

Statistically Determined Spatial Drift (SDSD) Program Documentation

(v7.1 ; 19 August 2008)



Content

Copyright & Contact	2
Introduction	3
Setup.....	3
First run – a quick step by step drift correction of images	4
Second run – correcting data with stored drift information.....	6
A variant: Drift correction of 3D images	6
Drift correction of images with different size/calibration	7
Output Options	7
Filtering Options	7
Calculate Options	8
Expert mode	8
Filtering Options – Expert mode	9
Calculate Options – Expert mode.....	10
Acknowledgments.....	11

Copyright & Contact

Please remember that the program is *Citation Ware*. This means you may use it freely, but the authors have to be cited by the citation given below when the software is used for scientific work. A short note (email) to the authors giving the citation of the publication would also be appreciated. If other people contact you to spread the files, please forward their request to me and do not spread the files. This ensures that I can keep track on who is using the SDSD packages. Additionally, users will always get the most up-to-date version.

Automated Spatial Drift Correction for EFTEM Image Series
B.Schaffer , W.Grogger , G.Kothleitner
Ultramicroscopy Vol.102/1 (2004) pp.27-36

For questions, suggestions, and bug reports the authors can be contacted

by email: bschaffer@superstem.org
by mail: SuperSTEM Laboratory
STFC Daresbury
Keckwick Lane
Warrington WA4 4AD
United Kingdom

Web page: <http://www.superstem.dl.ac.uk/>

Please note, that the SDSD version 6 and higher make use of FFT-routines written by Christoph Koch. The routines are stored in the file *Transforms.dll* which can be downloaded in its newest version from the DM-Scripting-Database. http://www.felmi-zfe.tugraz.at/dm_scripts/

Introduction

Energy-filtering transmission electron microscopy (EFTEM) is a widely used technique in many areas of scientific research. Image contrast in energy energy-filtered images arises from specific scattering events such as the ionization of atoms. By combining a set of two or more images, relative sample thickness maps or elemental distribution maps can be easily created. It is also possible to acquire a whole series of energy-filtered images to do more complex data analysis.

However, whenever several images are combined to extract certain information, problems are introduced due to sample drift between the exposures. In order to obtain artifact-free information, this spatial drift has to be taken care of.

Manual alignment by overlaying and shifting the images to find the best overlap is usually very accurate but extremely time consuming for larger data sets. When large amounts of images are recorded in an EFTEM series, manual correction is no longer a reasonable option. Hence, automated routines have been developed that are mostly based on the cross-correlation (CC) algorithm.

A major problem with most CC-based routines is that they are not very robust, as a single CC often yields inaccurate or even incorrect results since a CC approach assumes the two images to be (almost) identical in content but only shifted with respect to each other. This, however, is clearly not the case for a series of energy-filtered images where elastic and inelastic contrast varies as a function of energy loss.

The Statistically Determined Spatial Drift (SDSD) correction program tries to overcome the problem by a three-step approach: First, a set of digital filters is used to enhance the relevant information (edges, features, etc.) while at the same time misleading information (noise, contrast changes due to different energy losses, etc.) are reduced. As a result, the images used for the drift correction become ‘comparable’ in contrast and therefore suitable for a CC algorithm. In a second step, the user is asked to ‘check’ the data and select the best region in the images for the drift correction. Areas with misleading contrast can therefore be avoided. Finally, the algorithm increases its robustness by cross-correlating *all* image pairs with each other and the obtained information is evaluated statistically. Outliers can thus be identified and excluded resulting in a much more accurate drift determination with even sub-pixel precision.

The SDSD routine has a user friendly interface which allows correction of both, series of EFTEM images and 3-dimensional image stacks (spectrum images) ‘by two clicks’. It is possible to correlate images with different scaling (nm/pixel) as well as different size (pixel-x-pixel). A range of parameterized digital filters is offered for enhancing image contrast and suppressing noise. On demand, the drift correction can be performed with sub-pixel precision. The final drift information can be stored as a (small) file for later use, so that it is not necessary to store both raw and drift-corrected data which can easily fill up disk space.

