

July, 2026

SmartAlign *for DigitalMicrograph*

Image Registration and Scan-distortion
Compensation Plug-in

SmartAlign Manual 4.0 (Provisional)
with
Spectrum Image Acquisition
Smart Search

and

Template Matching Module
4D-STEM Module

HREM Research Inc.

Preface

This user manual describes how to install and operate the *SmartAlign* for STEM image/spectrum-image/4D-STEM registration and scan-distortion compensation plug-in (version 3.0) for DigitalMicrograph. The software described builds on the work from three manuscripts relating to:

1) image series registration and averaging:

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](#)

2) multi-frame spectrum-image registration and summation:

Jones, L et al.: “Managing dose-, damage- and data-rates in multi-frame spectrum-imaging”. *Microscopy*, **67** (2018). [Link](#).

3) multi-frame 4D-STEM dataset registration and summing:

O’Leary et al.: “Increasing Spatial Fidelity and SNR of 4D-STEM Using Multi-Frame Data Fusion”. *Microscopy & Microanalysis*. 28 (4), (2022). [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(s) such as:

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

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 tunnelling 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 manuscripts referenced above, please contact Dr Lewys Jones at:

lewys.jones@tcd.ie

Copyright Statements

© Copyright 2015-2026 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	7
Software Requirements	7
Software Installation	7
Typical Workflow	8
Plug-in Menu	9
Supported File Types	10
Data Acquisition	11
'Control' Tab	11
'Signal' Tab	14
SmartSearch	15
Data Preparation	18
Data Shaping Options	18
Data Scaling Operations	19
Scan Metadata Options	19
Rigid Registration	21
Selecting the Area for Rigid Registration	21
Selecting the Rigid Registration Options	23
Data pre-processing Options	23
Registration Method	24
Output Options	27
Non-rigid Registration	29
The Types of Iterative Registration Cycles	29
The Number of Iterative Cycles	30
Types of Non-linear Cycle	30
Smoothing Kernel Size	31
Registration Convergence - the Target, Acceleration & Number of Cycles	31
Storing the Transformation Data	32

Choice of Reference Image	32
Output Options	33
The Progress Bar	33
The NR Outputs and Image Cropping	34
Aligning Extra Image Signals	35
Repeat Alignment of Similar Data-sets	36
Description of Options	37
Aligning Spectrum Image Series	38
Aligning 4D-STEM Data Series	40
Additional Menu Items	42
Output Display Options	42
Advanced Options	43
Spectrum-Image Acquisition Menu	45
Special Case Examples	51
Digital Super-Resolution Example	51
Processing STM & AFM Data	53
Advanced Filtering for Rigid Registration	55
Template Matching Module for SmartAlign	56
Introduction	56
Template Matching across a Series of Images	57
Name convention for a series of images	57
Pre-treating the Data	58
Defining the Motif	58
Detecting Occurrences of the Motif	62
Non-rigid Registration of the Extracted Tiles	65
Template Matching a Second Signal / Spectrum Image	66
Troubleshooting	69
Auto-save function not working with some detector channels	69
References	70

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 (or multi-frame spectrum-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:

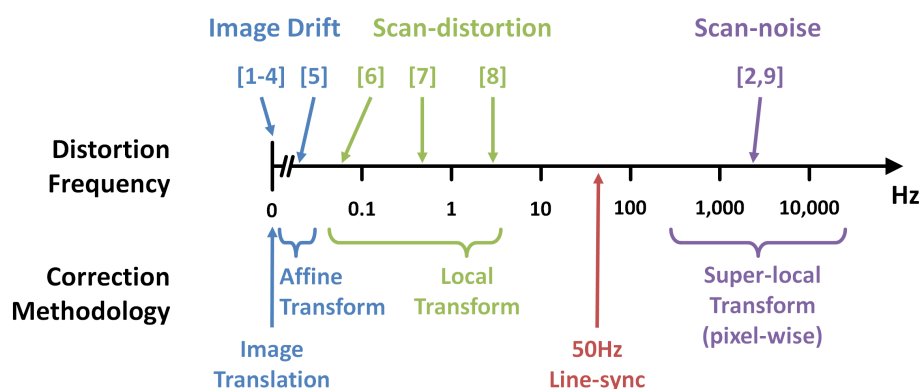
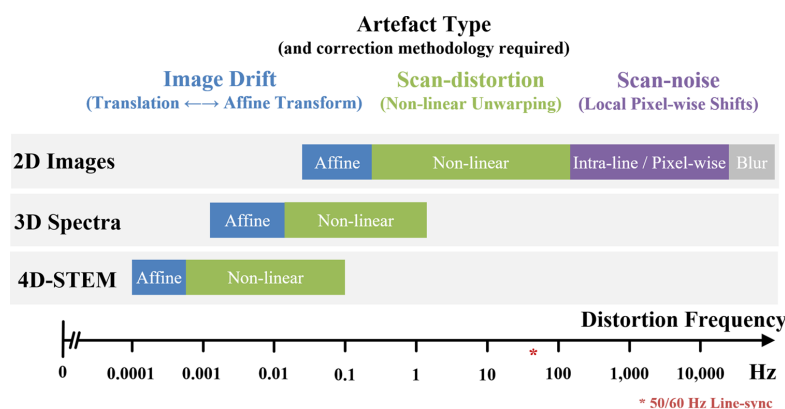


Image distortion frequency domains. For a description of the references above, please see the associated academic manuscript [1].

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 [2].

For spectrum-images and 4D-STEM, where pixel times may be 100 to 10,000 times slower than basic STEM imaging, the frequency ranges that may disturb experimental data are moved lower commensurately.



Examples of possible frequency ranges of image/data distortions depending on typical pixel times used in each experiment type [3].

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)

If you don't have Gatan SI Picker Tool, the HREM Picker Tool is necessary to run the *4D-STEM* module:

- HREM Picker Tool Plug-in (Free-ware downloadable from www.hremresearch.com)

For the spectrum-image acquisition options, this also requires:

- DigitalMicrograph (GATAN™) Spectrum Image (SI) base module, EELS module, EDS module.

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 HREM Mouse Tools, IPU and Acquire Image Series Plug-ins

All the files relating these tool plug-ins, can be installed by drag-and-drop copy to the folder "PlugIns."

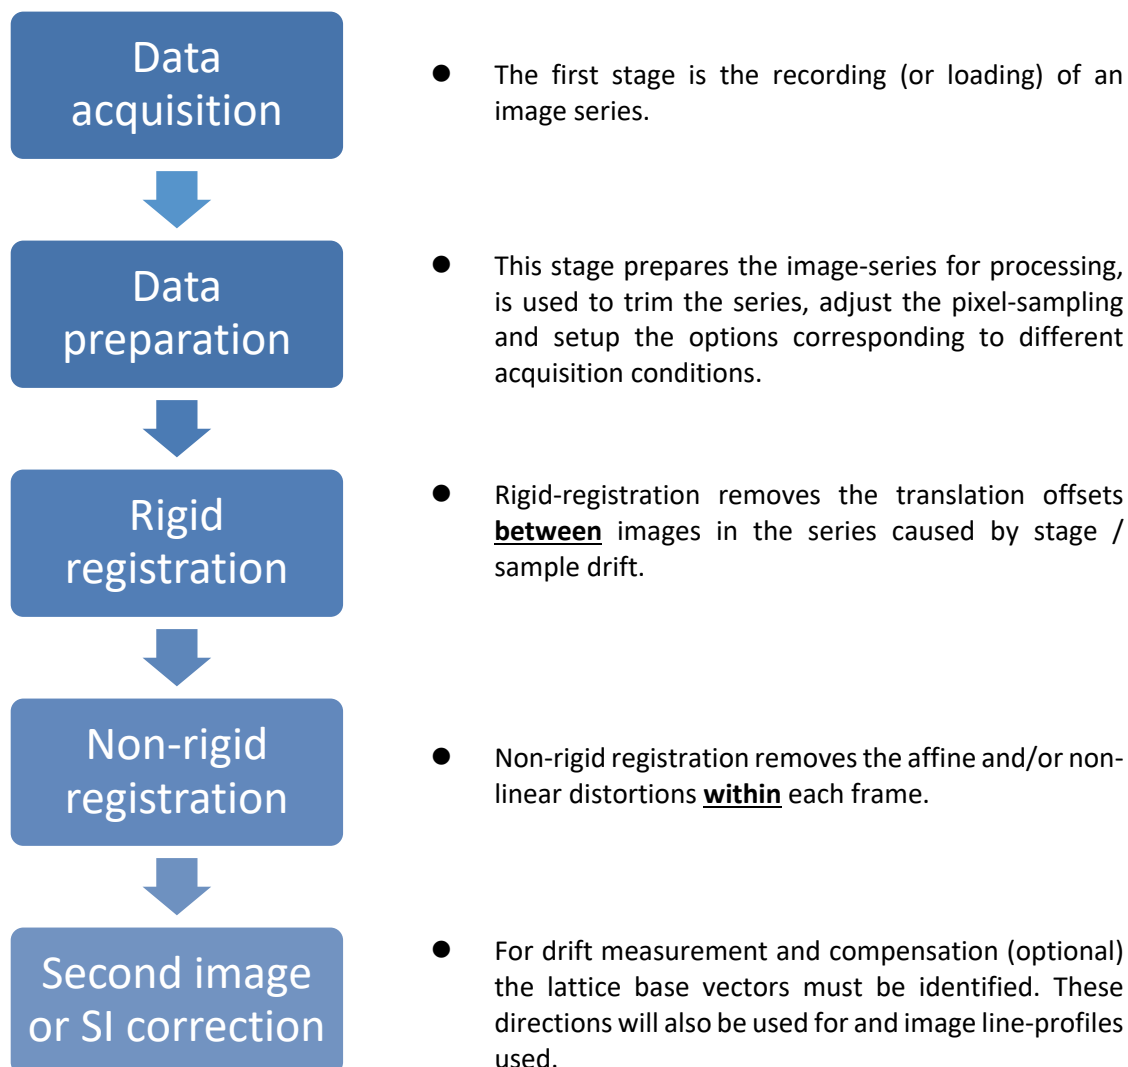
Installing Template Matching Module and/or 4D-STEM Module

These modules can be installed by drag-and-drop copy to the folder "PlugIns". *Please note that these modules require additional licenses.*

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.

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.

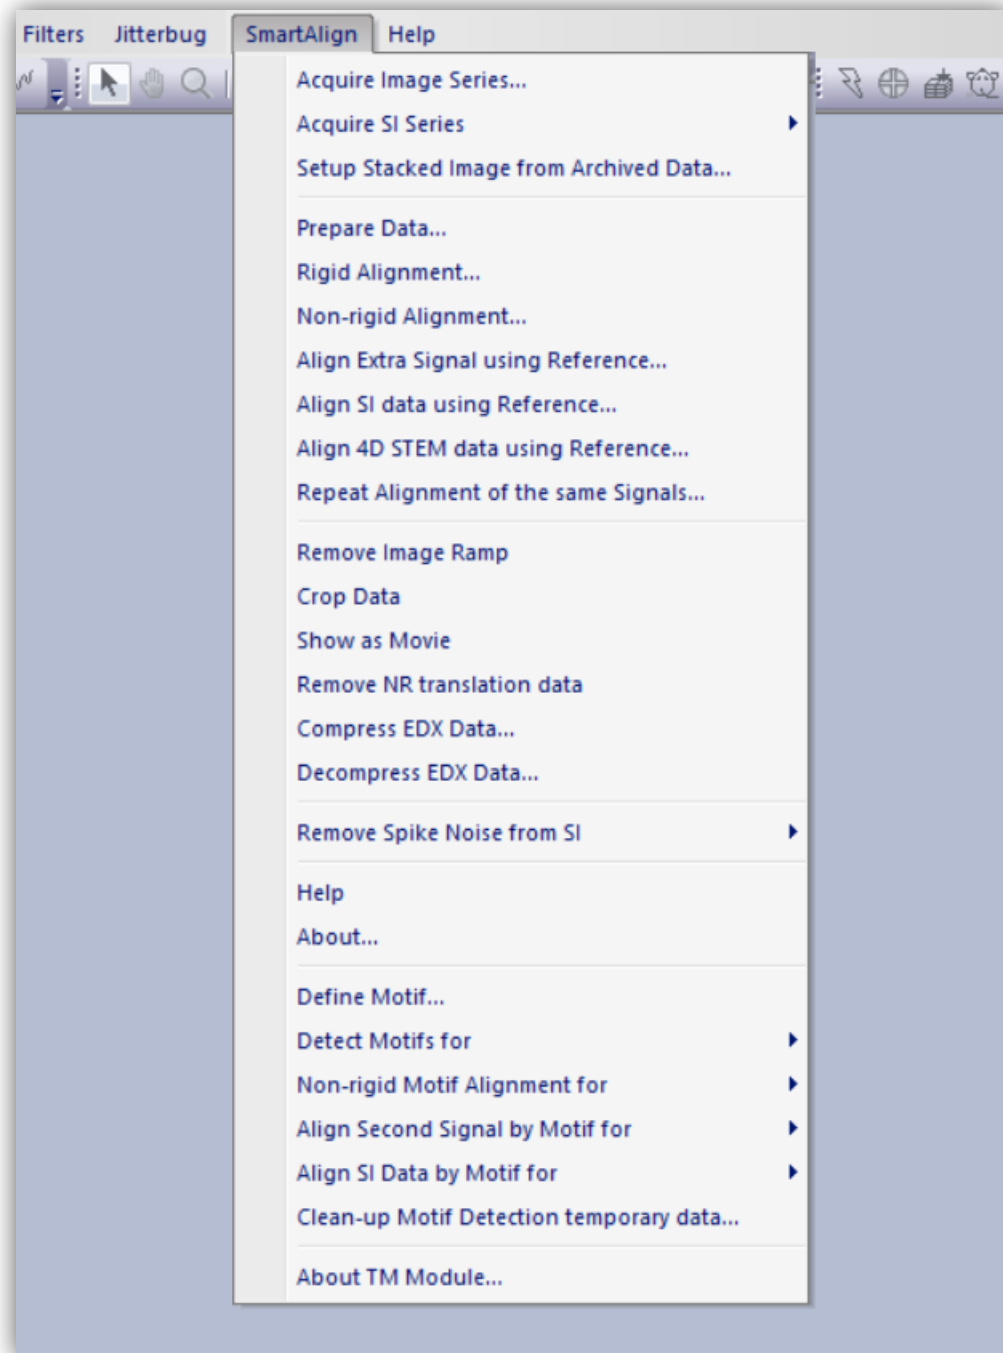


The plug-in is designed so that each stage in the processing creates a new data window for its results. If for some reason the calculation fails, or the user wants to go back one stage and choose a different setting, the new window can simply be closed and the result from the previous stage will still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



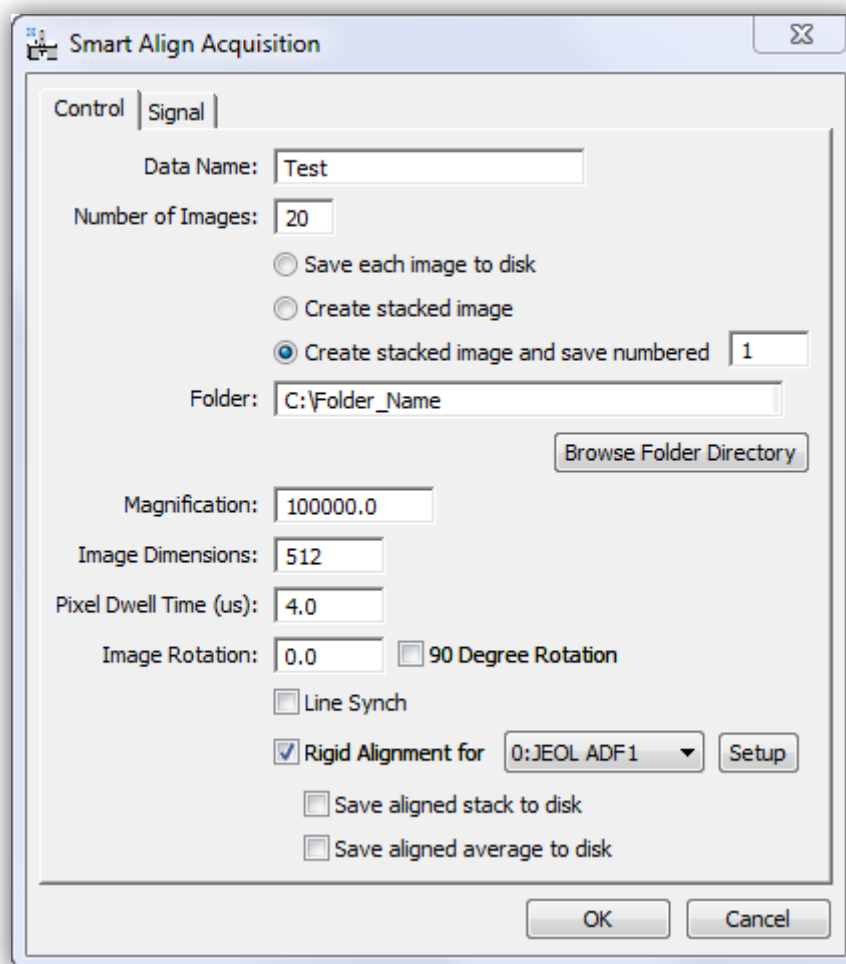
Supported File Types

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.

