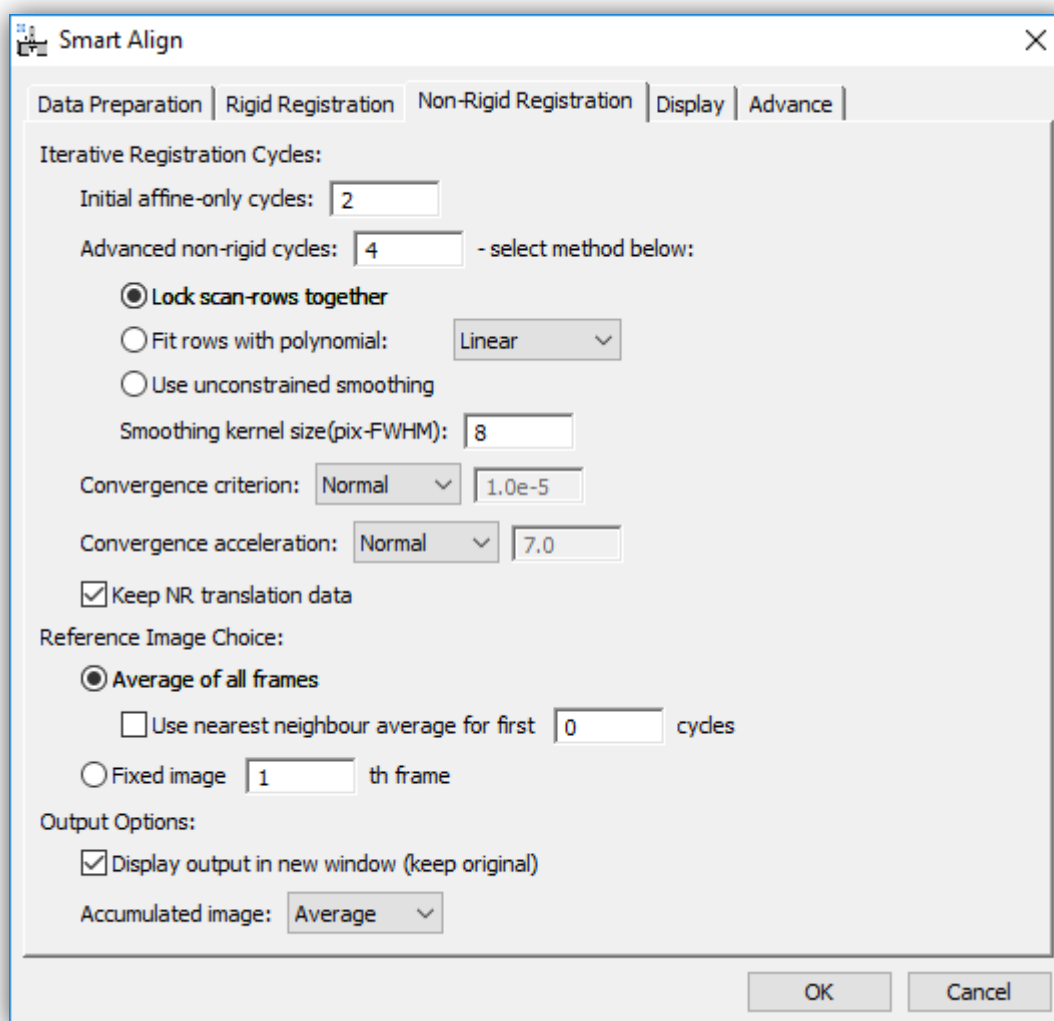
If the user selects not to crop the data, they will expand the canvas; in this case there is an option of what to fill the extra area with. This can be filled with blank (zeros) or it can be filled with the average of all non-zero observations from that position. The former option is useful for viewing the image drift and observing changes in the illuminated region of the sample; the latter is useful for a more ‘panoramic’ option, attempting to maximise the field of view.

In addition to the image series being output, the accumulated image is also shown after processing. The final option in this section determines whether the average of the frames is shown, or the summation.

These rigid-registered results can be saved at this point, it is not necessary to continue to the non-rigid calculation, as this can be resumed at a later time.

Non-rigid Registration

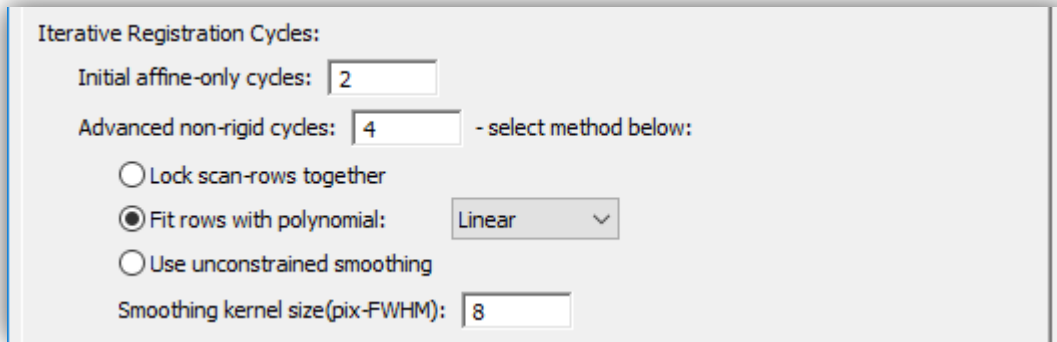
After finishing the rigid-registration, non-rigid registration can be performed. Before beginning the non-rigid registration, if it is not already, the results from the rigid registration should be selected to bring it up as the front window.