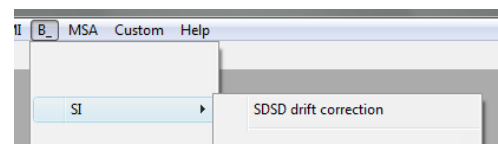
Setup

The program package includes three files:

Two DigitalMicrograph™ (DM) plug-in packages called *b_SDS.D* and *Transforms.dll* and this documentation PDF-file called *SDSD documentation.pdf*. To install the program, copy the two plug-in packages into the *PlugIns* folder of DM. The folder typically is located at:

```
C:\Program files\Gatan\DigitalMicrograph\PlugIns
```

After copying the file and starting DM, the SDSD drift correction is available from an additional *B_* menu. Please note that the SDSD packages were programmed and compiled in DigitalMicrograph v3.10.1 for GMS 1.5.1.

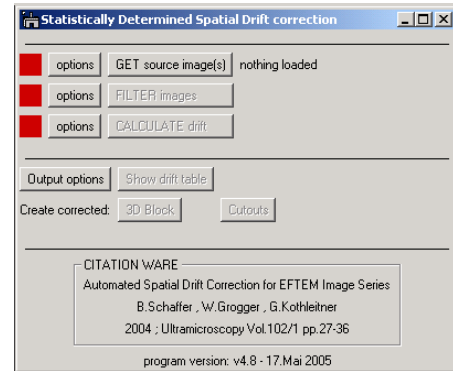


First run – a quick step by step drift correction of images

In this chapter we will quickly drift-correct a set of images. Please close all open images and then open all images which should be drift-corrected. Start the SDSD program from the $B_- \rightarrow SI$ menu.

When the program is launched, a floating window appears. This window is split in three sections, one for determining the drift, one for creating the drift-corrected output and one for citation¹ and program version information.

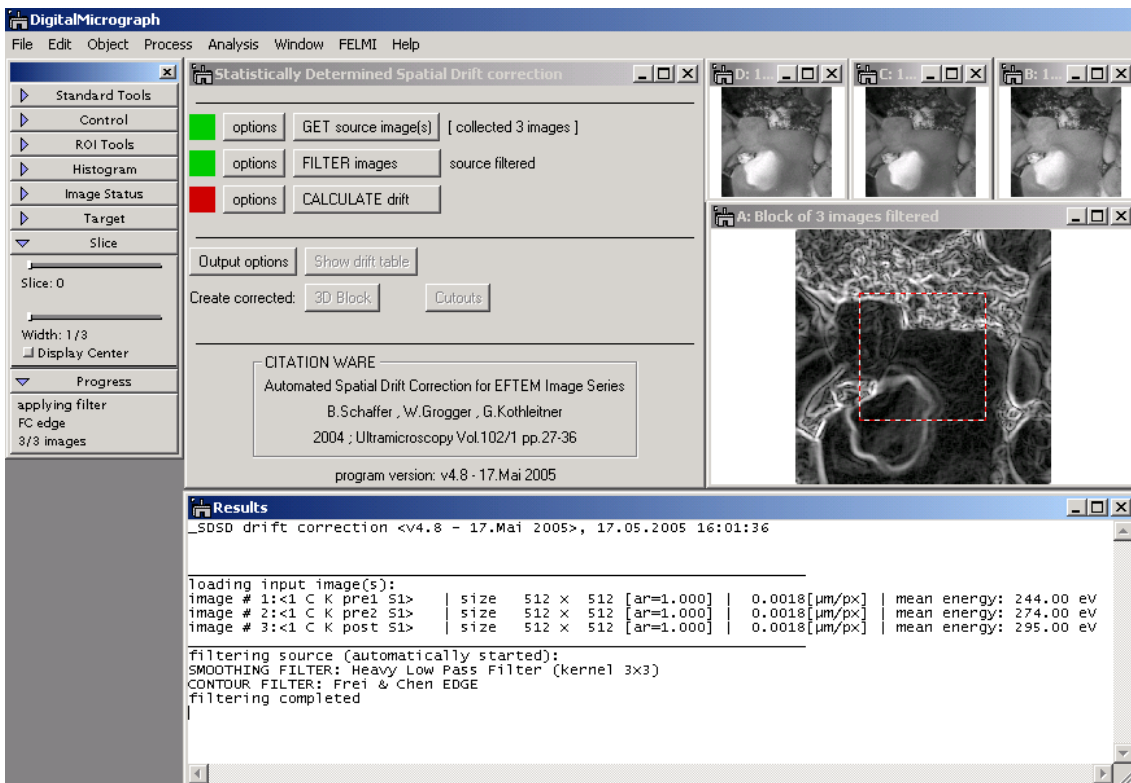
Each line in the drift determination stands for one of the necessary algorithm steps and consists of a colored status light, a push button, an options button and a status text line.