For a detailed discussion of how to optimise the number of frames in an experimental acquisition for a fixed electron-dose budget, see the manuscript [4]

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.

Rigid Alignment for

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.

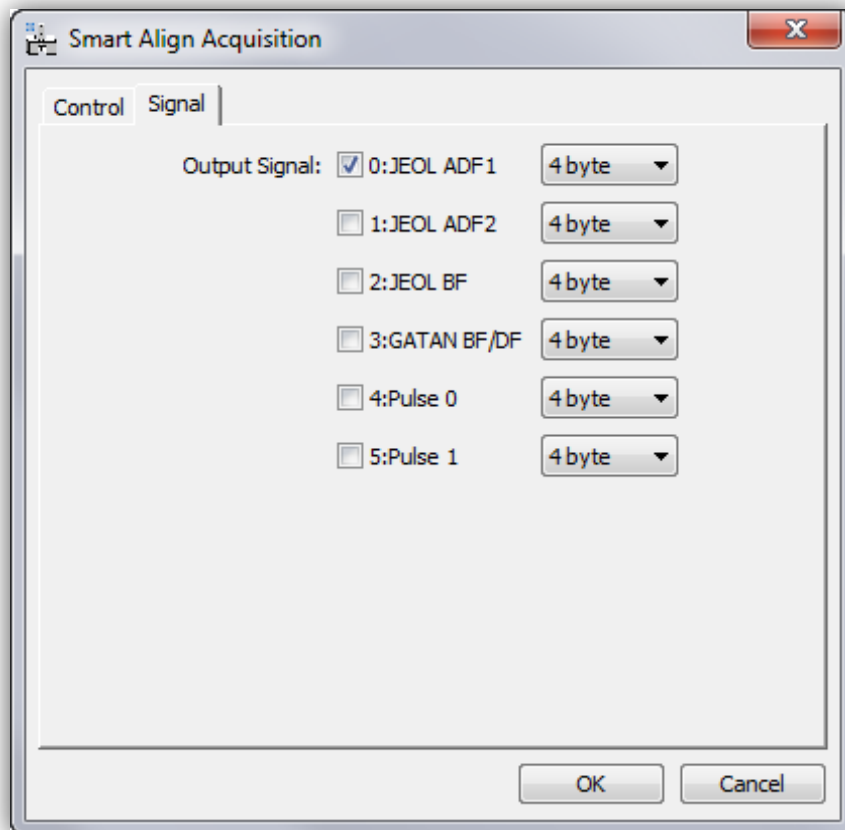
If this option is selected, then one signal can be selected as the inspection signal. Other signals can be registered later. The 'Setup' button allows for the rigid-registration settings to be adjusted.

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.

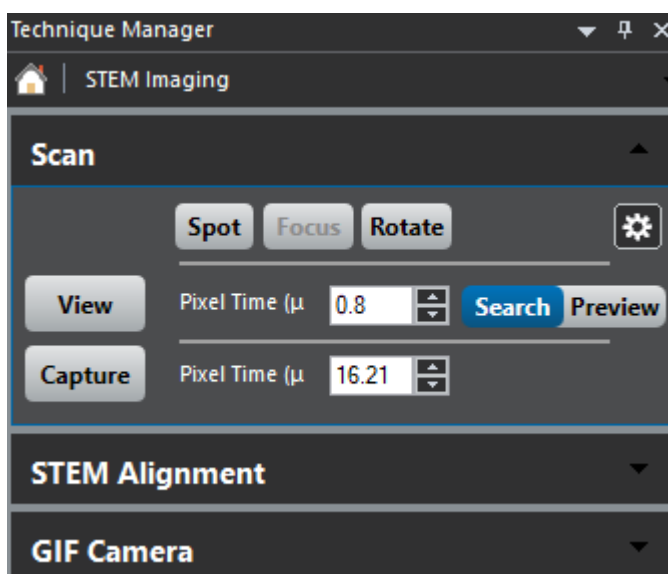


SmartSearch

SmartSearch performs the real-time Rolling Average of the View mode image(s) to improve good signal-to-noise ratio, while automatically records raw data for a specified period of time into the Frame Buffer for later off-line processing. This feature is especially useful when image quality is poor due to low electron dose. **Note:** Rolling average and raw data recording are performed for all signals displayed in View mode.

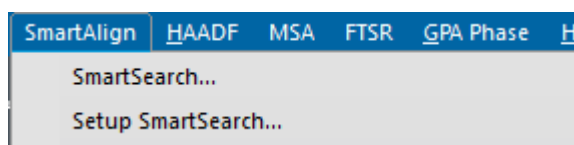
The recorded raw data can be used for rigid and non-rigid alignment of SmartAlign. Furthermore, the automatic raw data recording is especially useful for samples where beam damage is a concern. This is because it makes possible to cut out the undamaged portion from the saved data after data acquisition.

Before using SmartSearch, select the STEM Imaging Technique using the Technique Manager, and adjust the Pixel time in View mode of the Scan palette. If possible, select the shortest available time.



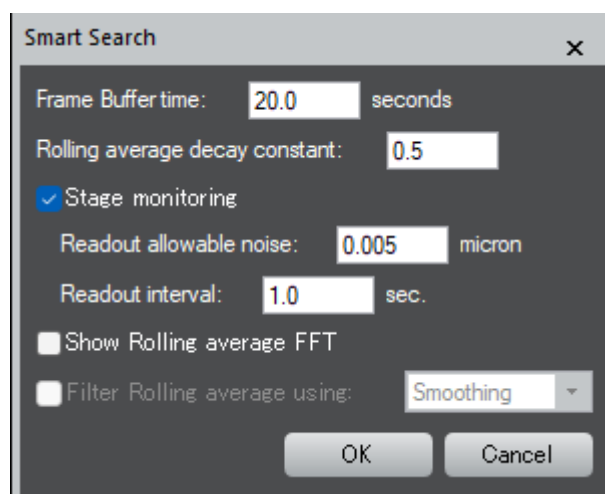
Additionally, please select the signals to be displayed in View mode using the "Scan" settings (the gear icon). For all selected signals, the Rolling Average will be displayed and the raw data will be saved in the Frame Buffer.

Selecting SmartSearch from the menu will launch View mode, displaying the Original Image and its Rolling Average for each signal. At the same time, automatic data recording to Frame Buffer will start.



If sufficient observations are made, finish View mode from the Scan palette. Then, each raw data temporarily saved in the Frame Buffer will be displayed as an image stack in the workspace, and you can process it freely.

SmartSearch has the Setup dialog shown below, where you can configure various parameters.



SmartSearch Setup Dialog

Frame Buffer time

This is the time length (in seconds) within which the raw images are saved in the Frame Buffer before View mode stops. Each series of raw images saved in the Frame Buffer will be output as a single stacked image.

Rolling average decay constant

This parameter controls how strongly the Rolling average improves the signal-to-noise ratio (S/N), and how quickly the changes in experimental conditions are manifested by the Rolling average. Lowering this value improves responsiveness, but decreases the signal-to-noise ratio (S/N). Conversely, increasing this value improves the S/N, but decreases responsiveness. (A value of 0 corresponds to no Rolling average, while a value close to 1 means weak responsiveness to the recent images.)

Stage monitoring

If selected, the stage movement is monitored, and when the stage movement is detected, the decay constant is temporarily decreased to improve the responsiveness of the Rolling Average.

Readout allowable noise:

The measured stage position may fluctuate due to noise. Therefore, the range to be ignored as noise is specified here in micron.

Readout interval:

This is the time interval for monitoring the stage movement.

Show Rolling average FFT

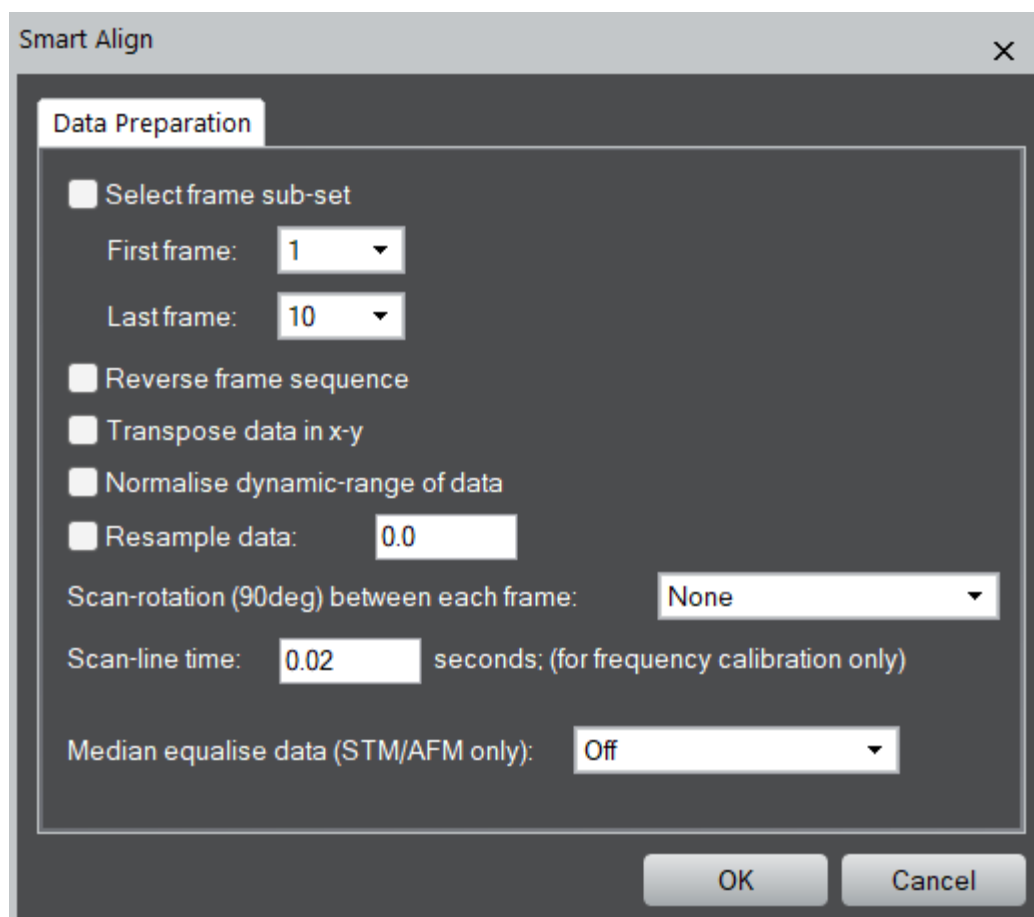
If selected, the FFT of the Rolling average is also displayed for each signal.

Filter Rolling average using

If checked, the noise filter used for Rolling average can be select from this list. This feature will be implemented in future.

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

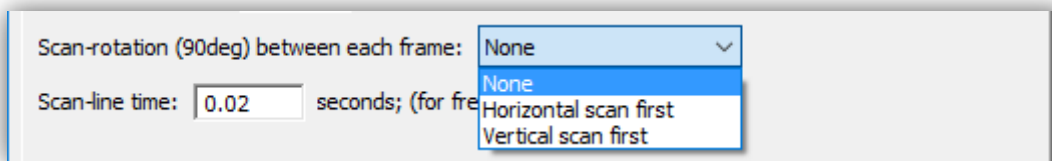
This option allows the user to resample the data in x-y before any onward processing. Resamplings smaller than one will reduce the number of pixels in the image (more coarse sampling), and values larger than one will increase the number of pixels in the data (more fine sampling). Reducing the sampling 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.

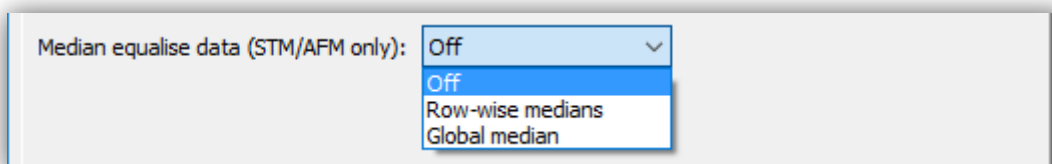
If either orthogonal type option is used, then the user should select whether the fast-scan direction is horizontal or vertical in the first image of the series **after** any trimming of the series and/or transposing has been applied from the earlier options.

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).

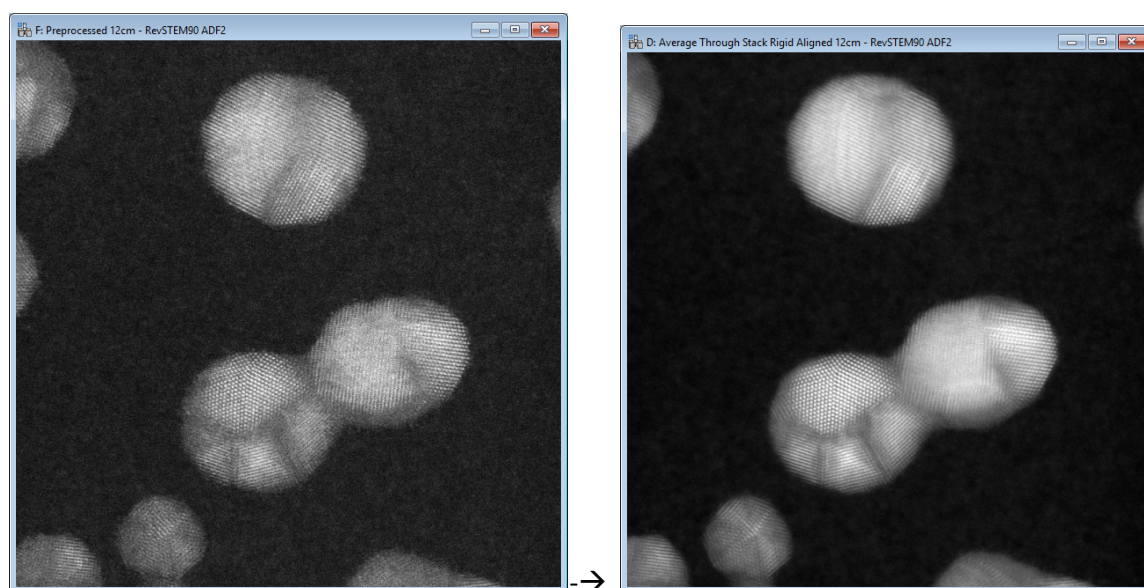
Rigid Registration

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):

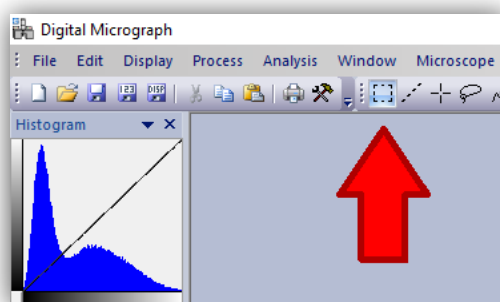
To use the full field of view, simply select the pre-processed data as the front window and then select "Rigid registration" from the *Smart Align* menu.



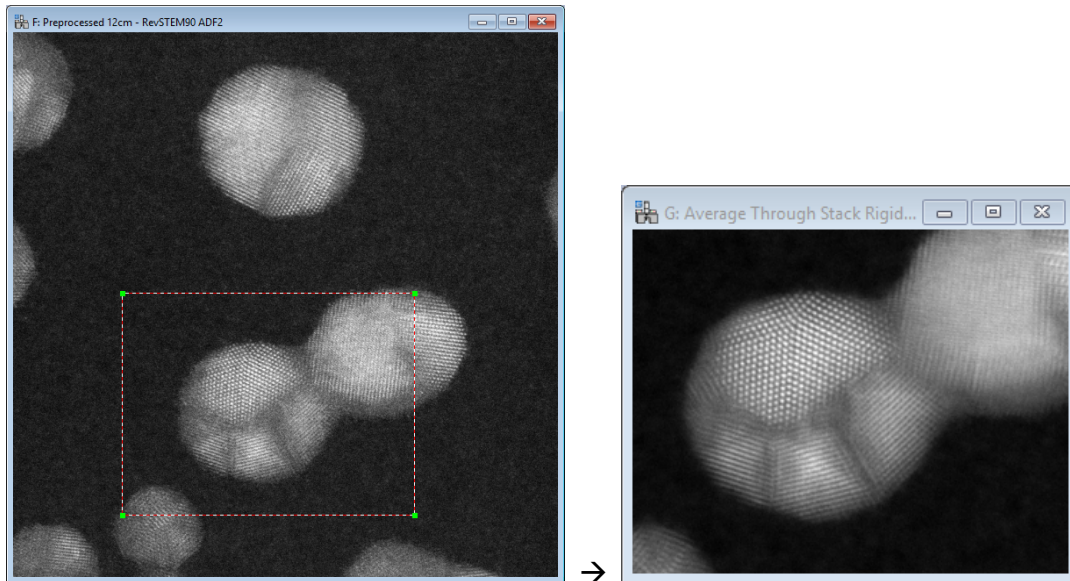
Using the entire image will yield an aligned image with the biggest possible field of view, but it may take longer to compute. For low-dose data, where crystal fills the field-of-view, then using the whole area is likely to increase alignment precision.