This menu is divided into three sections, with options related to the iterative registration cycles, the choice of reference image to use, and for the outputting of the result.

The Types of Iterative Registration Cycles

Since the *SmartAlign* algorithm does not enforce any prior knowledge on the expected symmetry or appearance of the sample, the registration runs iteratively, converging towards the ground truth. These iterations are broken into two types and can be selected separately or together.



The first type is to correct the scan-distortion using only affine transformation. This type of transformation is a 'linear non-rigid transformation'. Linear, because any straight line feature before correction will be preserved as a straight line afterwards, and non-rigid because different pixels can move by different amounts in x and y. Using this type of iteration the user can correct for image shear, scale, and (small) rotations.

As every pixel in the field of view is used when calculating the affine transform for every frame (6 parameters), this type of correction is very robust to image noise.

The next type is the "Advanced non-rigid cycles"; these are the fully non-linear non-rigid iterations. 'Non-linear', because now straight lines are no longer required to be preserved. Unsurprisingly, non-linear transformation is necessary to correct non-linear distortions.

The Number of Iterative Cycles

In the majority of cases few affine type cycles will be required. However, if scan-rotation between frames is enabled or if the drift rate of the stage is varying over time then more may be necessary. Fortunately, these affine calculations are computationally very efficient so a default of say 2-cycles is a good starting value to try in most cases.

The full non-linear calculation is more computationally intensive. For quick calculations to check the quality of data two cycles may be sufficient. For a more final analysis a typical number might be four cycles.

Types of Non-linear Cycle

The *SmartAlign* algorithm offers three types of non-linear iterative cycle. The first of these is the so-called 'scan-locked' mode. This is described in more detail in the *SmartAlign* academic manuscript. In brief, this mode maximises the artefact robustness of the algorithm by utilising

the prior knowledge that the STEM is a scanned instrument. This mode is also applicable for STM/ADF/SEM etc. For serial microscopes this is generally the most appropriate mode.

In some cases, where the sample is rotating slightly in the plane (not tilting) then row-locking will not describe this additional motion. For this case 'row-fitting' should be used. This allows for the benefits of row-locking, but also allows an additional degree of freedom to describe the rotation. A typical use case for this mode might be nanoparticles loosely attached to a substrate.

The final mode available is the fully unconstrained mode. This is most suitable for parallel acquired data such as conventional TEM and photographic data.

Smoothing Kernel Size

After selecting the type of non-linear cycle, the user must specify a smoothing kernel size.

There are two reasons why a smoothing kernel is used to constrain the diagnosed transformation fields. Firstly; it is not physically realistic for the diagnosed transformation fields to contain abrupt discontinuities and so a small smoothing is justified to remove noise.

Perhaps more importantly, and mostly when using the unconstrained (TEM) mode, it is important that no crystallography is observed in the transformation fields. Remember, the transformation fields are only a description of the environmental distortions affecting the data. As the distortions do not depend on the sample under observation, then these should not exhibit signs of crystal periodicity.

Remember adding a small scan-rotation, and using row-locking will improve STEM (row-locked) series processing.

Registration Convergence - the Target, Acceleration & Number of Cycles

As the *SmartAlign* software solves for the image distortions iteratively, success is determined by the measurement convergence. Three inter-related factors contribute to this convergence behaviour.

The Convergence Target

The convergence target specifies the exit criterion for the innermost iterative loop. It has units of the average magnitude of pixel shift in the frame. The preset options range from weak-convergence (average shift of 10^{-4} pixels) to a strong criterion of 10^{-6} pixels. The user can also specify a custom number.

The Convergence Acceleration

The convergence acceleration determines the aggressiveness of the iterative correction. A higher number will increase calculation speed but if the number is too high, and the data are noisy, then the convergence may become unstable. In this case a warning message will be displayed and the software will attempt to 'self-heal' by reducing the acceleration by 10%.

The weak setting yields an acceleration of 2 through to a value of 20 in the strong setting. The user can also specify a custom value.

The Total Number of Iterative Cycles

Adding together the number of affine-only and non-linear registration cycles yields the total number of iterative cycles. The final convergence criterion target specified above is not reached immediately, but rather the strictness of the criterion is approached gradually and reaches the ultimate target value at the final cycle.

Storing the Transformation Data

If more than one image signal was acquired, then the probe-offset data can be transferred from one diagnosis to correct other signals. If this is the case, then the user should check the 'Keep NR translation data' box. When this translation data is stored it is stored in the image metadata for later use, the data can be used to compensate a second signal immediately, or the file with its metadata can be saved and the translations used at a later date. In this case the file size will be larger by approximately a factor of three (the original image data, and two new metadata with equal dimensions describing the shifts in x and y).

Choice of Reference Image

The user can choose which reference data is used as the comparison image.

Average of all Frames

For the majority of cases the best reference data will be the average of all frames in the series. In this mode each image will be compared with the average when diagnosing the shifts. At the end of each cycle the updated frames will be used to make a new average reference image.

Fixed Frame

An alternative to using the average frame is to specify a particular frame to use as the reference for all others.