First load the images you want to correct by pressing the `GET source image(s)` button. This will load all currently open images into a stack of images. For standard settings, all images need to be of the same size in x and y and have to have the same calibration (nm/pixel). While loading the images, the result window will provide the image name, size, aspect ratio, calibration and – for EFTEM images – the mean energy loss value. The images are automatically sorted by their energy (prime criterion) and name (secondary criterion).

Now the next step (filtering) is automatically performed, because for standard settings ‘automatic filtering’ is enabled in the filter options. The filtering may take some seconds and the floating window ‘Progress’ will show you what the program is currently doing. Using the standard settings, the algorithm will first apply a ‘Heavy Low Pass’ filter of 3x3 kernel size and then the ‘Frei & Chen edge’ contour filter.

As a result you will be presented with a stack of filtered images with a selection in the middle. The first two status lamps will be green; the first status line will tell you how many files were collected and the second line that the filtering was okay.



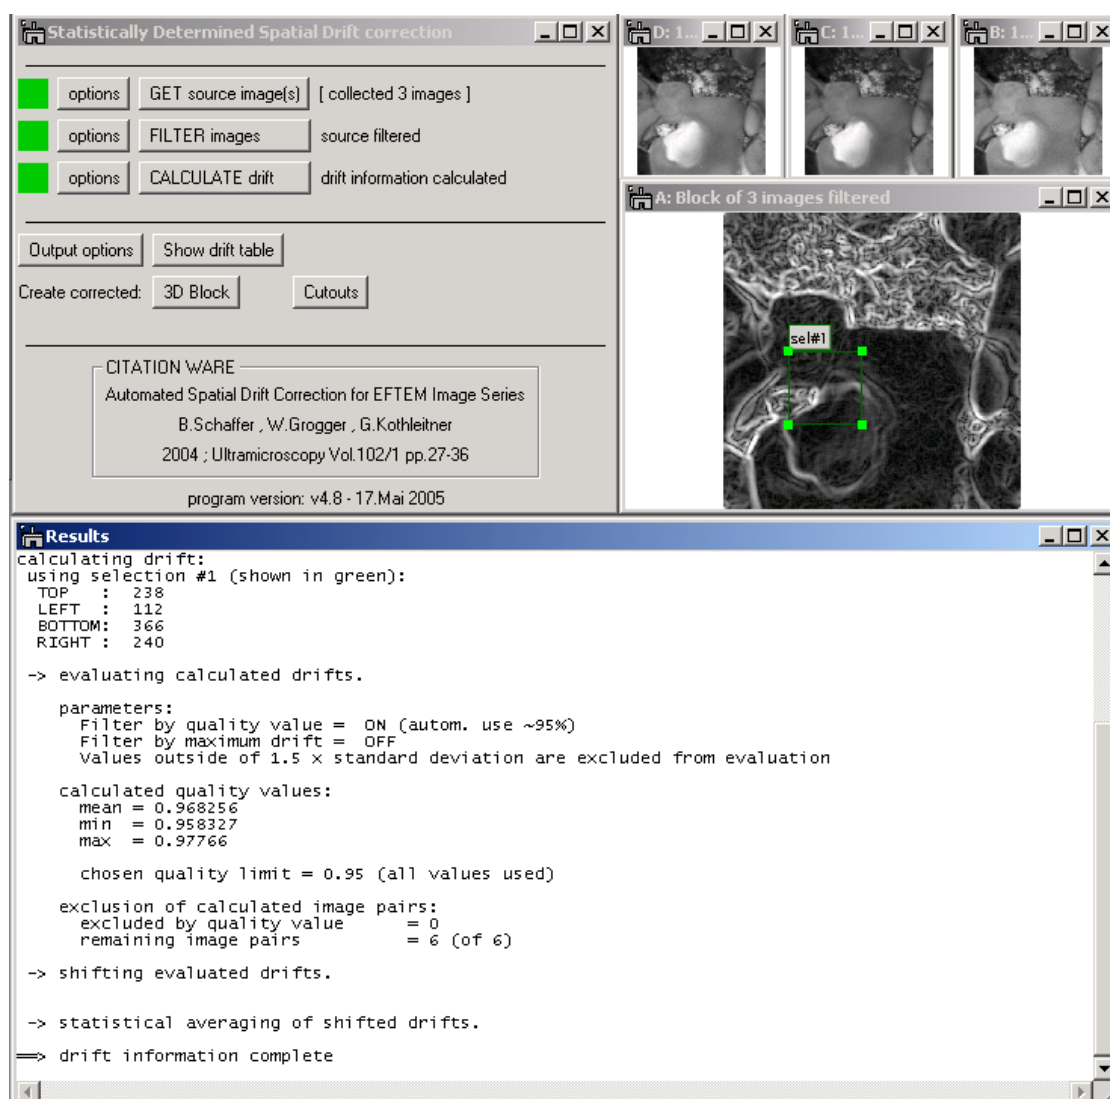
¹ Please remember that the program is Citation Ware. This means you may use and spread it freely, but that the authors have to be cited by the given citation when the software is used for scientific work.