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.

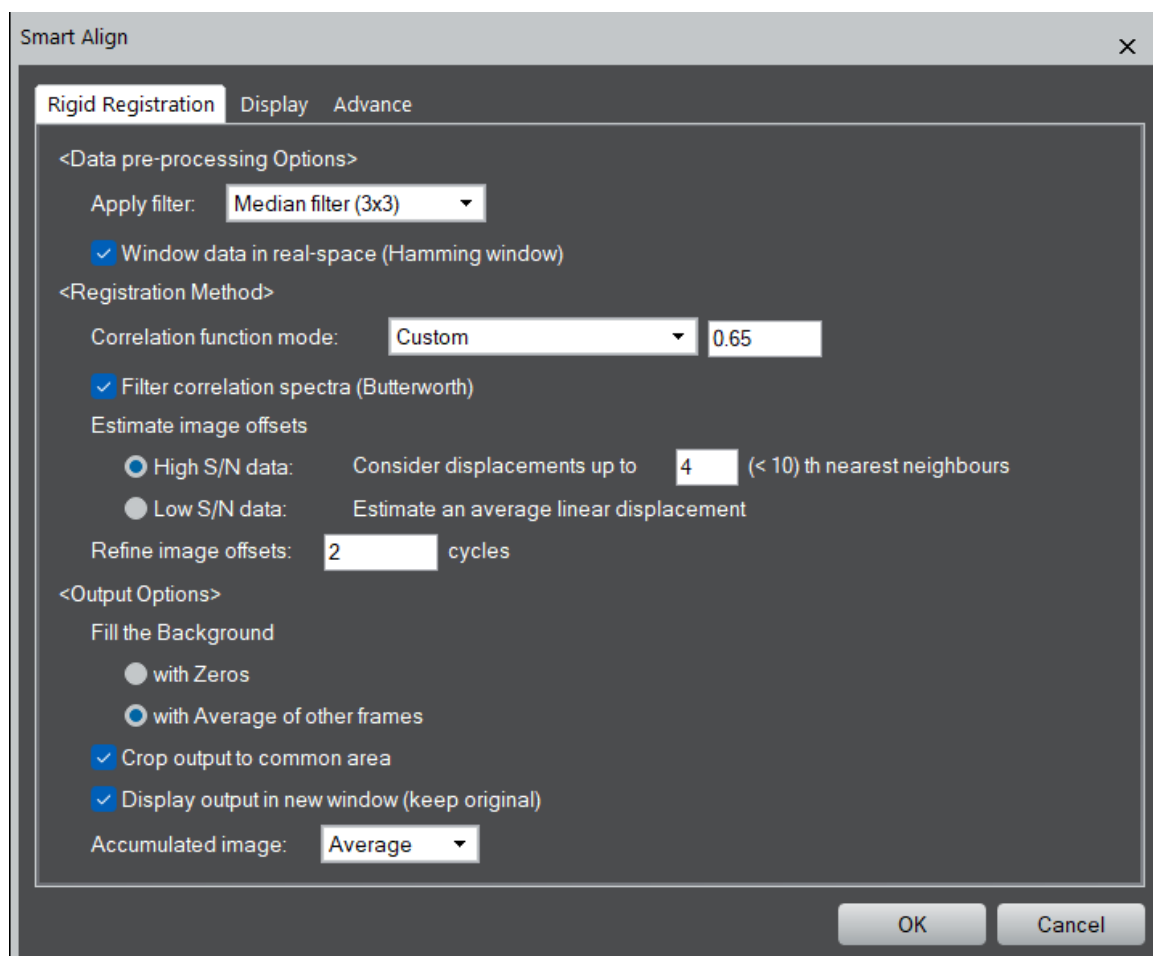


Where only one region of an image is of interest, then using a ROI rigid-align will generally be faster. This may also be useful if trying to align very large datasets where memory use would be an issue and a fiducial mark is present (see special FIB-SEM example at end of manual).

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:

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

The mutual-correlation function (MCF) [5]:

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

The phase-correlation function (PCF) [6]:

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

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 [6]. Hereafter, for generality this will be referred to as the normalised correlation function (NCF):

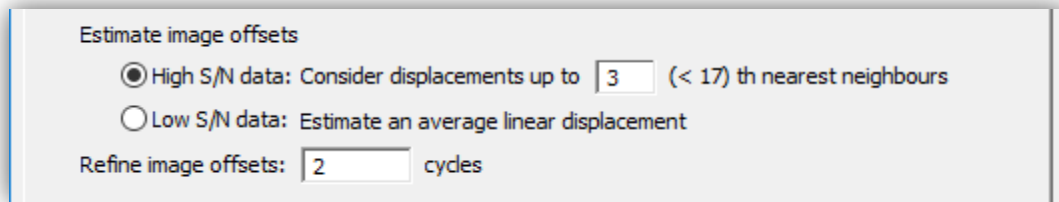
$$NCF(g, h) = F^{-1} \left(\frac{F(g) \circ F(h)^*}{|F(g) \circ F(h)^*|^\alpha} \right) \quad (4)$$

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



Estimate image offsets

☒ High S/N data: Consider displacements up to (< 17) th nearest neighbours

☐ Low S/N data: Estimate an average linear displacement

Refine image offsets: cycles

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 should be used, and only for the most noisy data is it necessary to use the low-SNR mode. If the image contrast is sufficient that the stage drift can be observed by the naked eye, then it is recommended to use the high-SNR mode. The low-SNR mode should only be used where the dose is so low that the stage drifts are not distinguishable by eye.

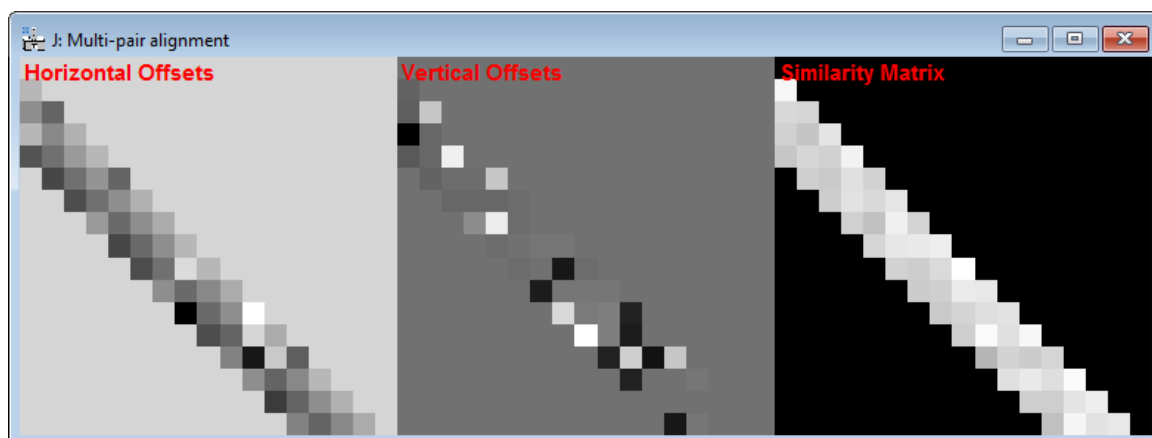
High SNR Mode:

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. If the rate of linear stage-drift is relatively constant, 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.

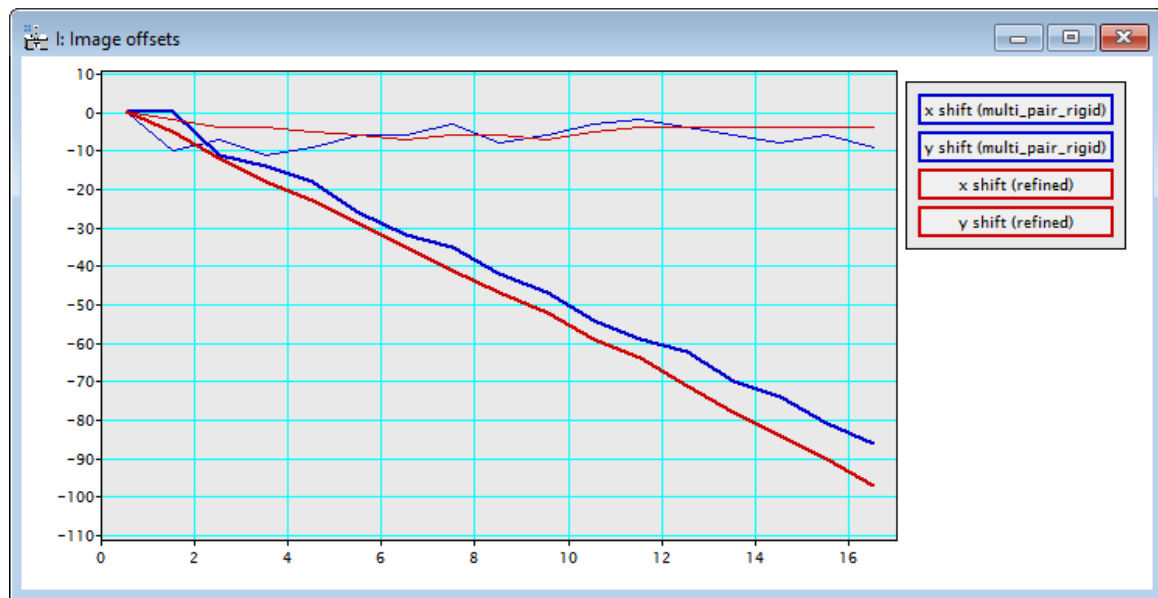
Together these two information describe all the shifts, the images are realigned and can be refined as described in the next section.

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.

It is also recommended to consider windowing the data in real-space (Hamming window option) if you have a very large offset of the scanning rotation origin.

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.

<Output Options>

Fill the Background

☐ with Zeros

☒ with Average of other frames

☒ Crop output to common area

☒ Display output in new window (keep original)

Accumulated image: Average ▾

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.

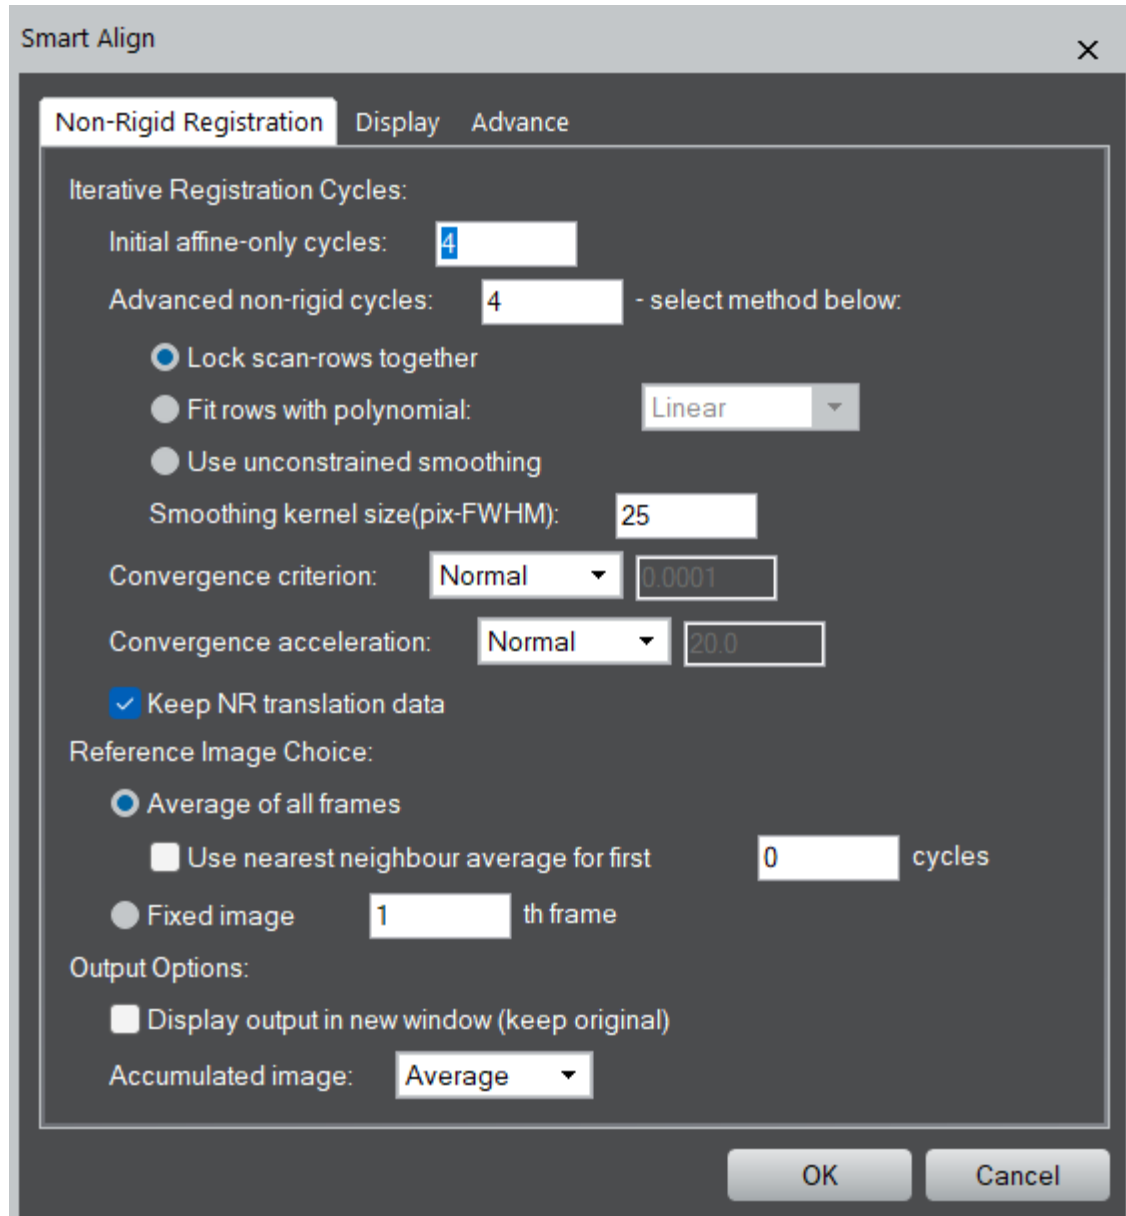
If the user selects not to crop the data, they will expand the canvas; in this case there is an option of what to fill the extra area with. This can be filled with blank (zeros) or it can be filled with the average of all non-zero observations from that position. The former option is useful for viewing 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

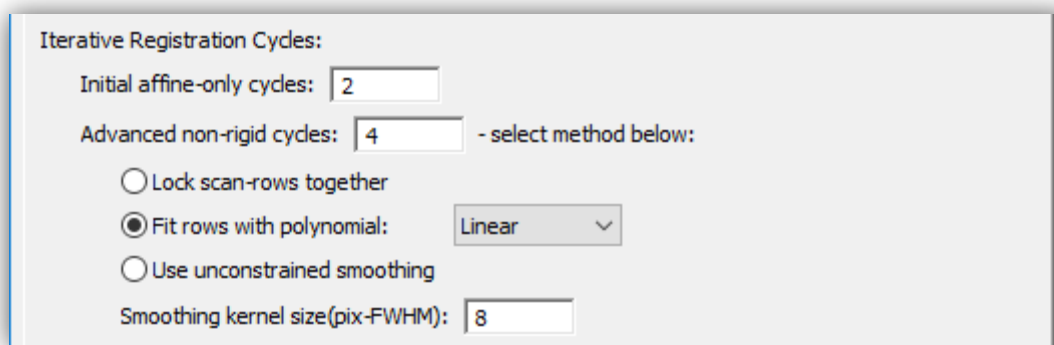
After finishing the rigid-registration, non-rigid registration can be performed. Before beginning the non-rigid registration, if it is not already, the results from the rigid registration should be selected to bring it up as the front window.



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.