Local Average for Initial Loops

A third choice, for special cases, is to use only the average of the images immediately before and after each image in the series as the reference – a local average. This local average can be useful if the image distortions are so large that averaging through the whole series would cause an excessive degradation in reference image quality.

An alternative use for this local average reference method is the processing of dynamic movies such as *in-situ* data, where the underlying sample is genuinely transforming and an average of all frames would not have a physical basis. In this case some distortion correction is still possible while maintaining temporal resolution.

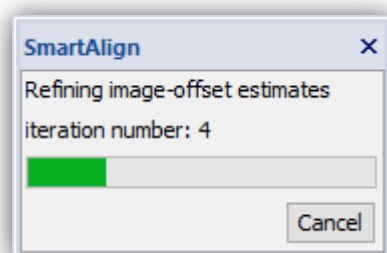
Output Options

Here the user can select to open the results in a new window and keep the original open (the default behaviour). In general, it is only necessary to deselect this in computers with very limited memory or with very large datasets.

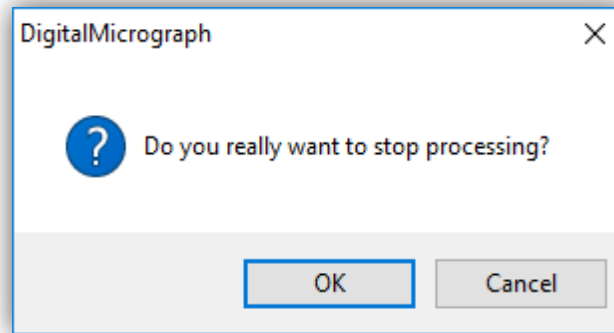
The user can also choose here how to output the compound image after processing. In general, the ‘average’ image is the most suitable (including for fractional ADF or STM/AFM height data), but a ‘summation’ image may be useful if the data represents counts or some other quantity you wish to accumulate.

The Progress Bar

After choosing all the appropriate options, pressing “OK” will start the calculation. During the calculation a progress bar will be shown, as well as some more detailed text based readouts in the Output window.



The calculation can be aborted if necessary by pressing the “Cancel” button. In this case a confirmation dialog will be displayed:



In this case only the currently running non-rigid calculation will be aborted. The existing rigid-registered result will remain in the workspace. The NR calculation can be started again, or the rigid results saved at this point and the calculation continued later.

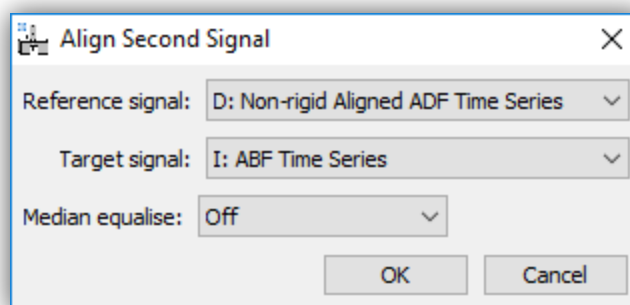
The NR Outputs and Image Cropping

After the NR registration is complete, and depending on the severity of the drift and distortions, pixels shifted from outside the scanned area will appear around the edge of the frame as grey pixels. It may be necessary to crop one or two pixels from around the perimeter. This can be done by selectin the ROI tool from the toolbar and then selectin the 'Crop' command from the SmartAlign menu.

Aligning Extra Image Signals

After an alignment has been performed (either rigid or non-rigid), then this can be further applied to another image signal so long as the data were acquired simultaneously. If the data were not acquired simultaneously then the drifts and probe distortions will not be equal and this is not applicable.

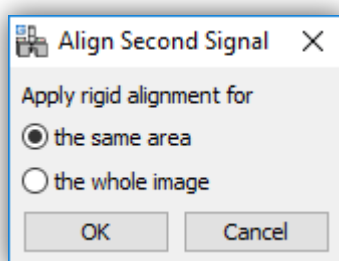
To do this select 'Align Second Signal...' from the main menu.



Here for example the diagnosis from the registration of an ADF series is copied to an ABF series. After pressing OK, both the new aligned signal will be displayed as well the average of the new series.

Rigid Registering Additional Signals

In the case of rigid alignment, if the alignment was performed on a reduced ROI, then the user will be asked whether to use these image shifts for the whole image, or to again return only a reduced ROI:



If the user selects “the same area” then the same area as the ROI used on the first signal is returned, if the user selects “the whole image” then the whole image is aligned and cropped to the largest possible common area.

Non-rigid Registering Additional Signals

In the case of NR registration, it is only possible to return the same field-of view, and the above choice will not be presented.

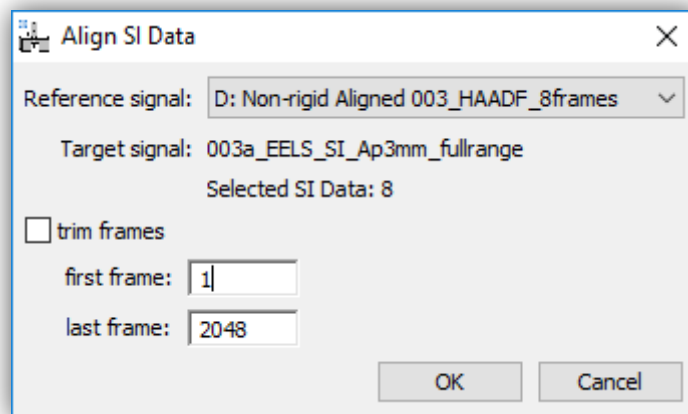
Aligning Spectrum Image Series

If the image series which has been aligned represent the simultaneous images frames accompanying spectrum image acquisition then the probe offsets can be copied from the image series to the series of spectral volumes.