In the next step you should use the 'Slice' tool of DM to quickly move through the stack of images and get a visual impression of the spatial drift. Then move (and resize) the selection to an appropriate region which will be used for the drift determination. Take care of the following points:

- A 'good' selection should be placed in a region which looks 'as similar as possible' in all images and holds (only) contour lines of sample features and no misleading diffraction contours. The more similar the selected regions look like, the better the drift determination will work.
- Smaller selections lead to faster drift calculation, but the selection must be at least twice as big as the maximum drift between two images!
- **The following limitation is valid, if the method "cross-correlation" is chosen. The standard setting "cross-correlation by C.Koch" does not have this restriction!**

Due to the FastFourierTransform involved, only selections with sizes of a power of two in each dimension are valid. You can force DM to create such selections by keeping the ALT button pressed during resizing the selection. If a selection of improper size is drawn, the program will only use a smaller, centered selection of proper $2^n \times 2^m$ size.

When you are satisfied with your selection, press the **CALCULATE drift** button. Now the program correlates all images to each other, performs the statistical analysis of the drift vectors and produces a final drift table. The 'Progress' window will tell you, what the program is currently doing and important facts are listed in the 'Results' window. When finished, the third status lamp will be green and the third status line will read 'drift information calculated'. The selection in the block will become green and labeled. It is of the size which was really used for calculation. The selection can not be moved or resized, but only deleted.



As the drift table is ready now, you can use the `SHOW drift table` button to present it as a spreadsheet. Each row stands for one image. The first two numbers are the drift (in pixels) for x and y direction, the third and the fourth number give the standard deviations of the two values, and the fifth number tells you, how many different correlations were finally accepted by the statistical treatment. The drift table can be saved as an image and later be used to produce drift-corrected output without redoing the drift determination. You can also manually alter some of the drift values if this is necessary.

	C0	C1	C2	C3	C4
R0	0.0	0.5	0.0	0.707107	2.0
R1	0.0	0.0	0.0	0.0	2.0
R2	0.0	-0.5	0.0	0.707107	2.0
	↑	↑	↑	↑	↑
	DRIFT in X	DRIFT in Y	Std. deviation of DRIFT in X	Std. deviation of DRIFT in Y	# of included correlation pairs

Finally, use the `3D Block` or the `Cutouts` button to produce the drift-corrected output.

When you produce *cutouts* they will be named with an additional 'AS' at the end, have the same calibration and image tags as the source and have an additional tag-tree called 'applied drift correction' holding information about the used source file and the selection used for the drift determination.

When you produce a *3D-Block* it will be named 'Block of ... images AS', has the same calibration as the *first* source-image and has a single tag-tree called 'applied drift correction' holding information about all used source files including their calibrations and tags.

Second run – correcting data with stored drift information

As explained in the previous chapter, you can display the drift information of a given set of images as a spreadsheet. This can be saved as a DM3 file. To use a stored drift for data correction, the following points have to be fulfilled:

- The drift information (spreadsheet) is open in DM
- All source images which the drift was determined from are open in DM
- The drift information (spreadsheet) is the *front-most* displayed image

If the `GET source images(s)` button is now pressed, the program loads the drift information and checks all source images. If all images are available, the first 'lamp' will become yellow and the last one green². The status lines will read: drift information loaded. You can now use the buttons `3D Block` and `Cutouts` to produce the drift-corrected output.

If not all source files are available, all status lamps will become red and the first status line tells the filename of the first missing image. Note, that the program uses unique 'image IDs' to indicate different images. Renaming the images does not influence this ID and will therefore have no influence on the SDS algorithm.