Iterative Registration Cycles:

Initial affine-only cycles:

Advanced non-rigid cycles: - select method below:

☐ Lock scan-rows together

☒ Fit rows with polynomial:

☐ Use unconstrained smoothing

Smoothing kernel size(pix-FWHM):

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 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.

When the image distortions are large, however, an average of all frames would cause an excessive degradation in reference image quality. In such a case, a local average of the images immediately before and after each image in the series can be used as the reference for first few cycles. This option can be activated by selecting 'Use nearest neighbour average'.

NOTE: An alternative use for the 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.

Fixed Frame

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

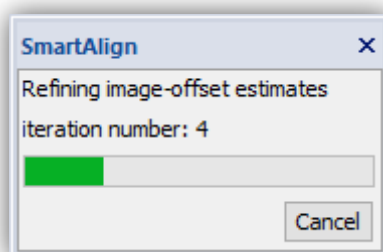
Output Options

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.

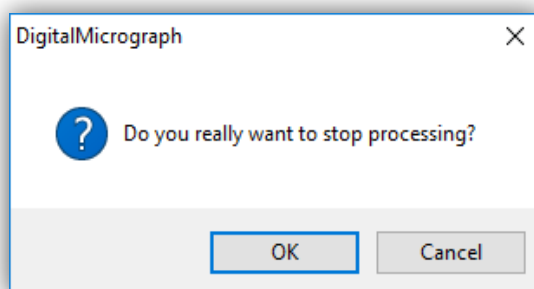
The user can also choose here how to output the compound image after processing. In general, the 'average' image is the most suitable (including for fractional ADF or STM/AFM height data), but a 'summation' image may be useful if the data represents counts or some other quantity you wish to accumulate.

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.

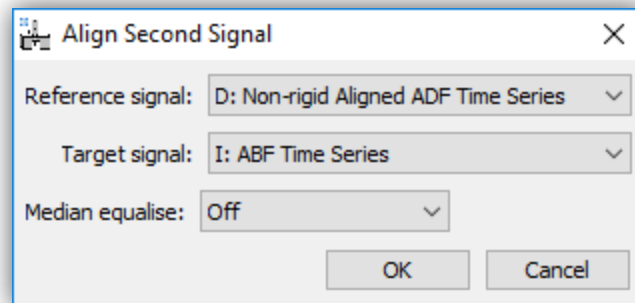
The NR Outputs and Image Cropping

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.

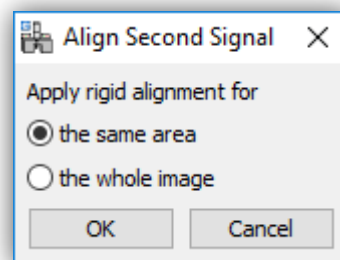
To do this select 'Align Extra Signal using Reference...' from the main menu.



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

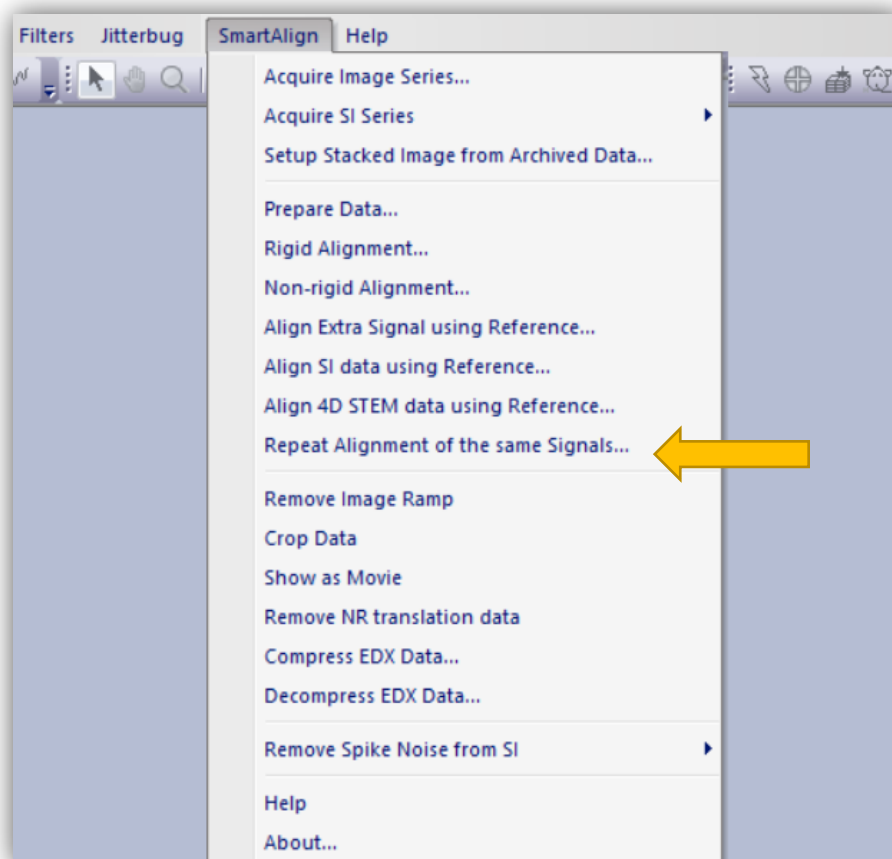
In the case of NR registration, it is only possible to return the same field-of view, and the above choice will not be presented.

Repeat Alignment of Similar Data-sets

Where many data-sets are recorded in an experimental session, it may be convenient to align all of these afterwards using the same settings. This will generally be suitable where the same (or very similar) magnification and pixel size has been used with the same detectors.

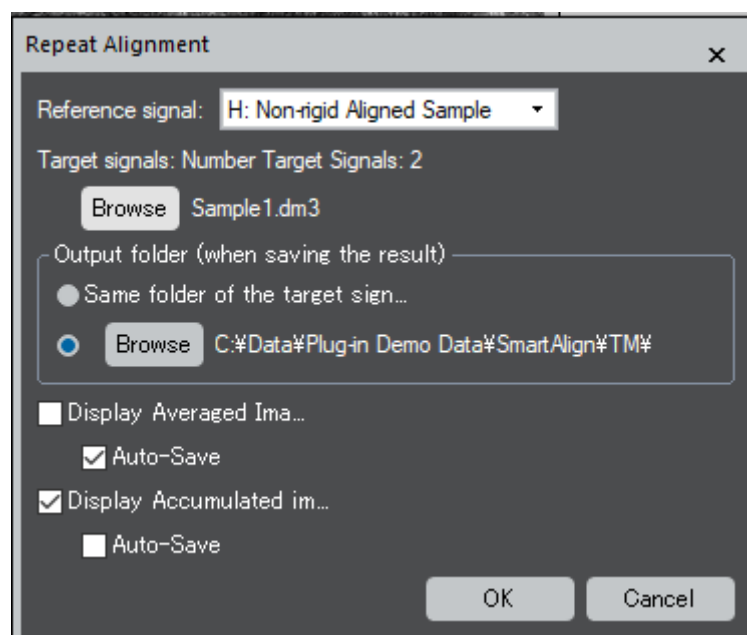
This automated processing can be left unattended and is accessed under the following menu:

Repeat Alignment of the same Signals...



The working principle here, is first that one data-set should be reconstructed as previously described above to determine the user settings as desired. These user settings will be stored in the metadata of the aligned image. On selection of this '**Repeat Alignment**' menu, the same alignment settings will be read from this metadata and then reapplied to the selected files.

The options needed to perform this repeated alignment can be specified in the dialog shown below:



Description of Options

- Here, the user should select the 'Reference signal' – this is the data-set that has already been aligned and from which the metadata will be read.
- Next, the user should click 'Browse' and navigate to the additional data to be aligned. CTRL+click can be used to add each clicked data, or SHIFT+click to select a range of multiple data.
- The next option allows the user to specify that the process data should be recorded back to either the same folder as the raw data or to specify a new location.
- The final options, also allow for the display and/or saving of either the averaged image or accumulated image.

Where a large number of data-sets are being processed (e.g., overnight), please ensure that the disk drive being used will have sufficient space for the outputs. Additionally, where the repeated alignment is wholly unsupervised or where a very large number of data-sets are being processed, the option to display the outputted images can be disabled to avoid unnecessary memory usage by DigitalMicrograph. The saved results can instead be reviewed later (e.g., the next morning).

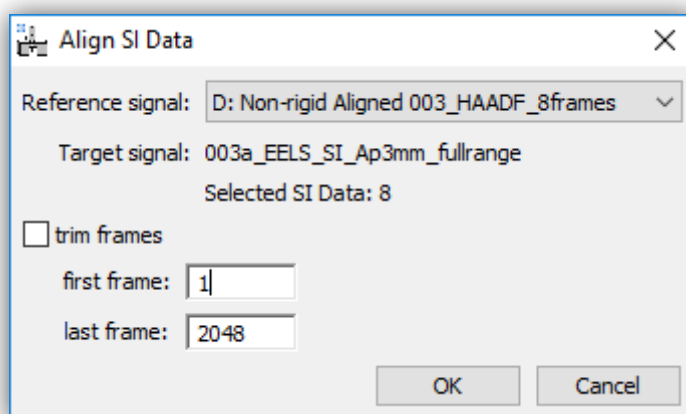
Aligning Spectrum Image Series

If the image series which has been aligned represent the simultaneous image 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 using Reference...' 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.



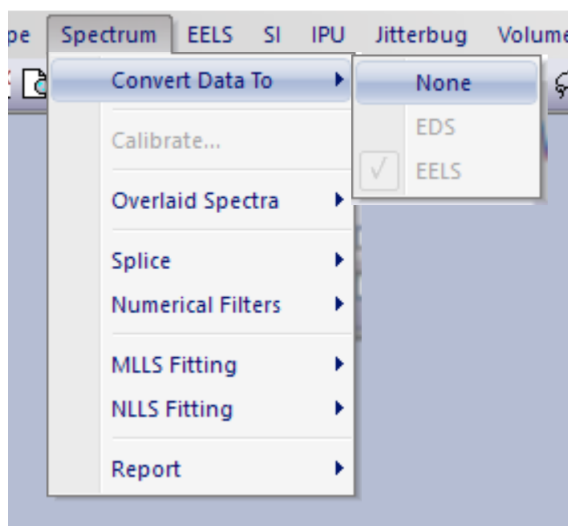
The size of the spectrum images in the third dimension (number of channels) will also be displayed, and if necessary, a reduced range can be reconstructed to minimise RAM usage by selecting the check box "trim frames" and specifying a range of channels (frames).

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.

Post Summation Data-type Verification

Depending on the type of multi-frame spectra that are ultimately summed, it may be necessary to check the data type. If the wrong data type is selected, this can be corrected by selecting: Spectrum > Convert Data To, and choosing your desired data type from Spectrum menu.



Aligning 4D-STEM Data Series

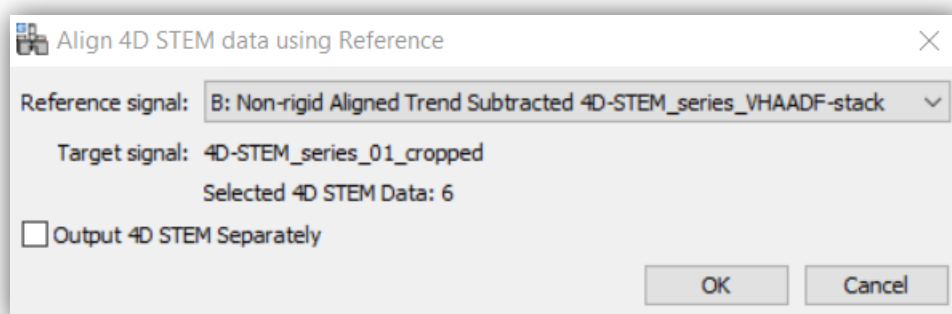
(This command requires the 4D-STEM module license)

If the image series which has been aligned represent the simultaneous image frames accompanying 4D-STEM acquisition, then the probe offsets can be copied from the image series to the series of diffraction volumes. Where a conventional imaging detector is not used during 4D-STEM acquisition, then an image series synthesised from the 4D data itself may be used as the diagnostic image series.

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.

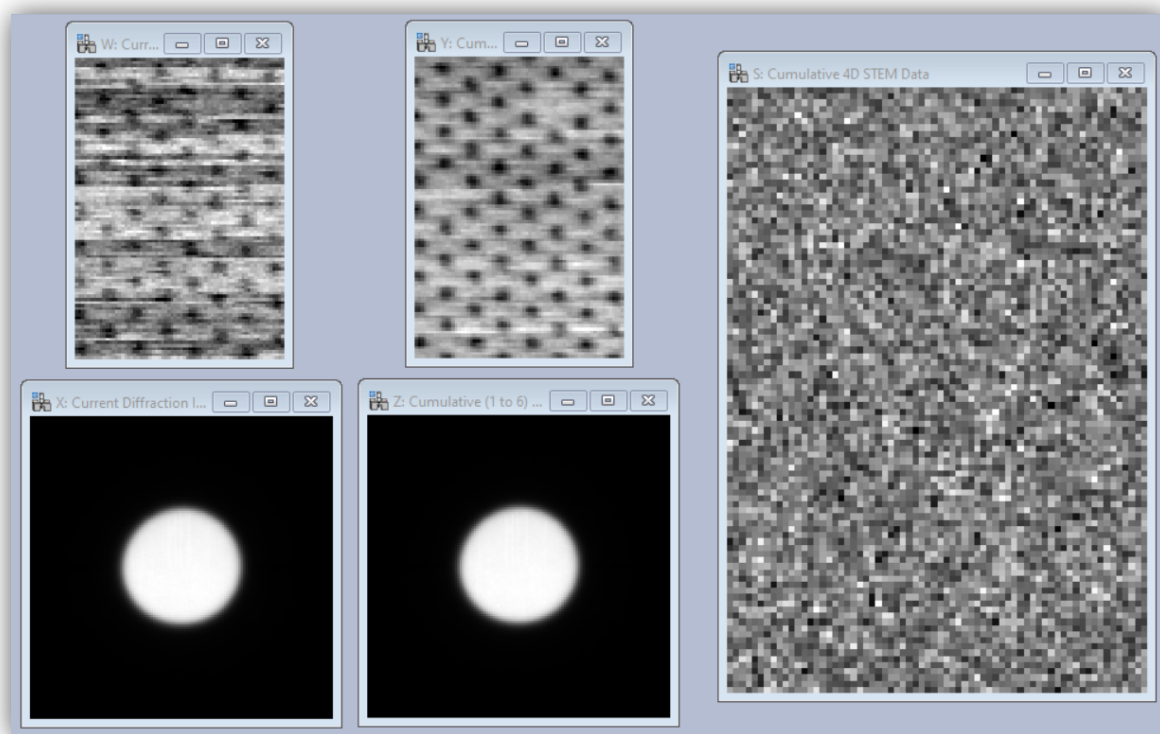
With the aligned image series selected as the front window, select 'Align 4D STEM Data using Reference...' from the *SmartAlign* menu.

Then, a file explorer dialog will open where you should select all the files that represent the 4D volumes associated with the series. The below example shows the result of selecting 6 data volumes.



A checkbox 'Output 4D STEM Separately' will be shown giving the user the option to 'Output 4D STEM Separately'. If selected, this will align and non-rigid register the 4D volumes in the spatial dimensions, but will then store them to disk separately.

During the registration process (which may take some tens of seconds depending on the PC performance), the user will be shown the current 4D data being processed as well as the cumulative data processed up to that point:



Once the processing is completed, the 4D volume data is displayed. This is a 4D data and may be inspected using the SI Picker tool. (If you don't have Gatan SI Picker Tool, the HREM Picker Tool can be download from HREM Research home page.)