First all the steps above should be completed; during rigid alignment cropping should be selected, and during non-rigid registration the option to keep NR translation data should be selected.

With the aligned image series selected as the front window, select 'Align SI Data...' from the *SmartAlign* menu.

A file explorer window will open and you should select *all* the spectral volumes that match the image series. The spectral volume series will be read into DM in the order based on their filenames. For example adding numbers or letters to the SI filenames will help them be read correctly. The example here shows an 8-frame image series and 8 spectrum images selected.



The size of the spectrum images in the third dimension (number of channels) will also be displayed, and if necessary a reduced range can be reconstructed to minimise RAM usage.

You should verify that the correct reference data is selected and press 'OK'.

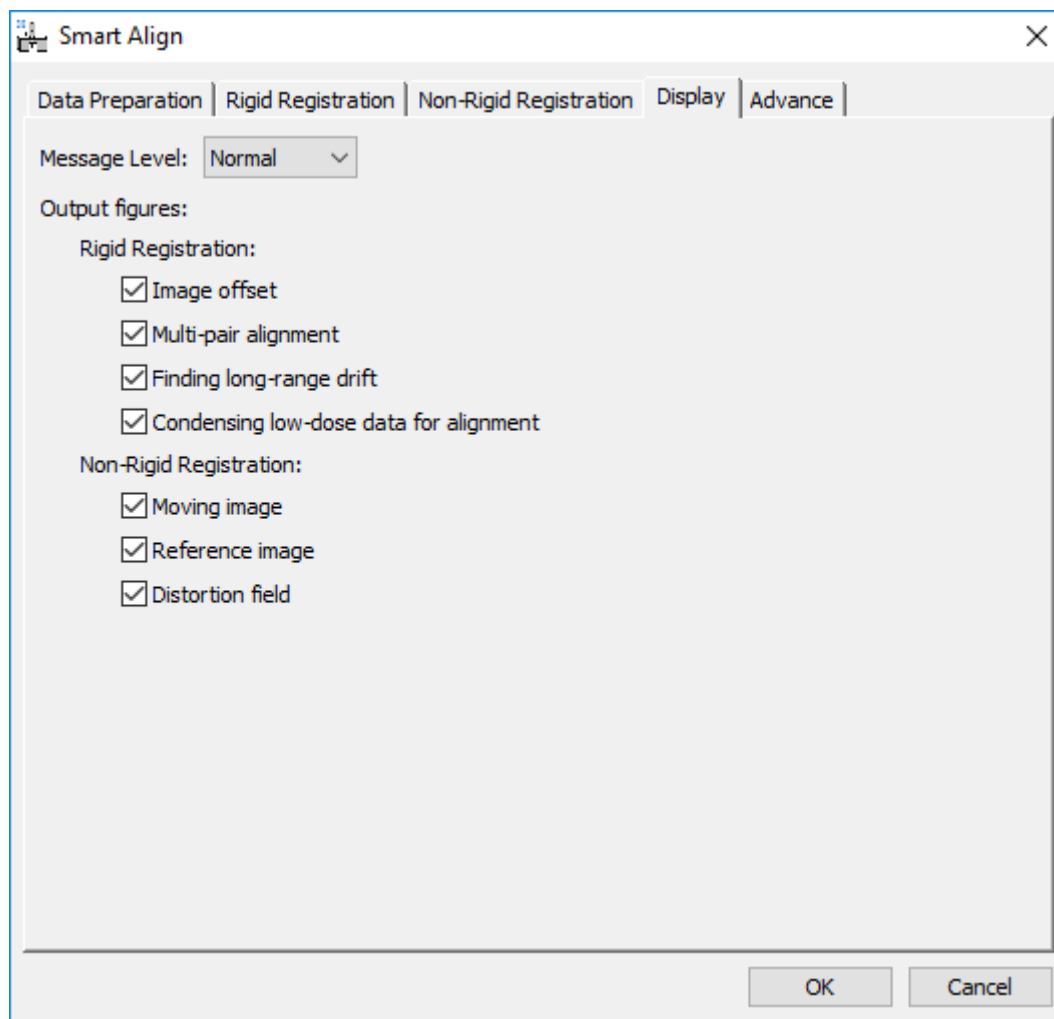
While the SI realignment and NR registration is running, an integrated spectral view will be shown as well as an integrated spatial view. These are not interactive and are just visualisations to monitor the progress. Once complete the corrected spectral volume will be displayed.

Additional Menu Items

There are two further tabs containing output display options and the advanced settings.

Output Display Options

This tab allows the user to display more or fewer additional diagnostic figures during processing.