A variant: Drift correction of 3D images

Instead of using individual images as input, you can also drift-correct 3D images as they are produced by the EFTEM imaging spectrum method. (e.g. by the EFTEMSI plug-in of Gatan)

To load a 3D images simply make it the 'frontmost' displayed image by selecting it and then press the `GET source images(s)` button. The status line will now tell the image name, but all other steps are as

² This is done using the *Global Image IDs* provided by Gatan and described in its documentation. The necessary IDs are stored in the drift table as tags.

in the previous chapter, just that the cutouts will be named by the plane number in brackets while the 3D block will now be identical to the source with respect to tags and calibration.

Drift correction of images with different size/calibration

The SDSD program also allows drift corrections of images of different size. This may be important if you either have a series of images acquired with different binning or if you want to correlate already processed images like different elemental maps. There are three options in a drop down of the 'Collect' – Options Dialog:

- Image of different sizes: **do not collect**
This is the standard setting. When images of different size are encountered, the program stops without doing anything.
- Image of different sizes: **scale to biggest**
All images are scaled to one common size – the size of the biggest image. If the images are of different aspect ratio, the algorithm will not load the images. The “drift measurement” is performed on the scaled images, but the resulting cutouts will be scaled back to their original size. Use this setting for drift correction of differently binned images.
- Image of different sizes: **fill with 0**
All images are scaled to one common size – the size of the biggest image. Smaller images are placed in the top-left corner and the empty space is filled with zero value. Use this setting for drift correction of differently sized images of same magnification.

The standard options will not allow loading images of different calibration (nm/pixel) in order to prevent correlation of images with different magnification. However, sometimes one may want to override this setting by enabling the **allow different calibrations** checkbox. Images of different calibration will then produce a warning message in the 'Results' window but not stop the algorithm.

Output Options

There is currently only one setting in the 'Output' Options Dialog. You can activate/deactivate the *sub-pixel precision* correction. When enabled, the drift correction is performed with sub-pixel drift values by bilinear interpolation of the images. If it is disabled, given drifts of the drift information (spread sheet) are rounded when the drift is applied. In the *expert mode*, you can set the precision value to which should be rounded in case of sub-pixel precision. However, the default value of 0.01 is usually a good choice.

Filtering Options

The 'Filter' Options Dialog is split into four sections.

If the **automatic filtering** is enabled, the filtering is automatically performed after the source images have been loaded. If it is disabled, the **FILTER images** button has to be pressed manually.

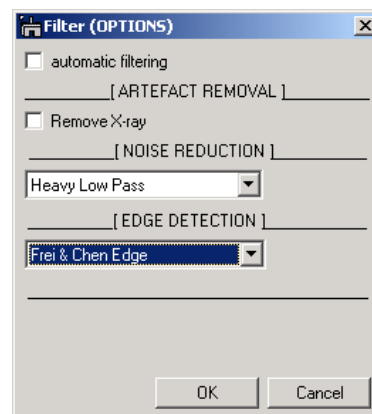
The **Remove X-ray** filter is part of Gatan's EFTEMSEI package.

If the package is installed, the menu command

EFTEMSEI/Artifact correction/Remove x-rays (GMS<1.5.1)

or *Volume/Remove X-rays* (GMS>=1.5.1)

becomes available and the SDSD program uses it to filter out disturbing x-ray artifacts. If the package is not available, no filter is applied but a warning is given in the results window.



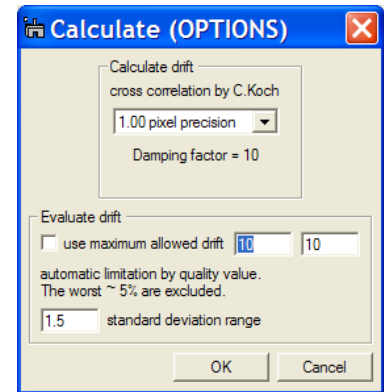
There are three noise-reduction filters offered. Two of them (**Butterworth** and **Hanning Window**) are based on frequency filtering and are much faster than the **Heavy Low Pass** or the **Kuwahara** filter which filter in spatial domain.

There are four edge-detection filters offered. The **Frei-Chen Edge** is somewhat more sensitive than the **Sobel** filter, but it is also slower. Both filters work in spatial domain while the **Butterworth** and the **Hanning Window** filters are based on frequency filtering.

Calculate Options

The ‘Calculate’ Options Dialog is split into two parts, one for determining the individual drifts of image pairs and one for statistic evaluation of all drift values.