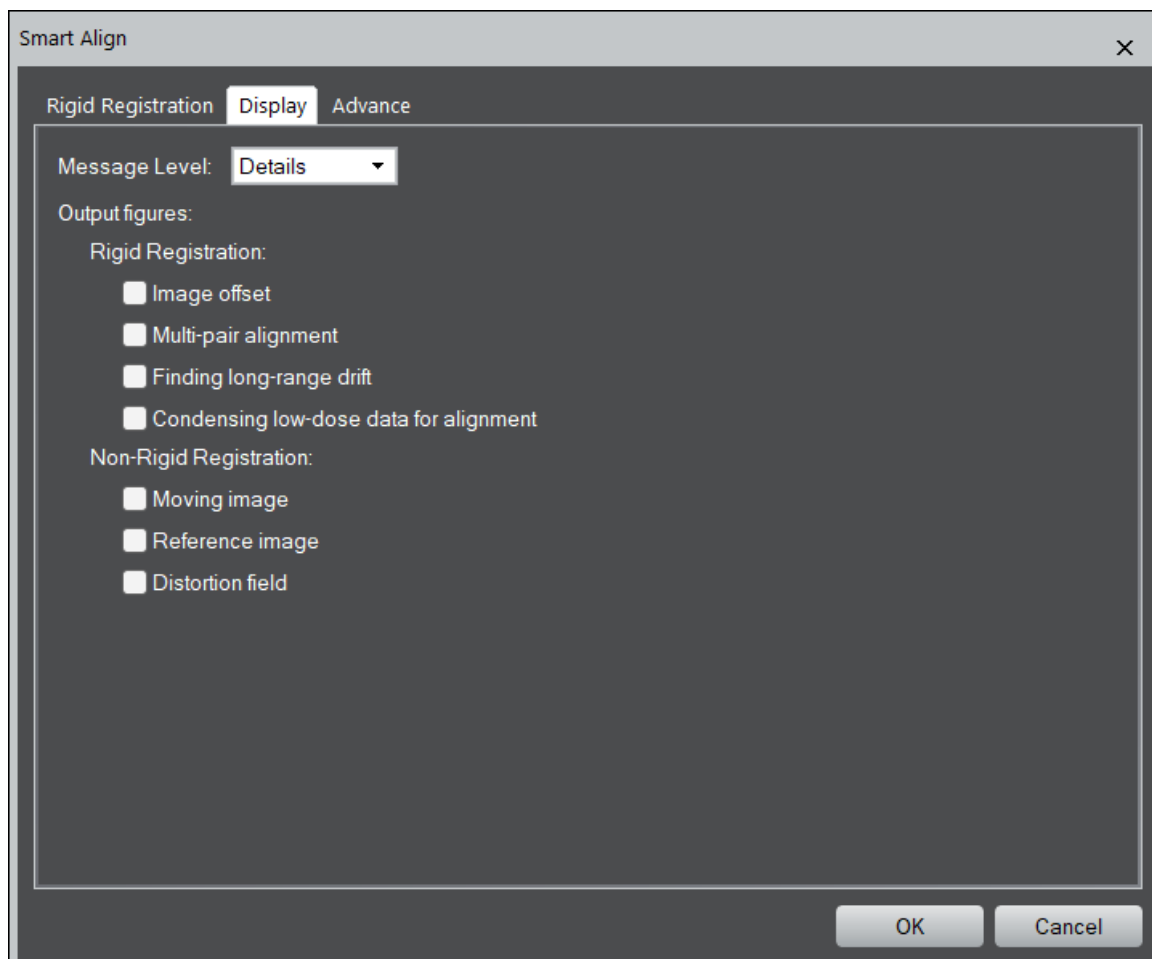
At this point, the 4D-STEM data may be saved and/or exported to another package for onward analysis.

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.



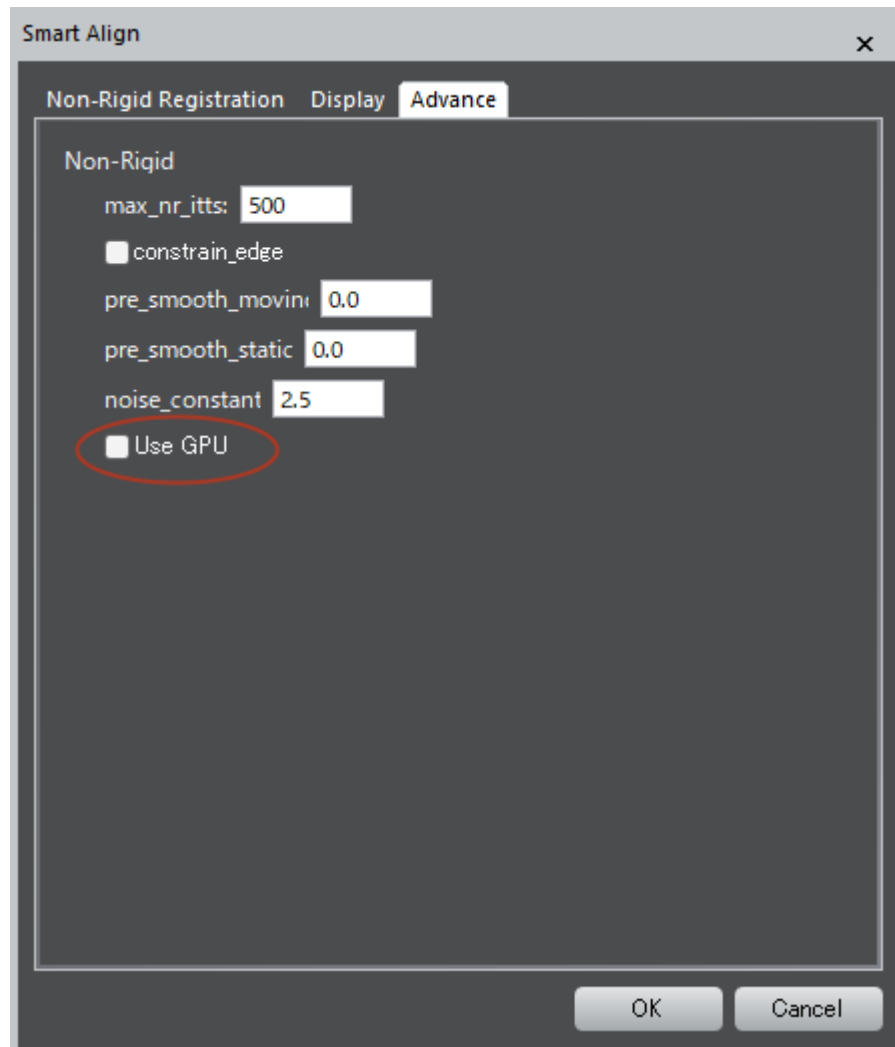
These options include a choice for the granularity of text message outputs (displayed in the DM results window) during processing, as well as the additional graphical display of variables.

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

NOTE: The same dialog will appear also for Non-Rigid Registration.

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 iterations 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.

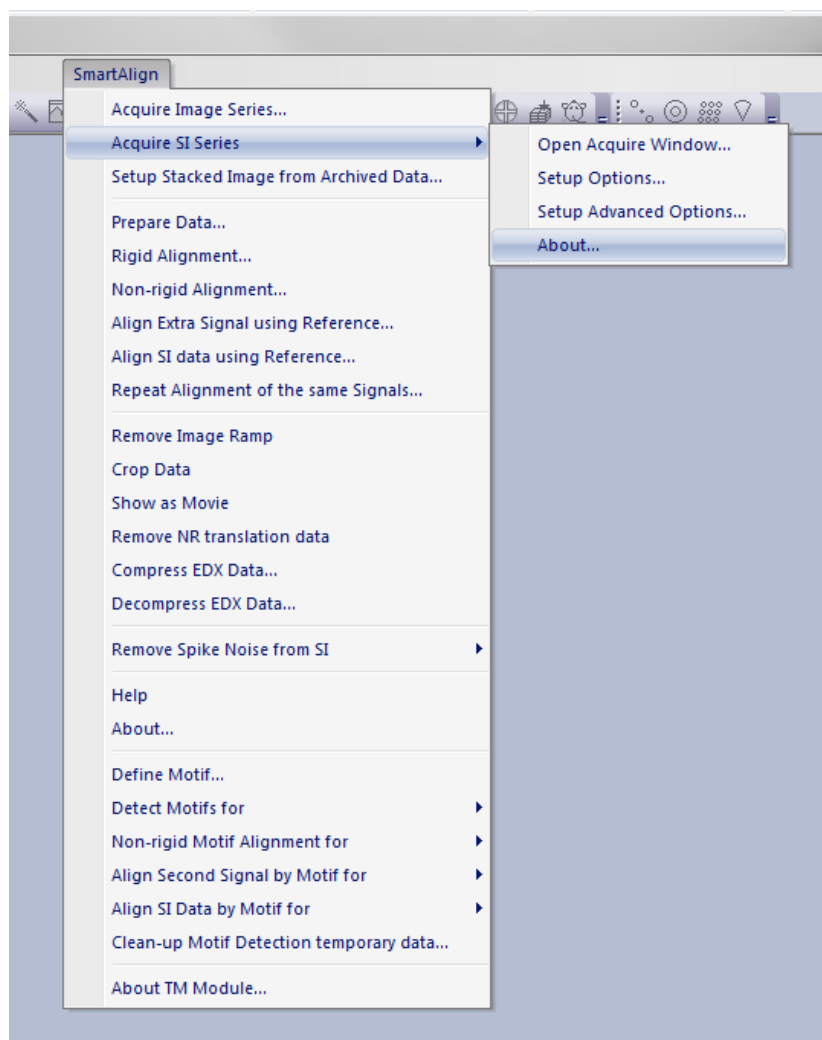
In the case of very noisy data, a small smoothing can be applied to the individual images or to the average reference image.

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.

Use GPU option (Non-Rigid Registration only): You can speed up Non-Rigid Registration if you have a suitable NVIDIA GPU (Graphic Processing Unit). **NOTE:** This option will not appear for Rigid Registration.

Spectrum-Image Acquisition Menu

From SmartAlign version 2.1 onwards, a menu to assist with the recording of multi-frame spectrum-image (MFSI) series is included. These options are found under the menu “Acquire SI Series”:



What the SmartAlign Spectrum-Image Module Is and Isn't

The spectrum-image module does not itself add functionality to the base installation of DigitalMicrograph. To capture spectrum-images in DM, users must already have those signals available. What the SI module does add, is a data-saving helper to automatically capture, save, and close from memory sequential SI scans.

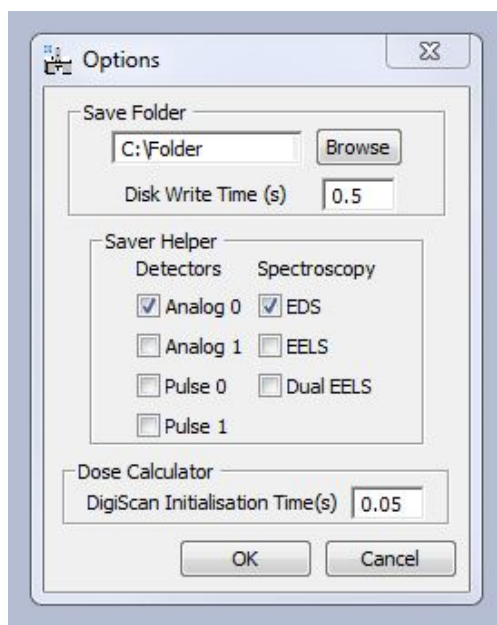
The Workflow of using the SI Module

The SmartAlign SI-module comprises of one main dialogue, plus two further settings menus. The main menu comprises of three buttons for data-recording, plus an acquisition planning tool that can be used to predict the experiment duration and electron-dose exposure.

Before using the MFSI acquisition menu, the settings for the experiments should be prepared.

Setting up a MFSI Acquisition

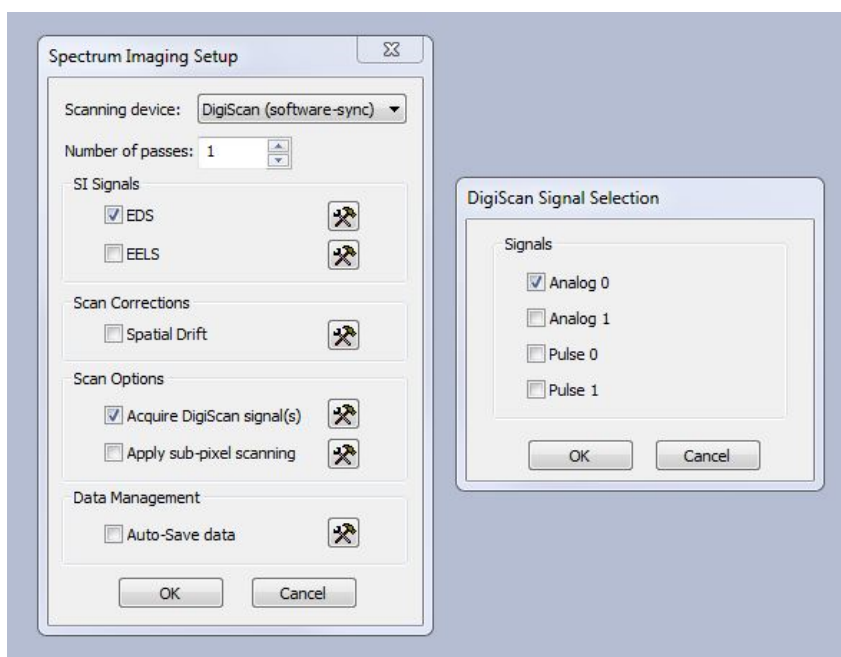
As the SmartAlign MFSI module automatically captures and saves data, the settings for the data saving (file location, file-type and signal choices) should be specified in advance. This is done using the “Setup Options” menu:



Using the “Browse” button the user should select the save location. You may create a new folder if a suitable one does not exist. It is recommended that a local drive (ideally with a good disk write speed) should be used. Use of a network drive with poor speed or latency may delay data-saving.

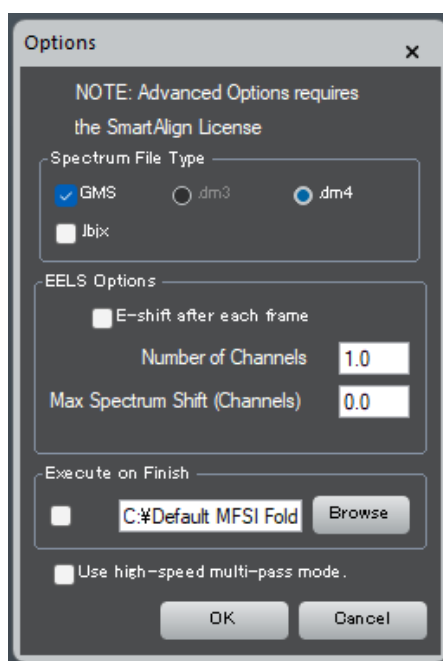
At the same time, the signals to be saved should be specified.

*Note – these settings are dictating which signals will be *saved*, and not which ones will be *acquired*. To ensure the data is available to be saved, the user should verify these settings match those selected in the Gatan SI menus (see following figure).



*Note – in the multi-frame SI acquisition, each data obtained for a single scan (passes=1) will be saved with the specific name. Therefore, in the Gatan SI settings, “Number of Passes” should be set to 1, and the Gatan “Auto-Save data” should be set to off.

The SmartAlign “Setup Advanced Options” dialog below contains less commonly used settings



Here, you can specify the file format to be used (Gatan file format or JBix format) , and EELS energy shift options to achieve either Binned Gain Average or custom dark-reference experiments [7,8].

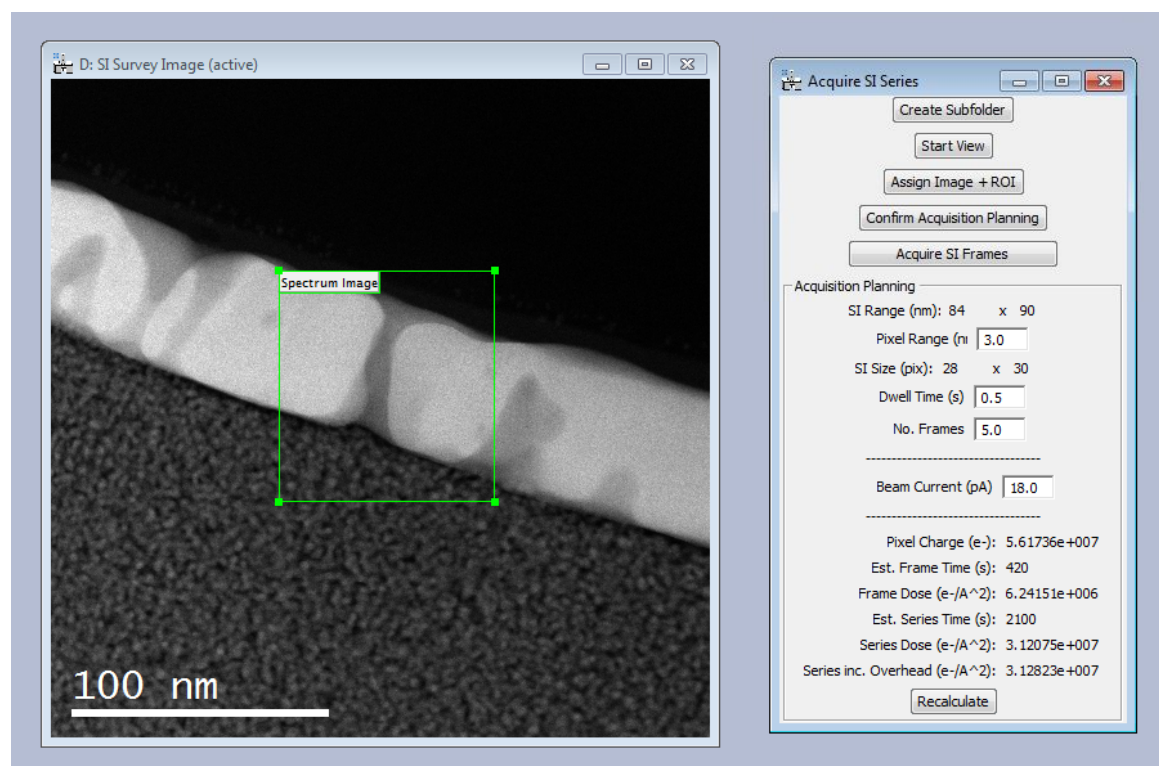
An interesting setting is "**Execute on Finish**". Any script selected here will be performed when the data acquisition is finished. For example, you can get an email through DigitalMicrograph. You can find an interesting sample script in the DigitalMicrograph Help.