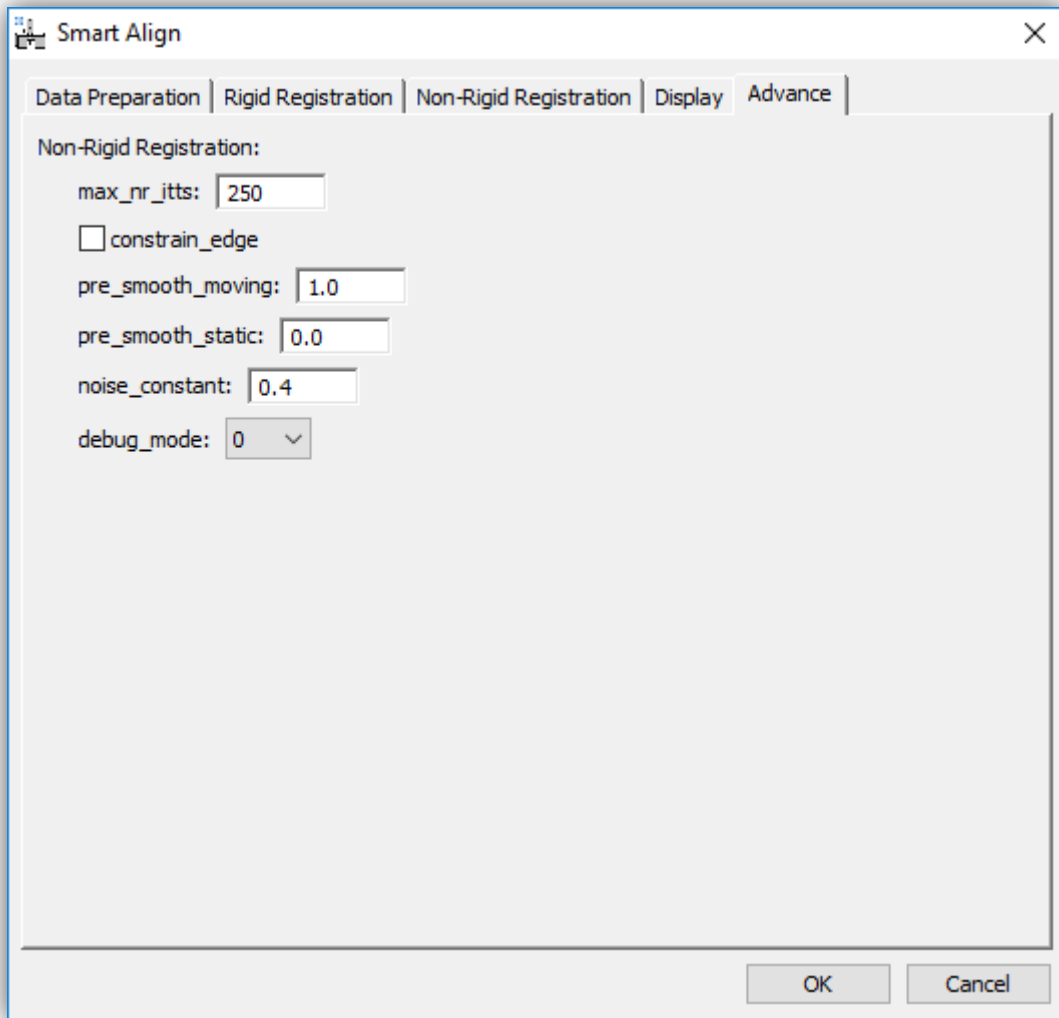


These options include a choice for the granularity of text message outputs (displayed in the DM results window) during processing, as well as the additional graphical display of variables.

These settings are saved when the registration is begun and will be preserved after DM is closed for the next time.

Advanced Options

This menu contains settings that rarely need to be changed.



These options can be used to set an upper limit for the maximum number of iteration per image in the innermost correction loop. Reducing this number may increase calculation speed but if reduced too far may stop the calculation from converging fully.

In the case of an image series with a featureless background, then constraining the edges of the transformation field may improve algorithm convergence.

In the case of very noisy data, a small smoothing can be applied to the individual images or to the average reference image.

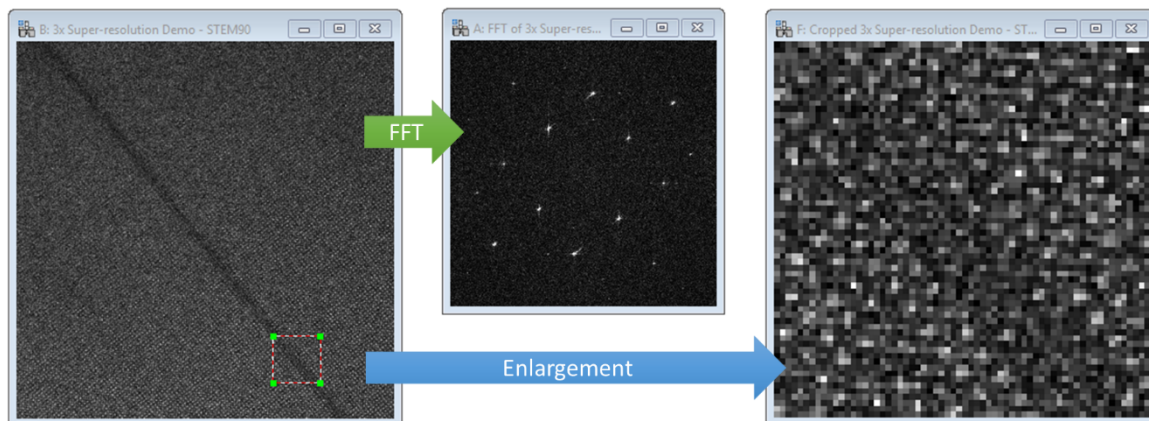
The noise constant controls the robustness to extreme outlier image pixels. A larger number will increase robustness but at the expense of reduced sensitivity. This number should in general not be changed.