In the first drop down menu you can select the *precision* of a single image pair drift calculation. Higher precision leads to much slower drift calculation as the images are scaled accordingly prior to cross correlation. Note, that due to statistics you can also achieve sub-pixel precision for the complete drift correction when large enough image series are used.



The second box “Evaluate drift” is used to select the exclusion criteria during the statistic evaluation process.

If the *use maximum allowed drift* checkbox is enabled, an individual drift vector is only allowed for further processing, if its value ([pixels]) in both X and Y is *less or equal* to the given numbers in the boxes to the right. (First is maximum allowed drift in X, second in Y.) After quickly checking and estimating a maximum drift by using the ‘Slice’ tool on the filtered stack of images you can enter some values and enable the checkbox. If the correlation of two images yields a bigger drift for some reason, this correlation is ignored.

The second line tells you about the current settings for the quality value exclusion criteria. You can only change these settings in the *expert mode*. The standard settings are an automated limitation which tries to exclude the 5% ‘worst’ drift correlations.

The *standard deviation range* determines which individual drifts are finally considered. All but the previously excluded drifts are used to calculate a mean and a standard deviation. Then, all values *outside* the mean plus/minus the given range in standard deviations are excluded and a new mean/standard deviation is calculated. This is continued until no individual drifts have to be excluded.

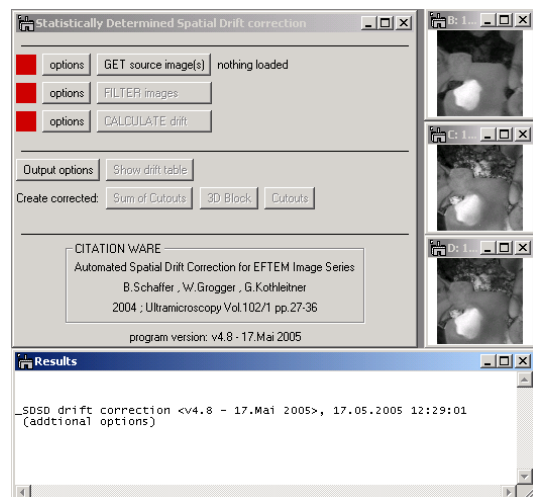
Expert mode

You can launch the SDSD program in *expert mode* by holding the [ALT] key pressed when selecting the menu entry “FELMI/SDSD drift correction”. You then have access to more detailed options. Note that changes in the options will be permanent and are also valid when the program is started anew in *standard mode*.

The *expert mode* also offers two additional output buttons. The first is called [Sum of Cutouts] and generates an image with summed intensity of all drift-corrected cutouts. Calibration and tags are taken from the first image of the series. This option can be useful for low-intensity acquisitions. Instead of one long exposure time several images of shorter time are acquired and then summed after drift correction as suggested by Aoyama et al.³

Additionally, the button [Show drift graph] is shown. This displays the calculated drift as XY-plot. If the button is pressed while the [SHIFT] key is pressed, the drift will be drawn as blue ROI within the source-block.

Another difference of the *expert mode* is, that right after “collecting” all images for drift by the [GET source image(s)] button, the [Show drift table] button becomes available. However, it will only present an empty drift table, but with all correct tags to be used. Such a “dummy drift table” can be used to enter drift values manually, or to combine values from different runs of SDSD or other drift correction schemes.



³ K. Aoyama, R. Matsumoto, Y. Komatsu, J. Electron. Microsc. 51 (2002) p257.

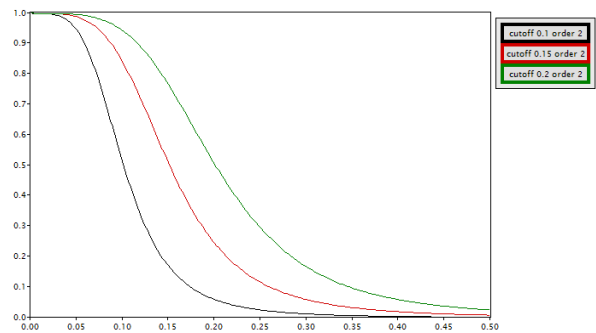
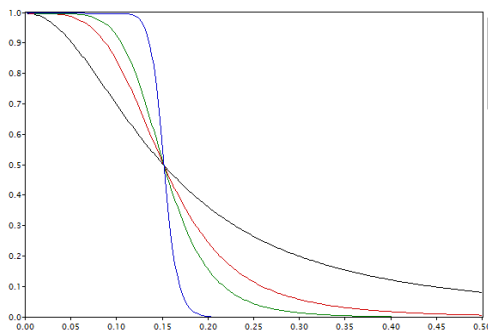
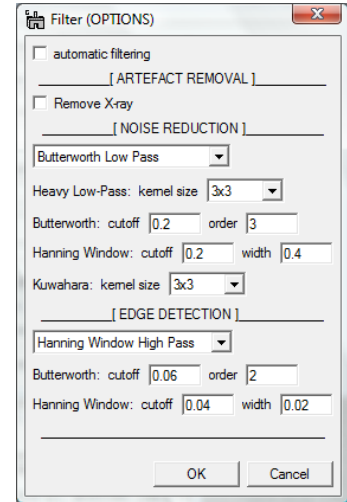
Filtering Options – Expert mode