When the last setting "**Use high-speed multi-pass mode**" is activated, a multi-frame SI data can be acquired at high-speed using the passes setting of Gatan SI acquisition. Here, each SI data can be saved individually using the hook-up capability, rather than a single accumulated SI data as in Gatan SI acquisition. This mode will be useful, especially when the number of frames is large and the exposure time for each frame is short.

*Note – 1) "Auto-Save data" in the Gatan SI settings should be set to off. 2) the "HQ Dark Correction" will not be applied to each data. 3) If your GMS does not support the hook-up capability, only the not-accumulated last data will be saved. In that case, please deselect this option.

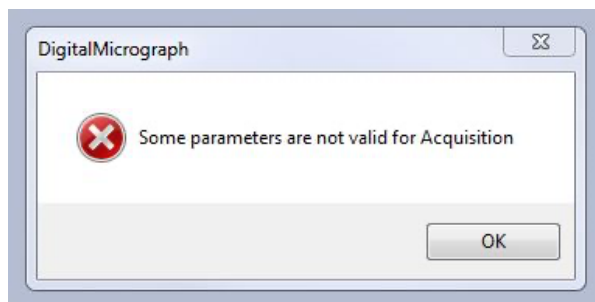
Performing a MFSI Acquisition

The MFSI acquisition workflow follows the following steps for each region of interest within an experimental session. The user should first have a Digiscan image open of their sample, and should use the Gatan SI menu to ensure this is assigned as the *active* image with an active ROI.

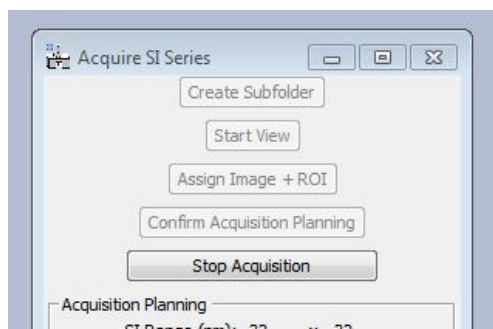


With this active image and ROI assigned, the three buttons guide the user through the process.

- 1) After navigating to a new region of interest in the sample, the first button prompts the user to create a new sub-folder to store data in. This will be a subfolder for the site of interest within the folder already specified via the settings menu. In this subfolder the numbered images and SI will be stored with automatic sequential numbering in the filenames. Users should avoid changing the filenames to avoid unintended reordering that may be relied upon later during the data registration. If the user creates a subfolder with a name that has already been used a warning will be displayed and the user should select a different subfolder name. Reusing the same name may lead to data being overwritten.
- 2) At this point the size, shape and position of the ROI can be adjusted. Next, the user should specify/confirm the pixel-step which will be used, as well as the dwell-time and number of frames. If any of these parameters are changed, the user should press the 'Confirm Acquisition Planning' button to update the settings for their acquisition. In this planning menu the pixel-size is used to determine the number of pixels needed to cover the desired ROI, and then using their dwell-time, to estimate the time needed per frame. With multiple frames, a total time for the whole series is also estimated. In addition, if the beam-current is specified, then various dose prediction are estimated as well. If any parameters are not suitable (or are left blank) the user will be warned before beginning the automated acquisition.



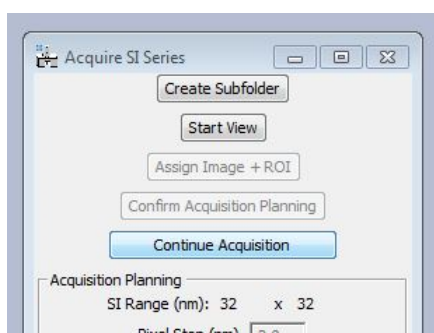
- 3) Once the acquisition parameters are confirmed, the user should press 'Acquire SI Frames'. This will begin the acquisition, saving and closing of the specified number of frames. Once the acquisition has begun, it is possible to stop the acquisition by pressing 'Stop Acquisition'. This will abort the current frame in progress:



*Note – both in terms of the time prediction and the dose prediction, these are more accurate where hardware synchronisation is used. Where software synchronisation is used, significant additional communication overheads may exist and this will increase the time and dose required to achieve the same results.

Recording Additional Frames

After recording the specified number of frames (5 frames in the example screenshot above), it may be that the sample remains undamaged / free from excessive contamination. In this case, additional frames can be captured with the images and spectra continuing to be added to the numbered series in the same folder.

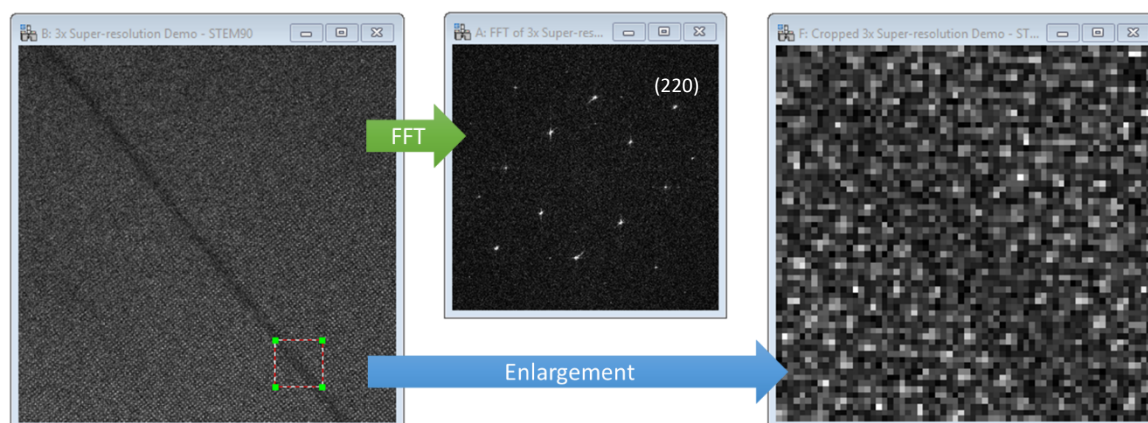


In this case, the user will see the button change to “Continue SI Frames Acquisition”. The series numbering will be reset when the user again presses “Create Subfolder”.

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.



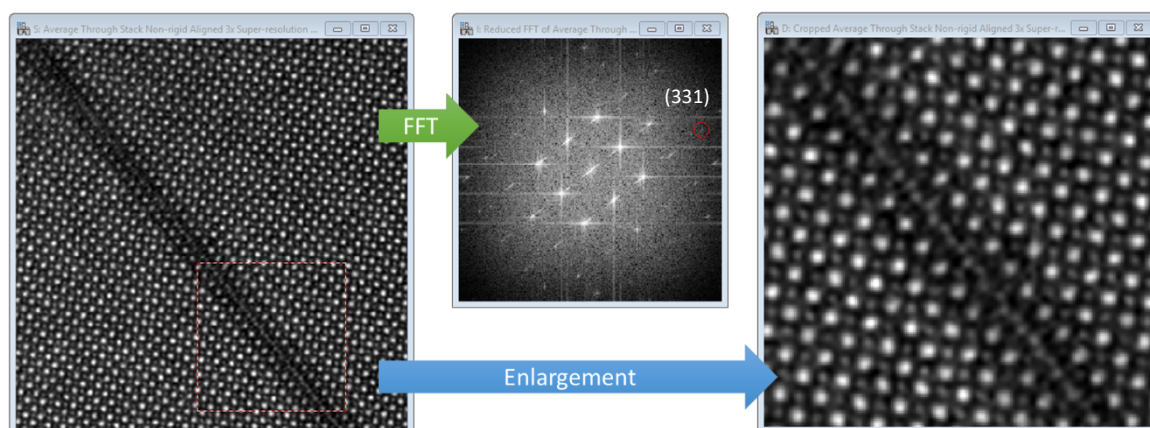
The FFT shows spots are visible right out to the Nyquist limit, and the highest observable spot is the (220) spacing (approx.. 1.38Å. A 64x64 pixel ROI is enlarged on the far right – again, not how the sampling is too low. Ordinarily this would be a poor choice of setting to record the image; we would 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*, where the resolution is extended beyond that of the original STEM probe, but rather this is *digital super-resolution* where the 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.

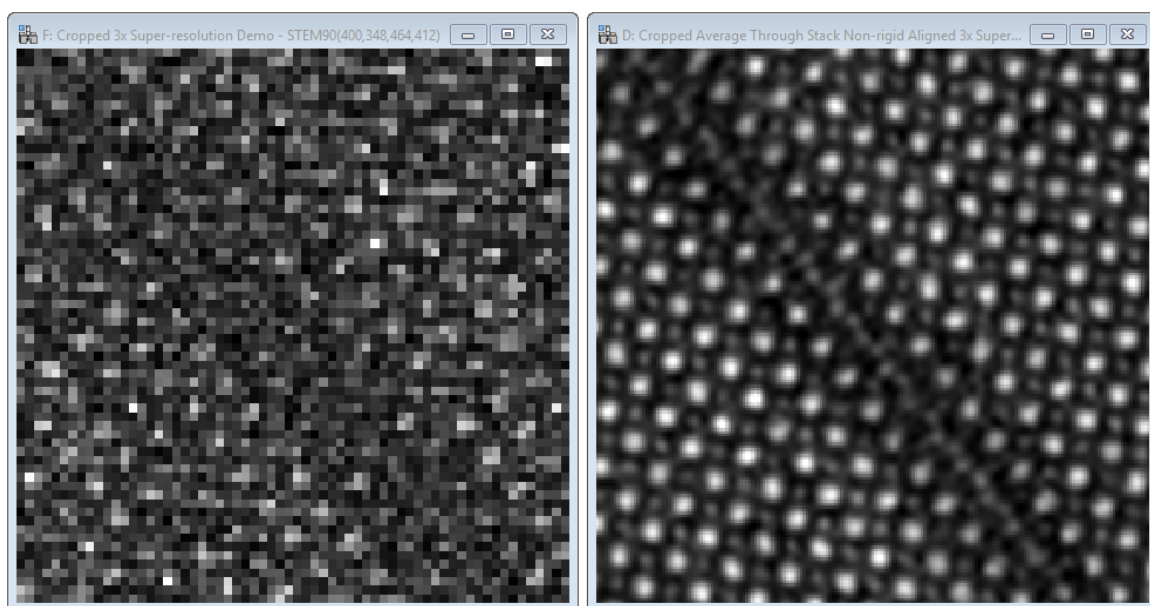
The results below show the lower-right part of the image after resampling and alignment, the reduced Fourier transform (central 50%), and a cropped enlargement of the same region as in the earlier figure above.



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.

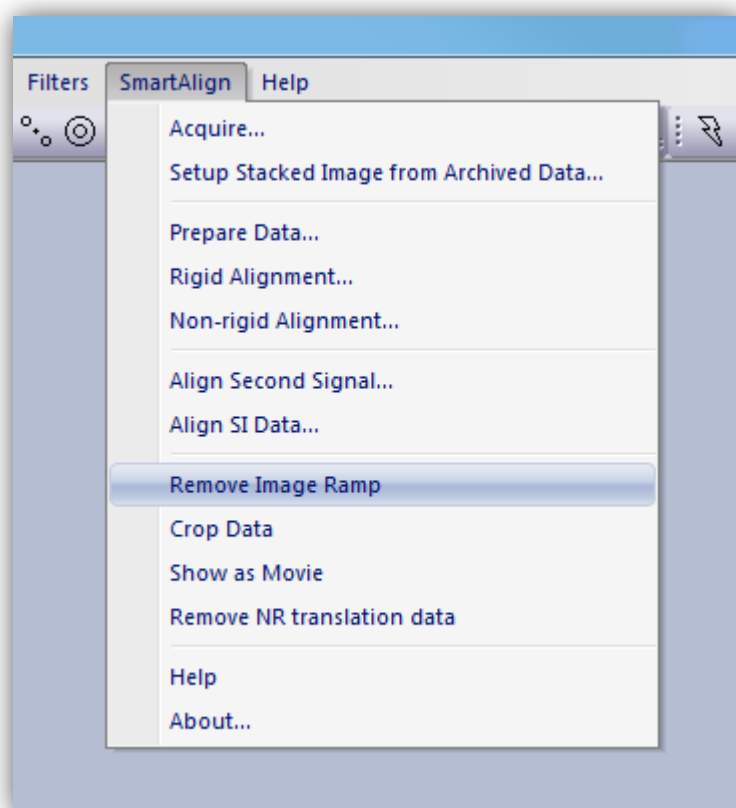


Going beyond 3x data-resampling will likely offer little benefit and will mostly just add to the calculation time.

Processing STM & AFM Data

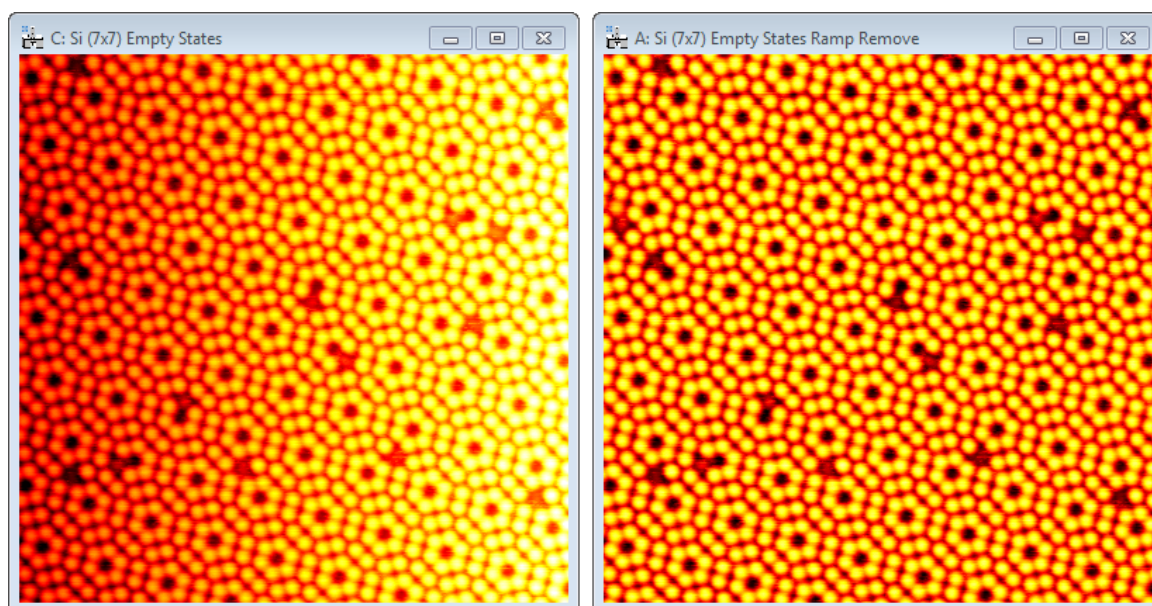
For many scanned probe microscopy (SPM) datasets, it is often necessary to first remove a planar ramp background. This tool is provided as a utility in *SmartAlign* and can be done at any point. However for best results, it is recommended to perform this step at the start of the data processing.