The debug mode setting is only used for remote support use.

Special Case Examples

Digital Super-Resolution Example

Consider the example ADF image-series below of a perovskite grain boundary with 16 frames with 512x512 image pixels each. This has been recorded at high-resolution *but* with insufficient pixel-sampling.



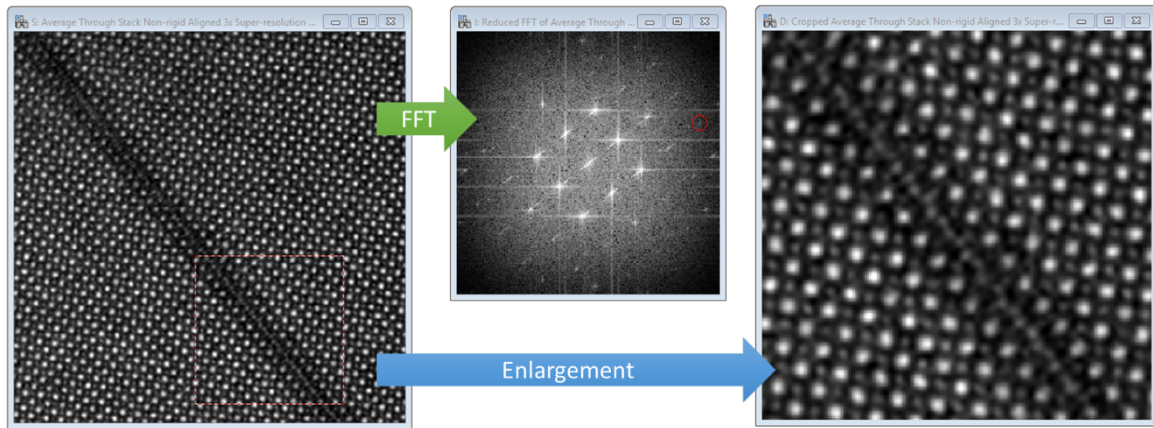
The FFT shows spots are visible right out to the Nyquist limit, and the highest observable spot is the (220) spacing (approx.. 1.38Å. A 64x64 pixel ROI is enlarged on the far right – again, not how the sampling is too low. Ordinarily this would be a poor choice of setting to record the image; we would usually increase the number of pixels or reduce the field of view (increase magnification). However, both these options would increase the total electron-dose over the sample.

With this data, angle-measurement, structure matching, peak-finding or strain-mapping would be very challenging. To overcome this, we will increase the sampling of this frame-series in post-processing using digital super-resolution.

Note: this is not *optical super-resolution*, it will not extend the resolution beyond that of the original STEM probe, but rather this is *digital super-resolution* where the restored image-series average will contain more *discernible* information than any one frame at the low sampling.

To do this select “Resample data” in the “Prepare data” menu. For this example, we will use 3x. Now the data should be rigid, and non-rigid registered as normal.

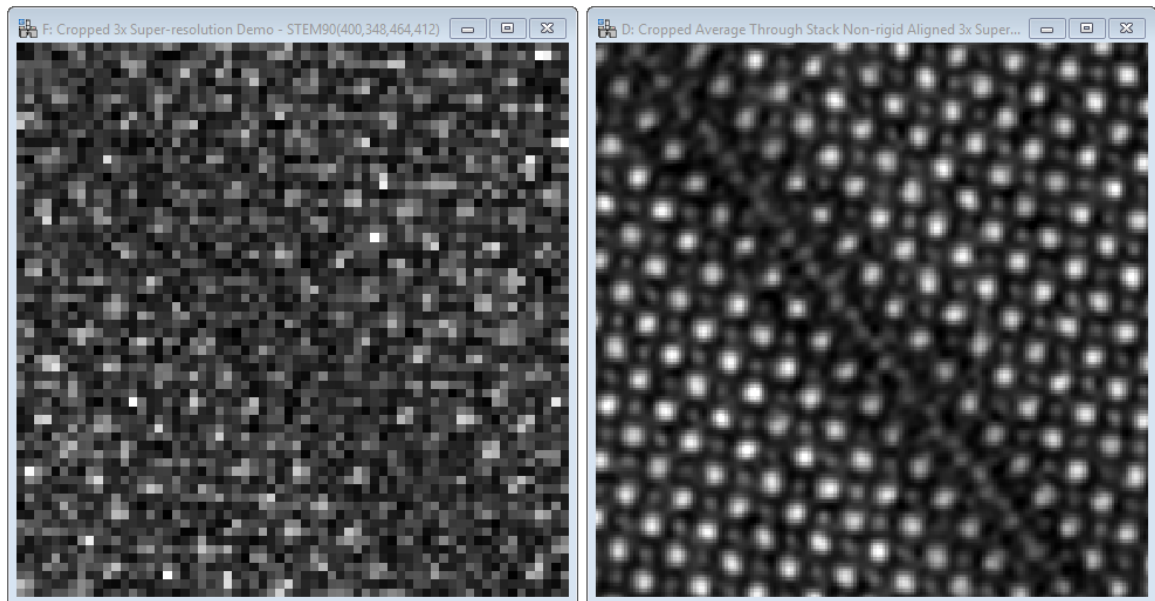
The results below show the lower-right part of the image after resampling and alignment, the reduced Fourier transform (central 50%), and a cropped enlargement of the same region as in the earlier figure above.