In expert mode the parameters of the digital image filters can be customized:

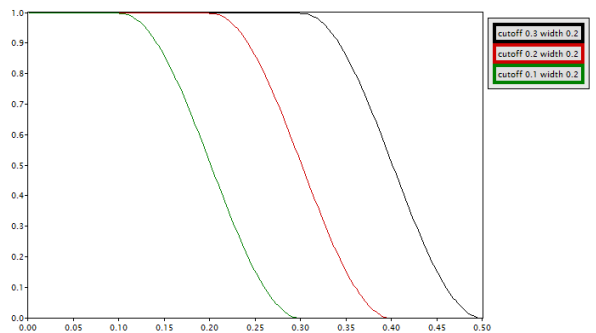
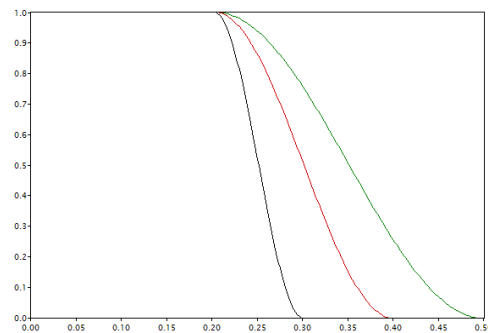
The *kernel size* of the *Heavy Low Pass* or the *Kuwahara* filter determines the strength of the filter. Bigger kernels create smoother images but slow down the filtering process.

The *Butterworth* filter is a damping curve in the frequency domain. A higher *order* will make the slope steeper; the *cutoff* frequency gives the point of inflexion. The modified *Hanning Window* filter is another damping curve in frequency domain. The *cutoff* frequency defines the point of inflexion of the damping curve, the *width* gives the transition with from 1 to 0 of the damping curve.

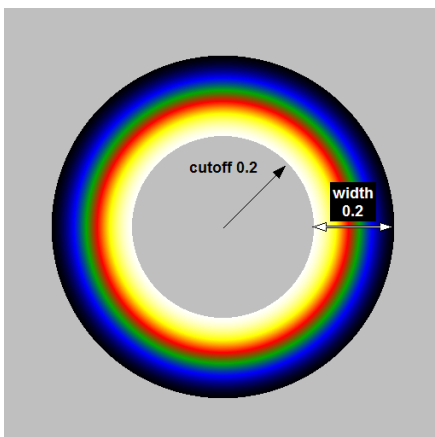
Both filters have a Low Pass (noise reduction) and a High Pass (contour enhancement) version. Those version are the inverse of each other. (HP+LP=1)



Butterworth Low Pass Filter with varying cutoff/order

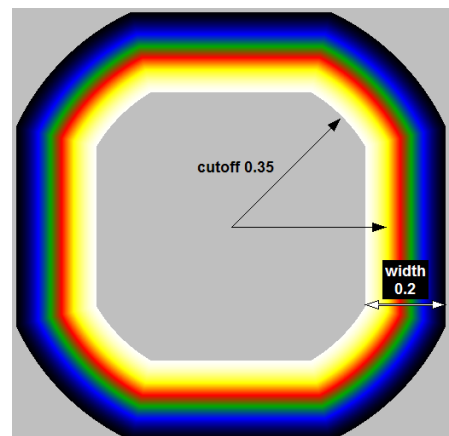


Modified Hanning Window Low Pass Filter with varying cutoff/order



← 2D version of →

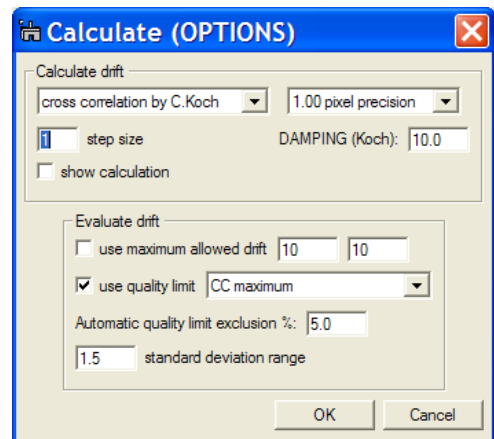
Modified Hanning Window Low Pass Filter. The transition width lies inside the Nyquist-Frequency range for all directions, if : width + cutoff > 0.5