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

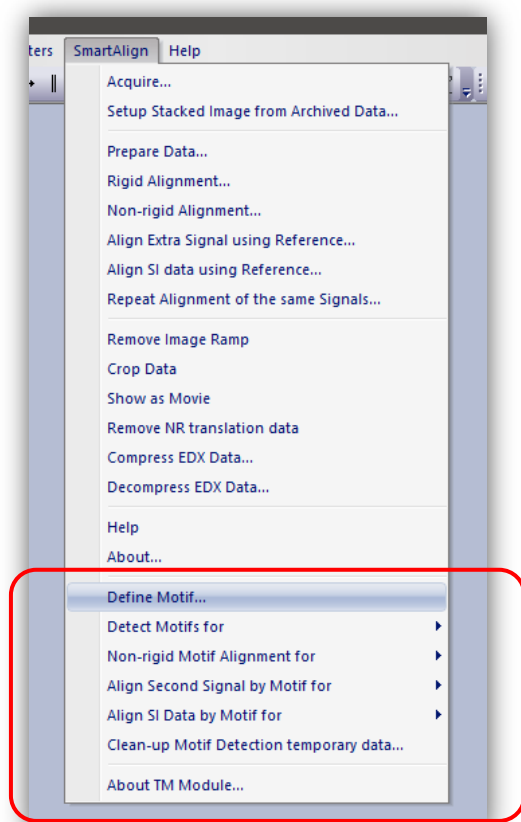
Template Matching Module for SmartAlign

(This module requires an *additional license*)

Introduction

Template matching is a common technique that involves searching for repeating regions within an experimental image of a crystal sample, and averaging them to improve image quality. However, the “Template Matching Module” (TM module) for *SmartAlign* adds further functionality, namely it includes non-rigid registration of the repeated image-regions. Thus, the TM module is extremely suitable for data that has been acquired by scanning techniques, such as STEM or SPM. Moreover, the same non-rigid registration information can be applied to other data acquired simultaneously, such as Spectrum Images (SI). Furthermore, the TM module enables the operator to easily perform template matching over multiple data (image frames) saved in the same folder to improve the data quality substantially.

This section describes the optional “Template Matching Module” (TM module) add-in for *SmartAlign*. When the license for the Template Matching Module is activated, the commands for the TM Module will be added below the *SmartAlign* menu commands as shown below:



Menus for
Template Matching

Template matching searches an experimental image for regions that repeat and extracts those regions to yield a series of matches. Here we will refer to the reference, the patch which we search for, as the “motif” and the extracted cropped regions as “tiles”. Once a series of tiles has been extracted from an image, this series of observations can be considered as if an experimental frame-series itself and is eligible for processing using *SmartAlign* in the usual way. This section of the manual then, describes the process of extracting this series and how to prepare it for onward processing.

Note: The TM module requires an independent license in addition to the SmartAlign license. The license for the TM module will be added remotely using Remote Update System (RUS).

Template Matching across a Series of Images

When using template matching on image-series data, the user can choose whether to align that series first and then extract the TM tiles, or to extract the TM tiles first and then align (distortion correct) those tiles. The two cases are not quite equivalent, and the choice of which to follow depends on the user’s experimental data and goal. These two work-flows are described in more detail below.

Template Matching followed by Non-rigid Registration (TM → NRR)

If your image series is of a good SNR and contains only a few repeats of the motif of interest, then performing the NRR across the full field of view would add calculation time for little benefit. In this case it is better to first use the TM to extract the tiles, and to only perform the NRR on those smaller regions. If the SNR size is poor, then it may be necessary to first register the images (see below).

Image Registration followed by Template Matching (RR → TM → NRR)

When the image-series SNR is very poor, such as low-dose image data, then correctly identifying the repeats of small motifs may become less precise. In this case, it is preferable to first align the series data by rigid registration and collapse this to a single average frame with higher SNR *before* performing the TM. This approach is also useful, when there are many repeats of the motif in each frame, and where extracting all these from every frame would yield an excessively long synthetic series.

Name convention for a series of images

The correct numbering structure is “Data_name_[xxx]”, where [xxx] is some increasing number series. This structure is the same as the DM menu “File > Save as series...” command. The image-series (as separate files) and the SI-series should be stored in the same folder.

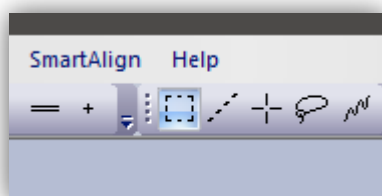
Pre-treating the Data

In the same way as for aligning image series, it can be useful to first remove any image-ramp intensity variation. This may help the TM module identify tiles more reliably. Similarly, if processing STM/AFM data, first rigid-aligning and tip-height correcting the data will help in the identification of more reliable motif matches.

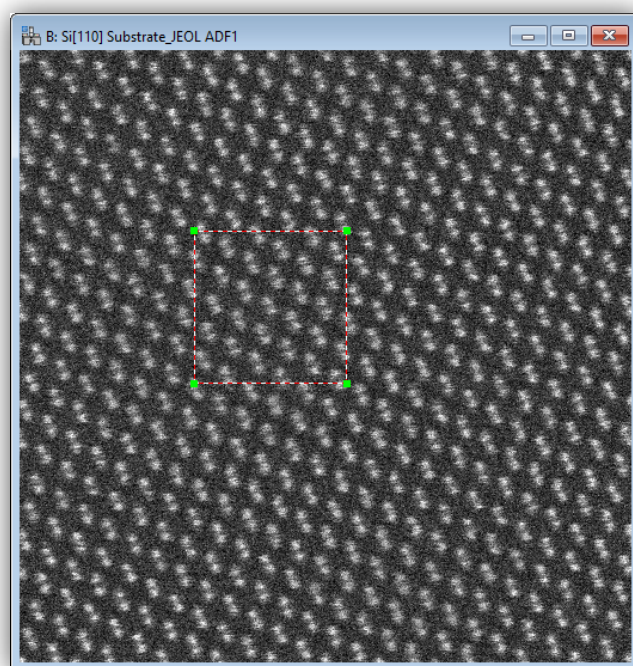
Defining the Motif

There are two approaches for defining the motif; from the image itself or from an external source (such as a simulation or some other high-quality image data). If using an external motif from a simulation or from another image, you can skip to the next section.

To identify a motif within the image that is of interest, simply draw a ROI using the ROI tool from the menu.

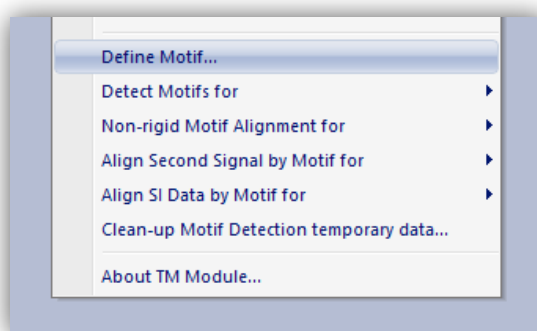


There is no restriction on the size or shape that this ROI can take, other than it must be smaller than the parent image. However, for images with a low SNR, selecting a very small ROI may lead to less reliable tile repeats being identified.

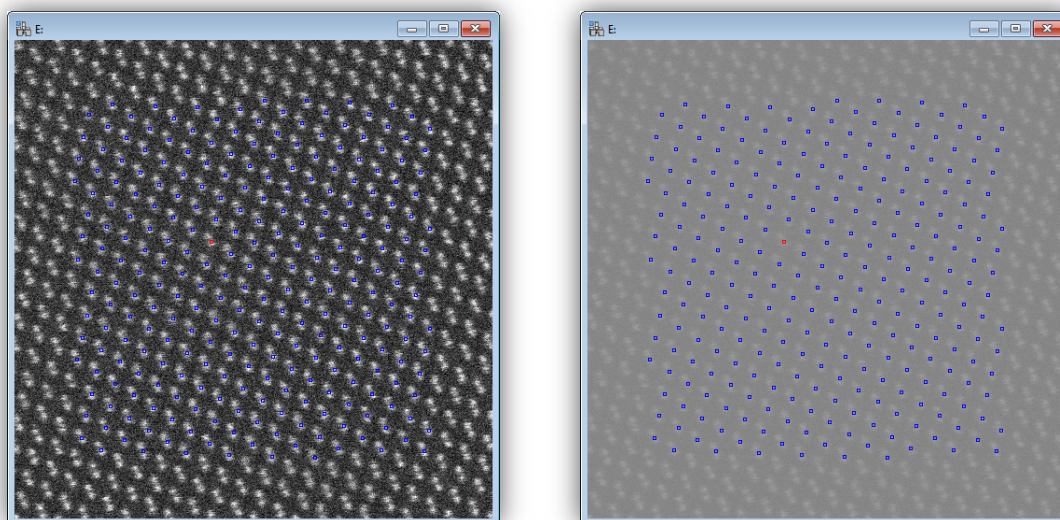


Note: the number of tiles that will be identified will depend on the size of this ROI. For example, if the ROI is small and contains only (say) one unit-cell of a crystal structure, then the number of tiles found will be approximately the same as the number of unit-cells in the image (excluding part-cells at the image edges). However, if the ROI is very large, the number of cells excluded at the edges increases and the number of tiles found will decrease.

We will now determine the motif from the image in a two-stage process; first we estimate the motif from a single noisy ROI, then this is used to search the image again. This two-stage motif calculation makes the template matching approach far more robust to experimental noise and distortion. With the ROI drawn select “Define Motif...” from the TM command:

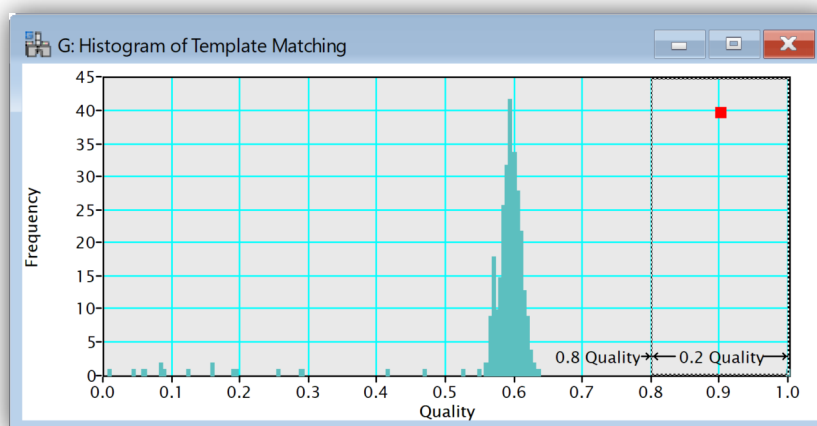


The image will then be searched for potential repeats of the motif, and these will be labelled over the image as blue dots*:

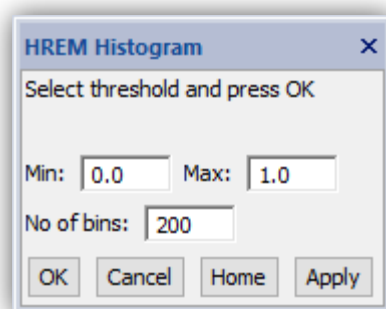


Tip: Adjusting the image contrast or gamma (see above) may make these markers easier to see.

At this stage you will also be presented with a histogram of the correlation-scores that indicate the quality of the matches that have been found. A perfect score would be one, whereas zero represents no-match at all.



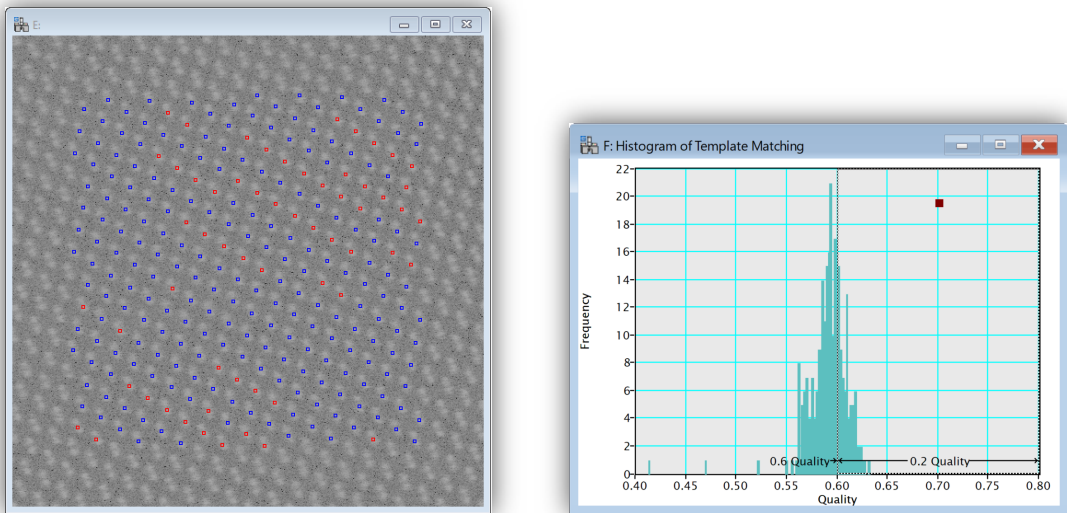
The display settings for this histogram can also be adjusted, for example to enlarge one part or adjust the number of bins, using the histogram dialogue.



After changing any of the parameters, simply press “Apply” to update the histogram.

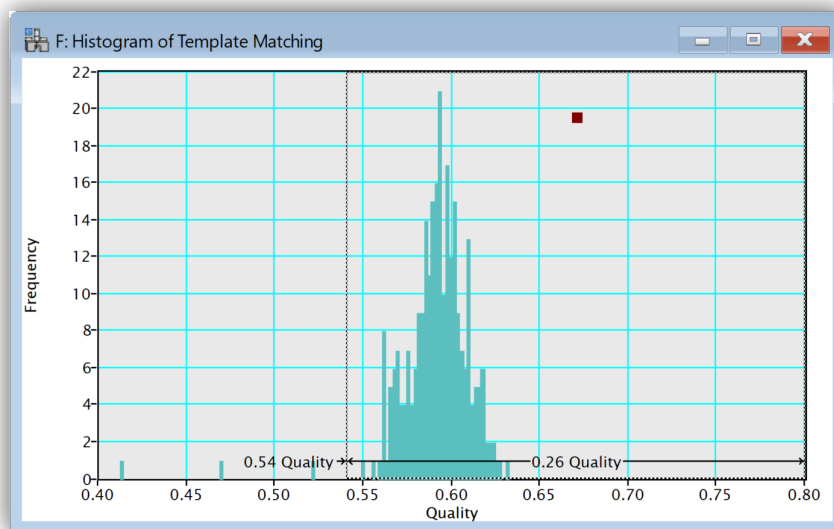
Next, the user must select the quality threshold of the tiles to keep. This is done by dragging the left edge of the dashed box displayed on the histogram. As you adjust the threshold the colour of the markers will change from blue to red if they are above the threshold. In this way you can interactively determine the appropriate threshold.

The figure below shows an example where the histogram scale has been enlarged to show the region 0.4-0.8, with a threshold set part-way through the main peak in the match-quality histogram (at 0.6):



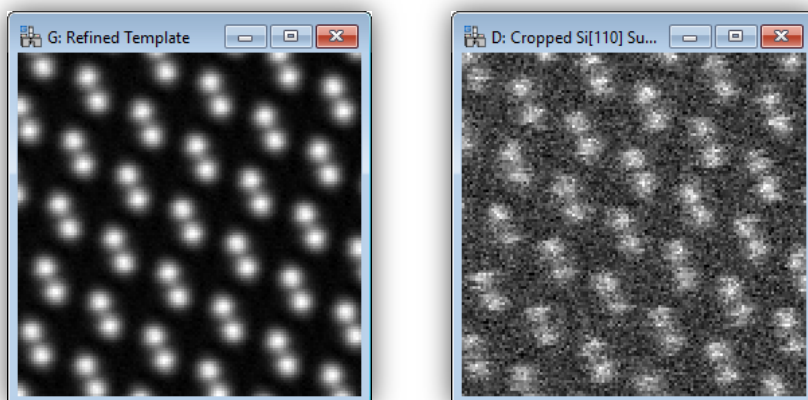
The result is that the highest quality matches (most similar to the motif) are coloured red and the other (non-selected) candidates remain blue.

In the next example below, selecting a quality threshold of 0.54 includes all the reasonable matches while rejecting poor ones around the image edges.



In general, the histogram will often show a peak around some value where the genuine matches are grouped, and then some outliers to the left side that are poor quality matches.

Once you are happy with the set threshold press “OK” in the histogram dialogue and the identified tiles will be extracted and averaged. This averaged tile then represents the same area as the manually drawn ROI but at a far higher SNR (compare with the single manual ROI below-right).

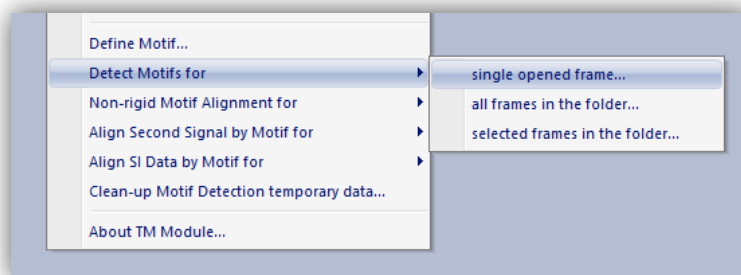


This image has the exact same pixel dimensions as the manually drawn ROI and is now available to be used for the final template matching (next section).