After processing the (331) spot is visible (approx. 0.9Å resolution), and the detail in the interface become clear. The recovered resolution is likely slightly poorer than had the image been recorded under optimal sampling conditions, but it is an improvement on the raw data.

Note: the data-resampling must be done *before* the rigid and non-rigid registration to achieve the digital super-resolution. Doing it afterward would just be a simple image scaling.

Below shows further enlargements of the before (single-frame) and after (average-frame) data.

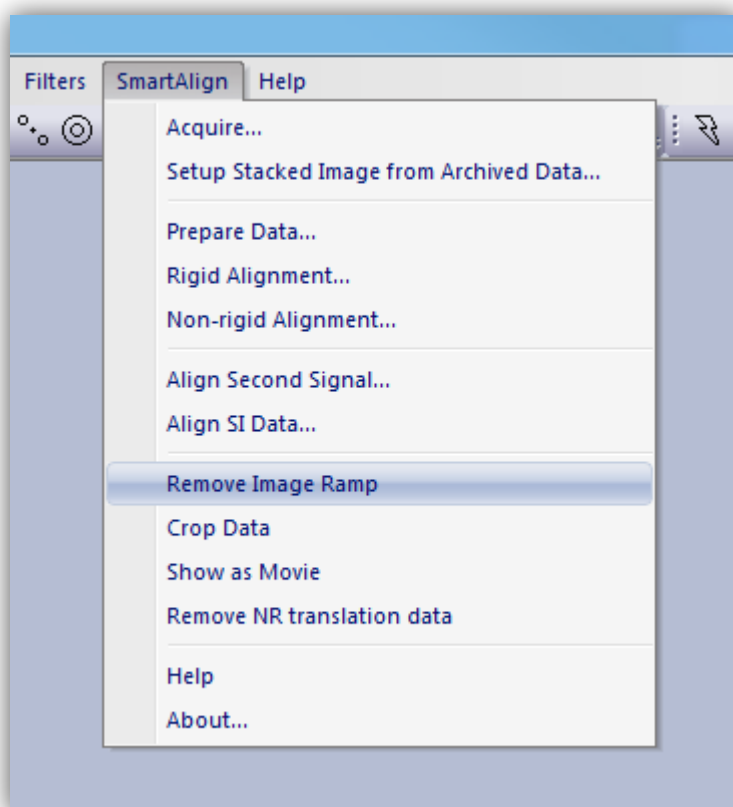


Going beyond 3x data-resampling will likely offer little benefit and will mostly just add to the calculation time.

Processing STM & AFM Data

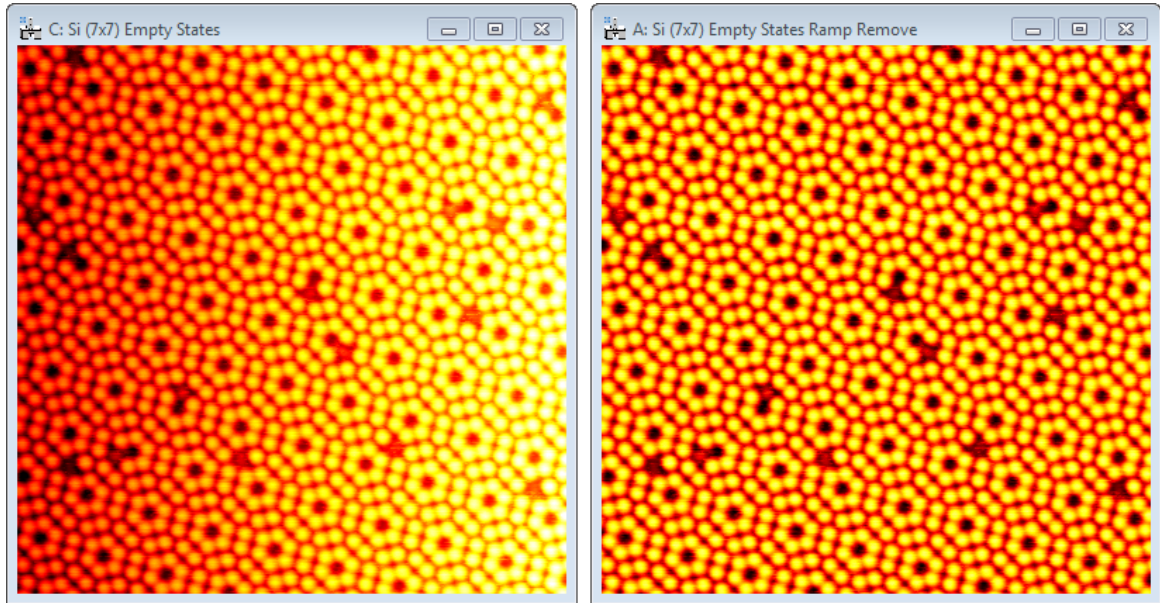
For many scanned probe microscopy (SPM) datasets, it is often necessary to first remove a planar ramp background. This tool is provided as a utility in *SmartAlign* and can be done at any point. However for best results, it is recommended to perform this step at the start of the data processing.

With the data loaded and stacked to create a series, select “Remove Image Ramp” from the main *SmartAlign* menu:



If a single image is selected this will operate once, if a series is selected, every image in the series will have a ramp calculated and removed.