Calculate Options – Expert mode

In expert mode the calculate drift section shows three additional elements:

The first drop down menu allows selecting the **method** used for the correlation of two images. Currently, cross correlation, cross-, phase- & mutual correlation by C.Koch are implemented⁴. The correlations by C.Koch use a different FourierTransform (stored in transforms.dll) to avoid the 2ⁿ restriction of standard FFTs. The difference between cross/mutal/phase correlation can be found in literature⁵. The standard setting “cross correlation by C.Koch” usually works best. The **damping coefficient** is used by the three routines by C. Koch. A value between 5-15 is recommended.



The *CUSTOM FUNCTION SDDS_Custom()* setting is special. Instead of using any inbuilt correlation method, the algorithm will search for a library function of the following prototype:

```
number SDDS_Custom(image &Img, image &Ref, number &driftx, number &drifty, number &qval)
```

If such a function is installed, it will be used to correlate the images *Img* and *Ref* with each other, returning the drift values in *driftx* and *drifty* and a quality value of the measurement in *qval*. This option can thus be used to incorporate your own drift measurement routines into the SDDS framework. Make sure that the *qval* value is between 0 and 1 with 1 being a “perfect measurement”. Ideally, the *qval* should increase linearly with the quality of the measurement.

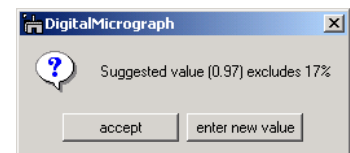
The **step size** determines how many images are correlated with each other. 1 means, that all images are compared, 2 that only every second image is used as reference, 3 that only every third is used as reference and so on. Normally, the value should stay at 1, but for large sets of images a less detailed but faster calculation may be sufficient.

If the **show calculation** checkbox is enabled, an image holding all individually calculated drift vectors is shown during and after drift determination⁶. This image can be saved and reused for different statistical evaluation of the same calculated individual drifts. Simply make it the ‘front most’ displayed image by selecting it and then press the `GET source images(s)` button. If all source images are found (see: ‘correcting data with stored drift information’), the evaluation is performed at once and the final drift is calculated.

The ‘Evaluate drift’ section shows the following additional options:

The limitation by the quality value criteria can be disabled with the checkbox **use quality limit**. The associated drop-down menu currently allows only the cross-correlation maximum to be used as a quality indicator⁷.

The field **Automatic quality limit exclusion %** defines roughly the percentage of ‘worst’ measurements which should be excluded from the evaluation. Note, that in *expert mode* you will always be presented with a dialog during the evaluation process, in which you are asked to approve the chosen quality limit.



⁴ If you have developed your own scripts/methods to determine the drift between two images you might want to contact the authors so that the methods can be implemented into the program.

⁵ R.R. Meyer et al, Ultramicroscopy 92 (2002), 89-109

⁶ The image consists of 5 planes of size “Nr.of.Img x Nr.of.Img”. Pixels at the positions (n,m) contain the information of the drift correlation of the nth and the mth image. The first two planes give the drift in x and y, respectively. The third plane shows the maximum value of the cross-correlation (= the ‘quality value’). The fourth and fifth planes contain information for the algorithm. There is more information in the tags of the image.

⁷ If you have developed your own scripts/methods to classify the “quality” of a determined drift of two images you might want to contact the authors so that the methods can be implemented into the program.

The value in the *exclusion %* field is only used as a starting point. In *beginners mode* the thereof estimated quality limit is automatically approved.

Acknowledgments

The authors would like to thank David Mitchel, Robin Harmon and Jürgen Plitzko for their help concerning the digital image filters and scripting in Digital Micrograph. This program also uses the transforms.dll by Christoph Koch, which allows correlations & FourierTransforms to be calculated for images of arbitrary size instead of $2^n \times 2^m$ dimensionality only.