Detecting Occurrences of the Motif

Single Opened Frame

The manually drawn ROI on the experimental data can now be deleted, and we use the result from the previous stage to search the image. From the SmartAlign menu, select “Detect Motifs for” -> “single opened frame...”:

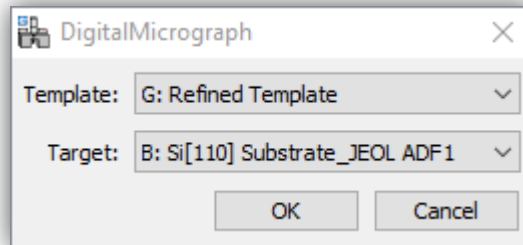


The user will be presented with the “Prepare data” dialog described in the earlier part of the manual with some options pre-set ready for template matching. At this stage the user can select up-sampling to achieve digital super-resolution, however as before it may take time to apply this if the number of extracted tiles is very large.

This option will operate on an image (or image stack), which is open in the DM window. The other options in this menu can be used to search through the images in a folder or a subset of those; note that the time take to batch process multiple images in this way scales with the number of images.

Note: It is strongly recommended to test the TM on a single opened frame first before processing an entire folder. However, in the case of “Detect Motif for...” it is a *must* to perform “single opened frame” before processing multiple frames.

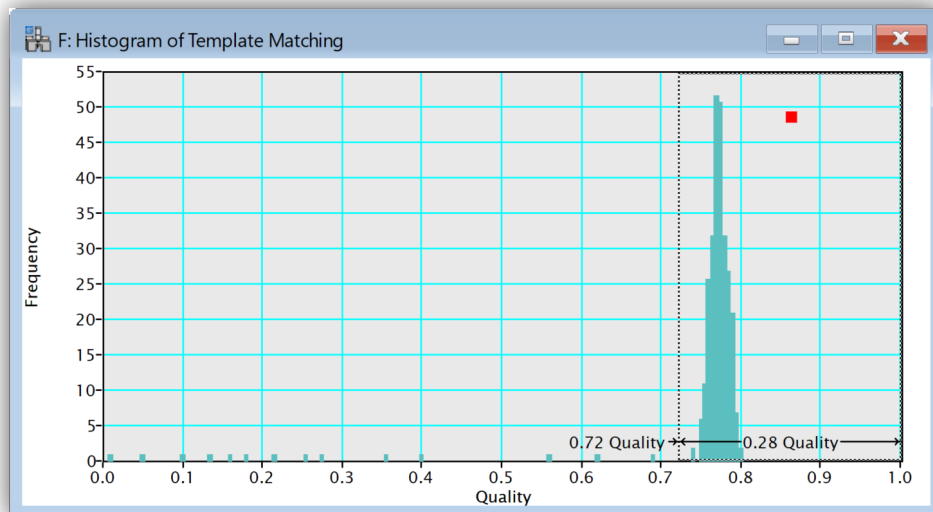
Next the user is asked to select the motif and the target images:



Having created a “refined template” in the previous section, this should be selected as the template and the experimental image selected as the target. Alternatively, if the template comes from some external source (such as another experimental image patch or an image-simulation) then this can be selected instead.

Note: when using an external reference template, the pixel-size (magnification), contrast type (bright-field versus dark-field etc.) and the scan-rotation must be the same as the target image.

Again, the histogram of match quality will be displayed:



As the refined template was used as the search motif, we now see that the peak has shifted (in this example case) from around 0.59 to 0.77. This is due to the two-stage motif definition process followed where the refined template had an improved SNR over the initial estimate.

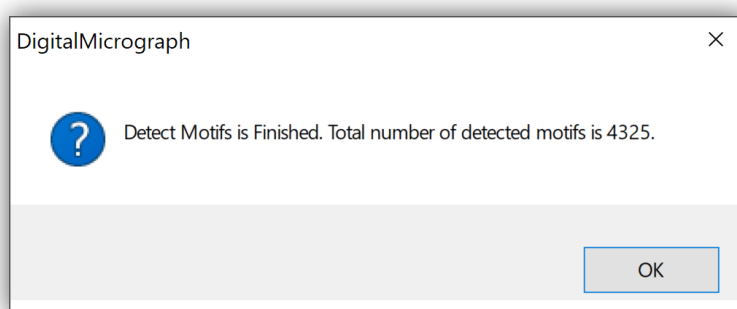
Once the threshold has been selected, the individual tiles will then be extracted and stacked to yield an image-series in the workspace.

Multiple Frames in the Folder

After detecting motifs from a single opened frame, the user can go on detect motifs from multiple frames saved in the same folder. To do this, select “Detect motifs for...” -> “all frames in the folder”, or “Detect motifs for...” -> “selected frames in the folder”. A file-open dialogue will open and you should select the first file from those with the sequential naming scheme. When you select “Open”, this will search for, and store all the locations of, all matches across the series of image-files. You will be asked whether to use a ‘fixed threshold’, this option will use the same match quality score between all frames and is usually the most appropriate choice (all frames will be processed unsupervised). Should the data suffer from irregular contrast changes or excessive beam damage, you may want to disable this option and choose the threshold manually for each frame (operator selection required for every frame).

Note: temporary data to support the calculation is stored in a sub-folder named “SA_temp”, These temporary files store the locations of all the identified tiles and should not be deleted/modified until the whole template matching is complete.

During the motif search you will be able to watch the progress in the Output panel. After all frames in the series have been searched a message will tell you the total number of motif matches that were found across the entire series.



In this example, across the 16-frames of the experimental series, a total of $\approx 4,300$ tiles were identified.

Non-rigid Registration of the Extracted Tiles

Once the tiles have been extracted they can be considered as a series of experimental observations. Where scanning-distortion is present, the non-rigid registration (NRR) may be needed. You can perform non-rigid registration from “Non-rigid Motif Alignment for...” command below:



Note: as the tiles were extracted using template matching it is not necessary to perform the rigid registration step normally needed in SmartAlign for experimental data series.

It is strongly recommended to test the NRR on a single opened frame first before processing an entire folder. To do this, select “Non-rigid Motif Alignment for...” -> “single opened frame” when the front image is the target. Then, you will get the dialog for the “Non-Rigid Registration” described in the earlier part of the manual with some options adjusted for the template matching.

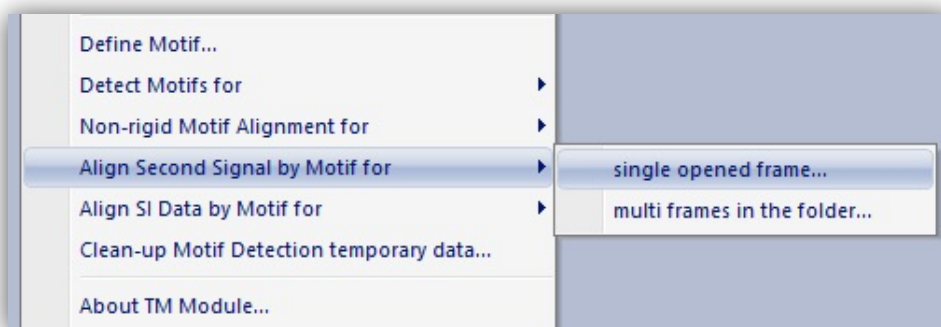
When the user has finished non-rigid alignment for a single opened frame, this can be performed for multiple frames saved in the same folder. To do this, select “Non-rigid Motif Alignment for...” -> “multiple frames in the folder”. A file-open dialogue will open, and you should select the same image as before (the first in the series of frames). Note that the same parameters which were selected for “single opened frame” will be used for the command “multi-frames in the folder”.

Note: temporary data to support the calculation is stored in a sub-folder named “SA_temp”. These temporary files store the transformation data and should not be deleted/modified until the whole template matching is complete.

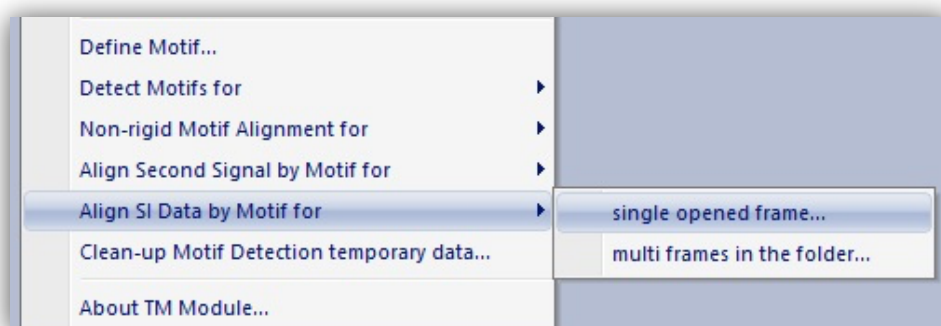
Template Matching a Second Signal / Spectrum Image

Once the repeats of a motif have been identified it is possible to extract the equivalent patches from another image signal or from a spectral image (SI). To perform this, select “Align Second Signal by Motif for ...” or “Align SI by Motif for ...” from the TM commands below:

For a second signal:



For SI data:



The required data is different for these two cases as follows:

TM Across a Second Signal(s)

For this process, the user should have (for example) an image/stack from a dark-field image and a simultaneous data from a bright-field signal. The user should choose whichever signal that has the higher SNR as the “master” signal and should work on this first. Once the TM is complete on the master signal, simply select “Align Second Signal by Motif for...” from the TM command.

As usual, it is strongly recommended to test the command on a single opened frame first before processing an entire folder. To do this, select “Align Second Signal by Motif for ...” -> “single opened frame” when the front image is the target.

TM Across an SI or SI-series

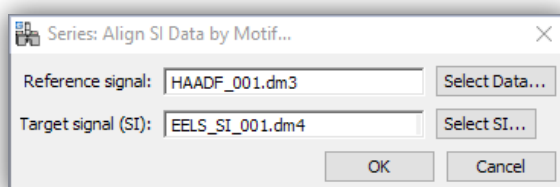
For spectrum image TM, the user must have a single image-SI pair, or a series of image-SI pairs. These must have been recorded simultaneously and with identical x-y pixel dimensions. The images should exist on the disk as separate files, numbered sequentially as for the SI*.

Note: the correct numbering structure is “Data_name_[xxx]”, where [xxx] is some increasing number series. This structure is the same as the DM menu “File > Save as series...” command.

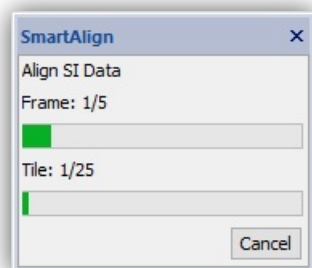
The image-series (as separate files) and the SI-series should be stored in the same folder.

In the case of *TM across a series of spectrum-images*, one image may have 100s of repeated motifs, tens of frames, and thousands of SI channels. Because of the scale of this calculation, it is important to test the behaviour on a single opened frame first before processing an entire folder. To do this, select “Align SI by Motif for ...” -> “single opened frame” when the front image is the target.

When you test the behaviour on a single opened frame, you can perform non-rigid alignment for a multiple SI data saved in the same folder. To do this, select “Align SI Data by Motif for...” -> “multi frames in the folder”. A dialogue will open for you to select the appropriate files; you again select the first image in the series, and the first spectrum-image in the series. Both should have the same sequential naming convention (such as below).



The plug-in will now use the stored motif locations and the stored alignment data to search each SI, extract the spectrum ROI, align it, and add it to a running summation with increasing SNR. This calculation may take some time for data with many frames and many tiles-per-frame. A progress bar will be shown to indicate the frame being processed and the tile within that frame being extracted and realigned.

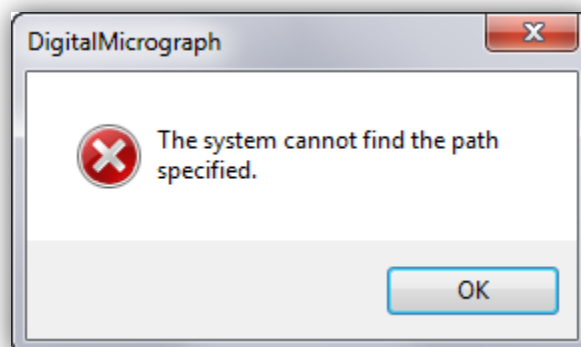


After the TM is complete across the series, the results will be displayed, and the user may choose to “Clean-up Motif Detection temporary data” from the menu and the temporary files will be deleted.

Troubleshooting

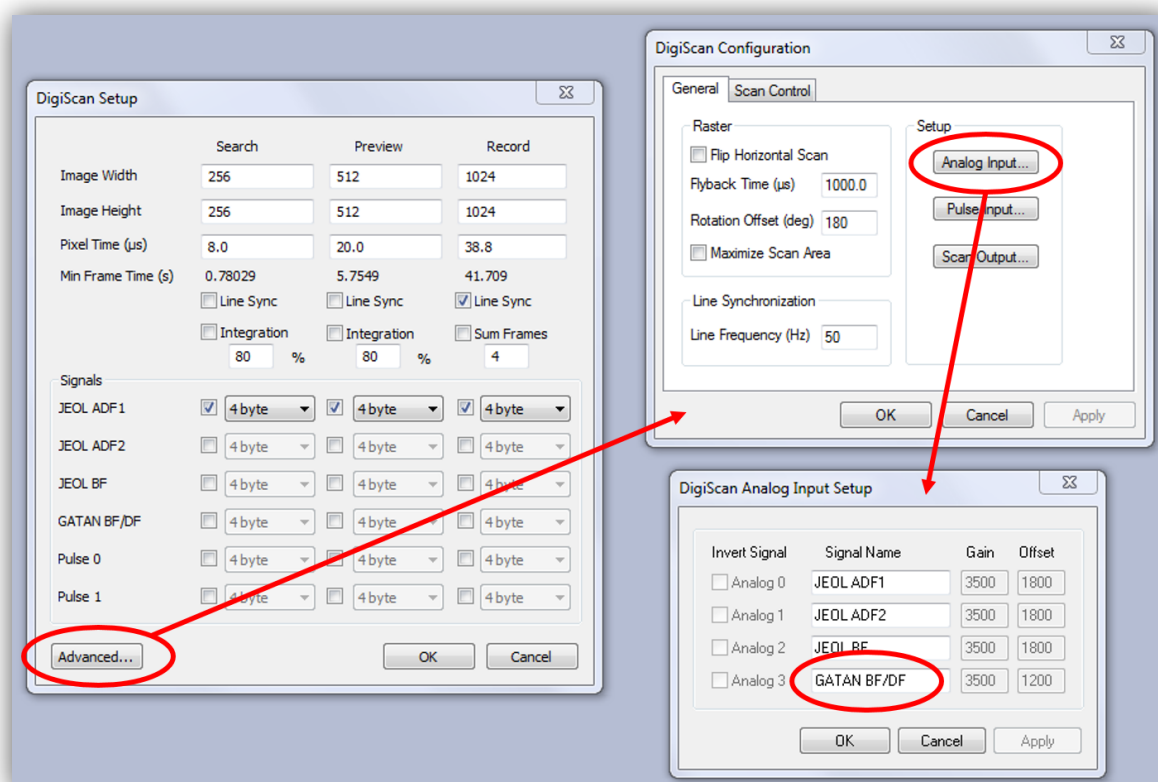
Auto-save function not working with some detector channels.

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

- [1] L. Jones, H. Yang, T. J. Pennycook, M. S. J. Marshall, S. Van Aert, N. D. Browning, M. R. Castell, and P. D. Nellist, *Smart Align—a New Tool for Robust Non-Rigid Registration of Scanning Microscope Data*, Adv. Struct. Chem. Imaging **1**, 8 (2015).
- [2] L. Jones and P. D. Nellist, *Identifying and Correcting Scan Noise and Drift in the Scanning Transmission Electron Microscope*, Microsc. Microanal. **19**, 1050 (2013).
- [3] C. M. O’Leary, B. Haas, C. T. Koch, P. D. Nellist, and L. Jones, *Increasing Spatial Fidelity and SNR of 4D-STEM Using Multi-Frame Data Fusion*, Microsc. Microanal. **1** (2021).
- [4] L. Jones, S. Wenner, M. Nord, P. H. Ninive, O. M. Løvvik, R. Holmestad, and P. D. Nellist, *Optimising Multi-Frame ADF-STEM for High-Precision Atomic-Resolution Strain Mapping*, Ultramicroscopy **179**, 57 (2017).
- [5] M. van Heel, M. Schatz, and E. Orlova, *Correlation Functions Revisited*, Ultramicroscopy **46**, 307 (1992).
- [6] C. D. Kuglin and D. C. Hines, *The Phase Correlation Image Alignment Method*, Proc. Int. Conf. Cybern. Soc. 1975 163 (1975).
- [7] M. Bosman and V. J. Keast, *Optimizing EELS Acquisition*, Ultramicroscopy **108**, 837 (2008).
- [8] V.-D. Hou, *Reduce Correlated Noise in EELS Spectrum with High Quality Dark Reference*, Microsc. Microanal. **15**, 226 (2009).