The result will be similar to the example below:



Note: the range (height data) in the image is preserved, and the ramp is subtracted such that the new image minimum is zero.

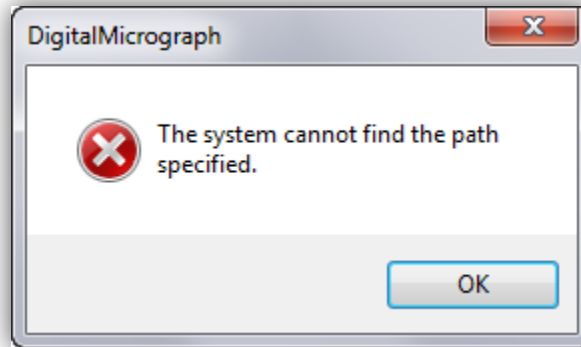
After the ramp removal is complete, you can continue to process the data as usual.

Advanced Filtering for Rigid Registration

Troubleshooting

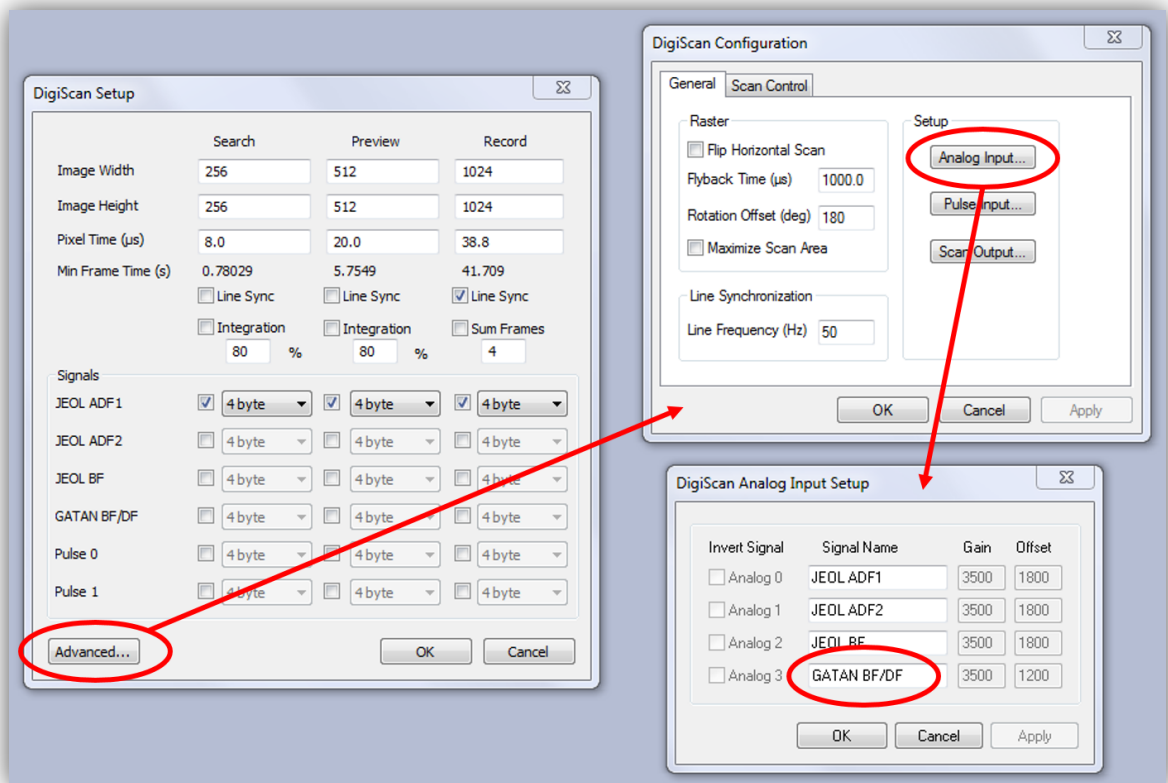
Auto-save function not working with some detector channels.

On some systems you may see the following error when trying to use the auto-saving feature for some detector choice settings:



This is caused when the detector name itself contains a slash ("/") character causing the filename to try to access a non-existent directory.

To correct this, go to the DigiScan Setup menu, select "Advanced" > "Analogue Input" and edit the name of the detector causing the error; in this case "Gatan BF/DF" can be changed, for example, to "Gatan BF-DF".



References

Heel, M. van, Schatz, M., & Orlova, E. (1992). Correlation functions revisited. *Ultramicroscopy*, *46*, 307–316.

Jones, L., & Nellist, P. D. (2013). Identifying and Correcting Scan Noise and Drift in the Scanning Transmission Electron Microscope. *Microscopy and Microanalysis*, *19*(04), 1050–1060. <http://doi.org/10.1017/S1431927613001402>

Jones, L., Yang, H., Pennycook, T. J., Marshall, M. S. J., Van Aert, S., Browning, N. D., ... Nellist, P. D. (2015). Smart Align—a new tool for robust non-rigid registration of scanning microscope data. *Advanced Structural and Chemical Imaging*, *1*(1), 8. <http://doi.org/10.1186/s40679-015-0008-4>

Kuglin, C. D., & Hines, D. C. (1975). The phase correlation image alignment method. *Proc. Int. Conf. on Cybernetics and Society*, *1975*, 163–165.