ThunderSTORM Manual v.1.0

Ann Wheeler, IGMM + ESRIC

ThunderSTORM is a complex image analysis package. For more information please see Ovesny, Hagen et al *Bioinformatics (2014) doi: 10.1093/bioinformatics/btu202*

At time of writing the plugin and other information about the software was being held in a GitHub repository here:

https://github.com/zitmen/thunderstorm

Whlist Thunderstorm is powerful and state of the art there is considerable functionality involved which this manual seeks to address in part:

Contents:

Introductory notes

Thunderstorm is a comprehensive platform for analysis of SMLM data. Within it are implemented several different localisation method and processing tools which enable an end user to try several different options when processing the data.

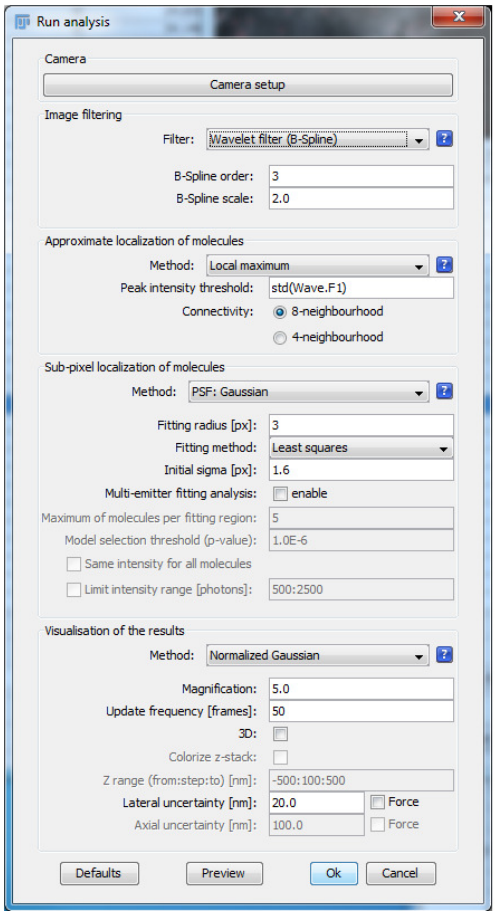

However having searched for an online manual documenting the tool information was slightly lagging. Here I write up a comprehensive manual featuring a short introductory section detailing how to get reasonable data from Thunderstorm quickly. Then a more detailed glossary discussing all of the options and the effect on processed data.

Open Thunderstorm and set the basic camera settings

(This is found in Plugins, Run, ThuderSTORM, Run analysis)

This Gui appears. Each part of it needs to be optimised. Suggestions for settings are made and mandatory ones –for the IGMM N-STORM system - are shown in black.

The Gui works as follows:

Section 1. The settings for the camera used

Section 2. Identification of photoswitching molecules to pixel resolution

Section 3. Sub Pixel localisation of molecule.

The bottom section is for on the fly visualisation of the results. This is very useful to see how the selected parameters are working.

Section 4: The mechanism of visualisation of results

Thunderstorm visualises data 'on the fly' to give an idea of what the final image will look like.

1. Guidelines for the choice of parameters

1.1. Camera Setup

Most of these parameters are fixed as they are determined by the

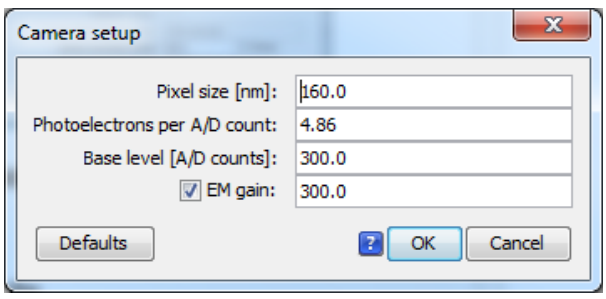

The Defaults for the STORM system are as follows: The system's camera is a – iXon897**. EM Gain. This should be set to 300.** But if a different value was used in the experiment it must be entered here. This assumes the camera is run at 17MHZ and has a gain of 3.

1.2. Approximate localisation of molecules. Our experiments indicate that the local maximum approach with 8-connected neighborhoods provides the highest F1-score compared to other methods.

Peak intensity threshold 2D The suggested threshold applied to the filtered images, in the case of the wavelet filter, usually ranges from 0.5 to 2 times the standard deviation of the 1st wavelet level, e.g., 1*std(Wave.F1). Such a value is recommended by Izeddin et al. *[1]* and works well for typical SMLM data. Increasing the threshold value will lead to less false positive detections at the expense of more missed molecules and vice versa.

Peak intensity threshold 3D: As the signal to noise ratio is usually higher in the 3D calibration data with fluorescent beads, users should set the threshold, in the case of the wavelet filter, to 5 to 8 times the standard deviation of the 1st wavelet level, e.g., 6*std(Wave.F1). The rest of the settings are the same as in 2D data analysis. Use the Preview button to see detections of the calibration beads with the current settings.

1.3. Subpixel Localisation of molecules method: It has been shown *[4, 3]* that the Gaussian function provides a very good approximation of the real PSF of a microscope. This is mainly due to pixelation effects and the presence of noise, which makes the difference between the Gaussian function and the real PSF negligible. The advantage of Gaussian PSF models are their simplicity, robustness, and computational efficiency.

Subpixel Localisation Fitting Method: Fitting PSF models by *maximum likelihood estimation* generally gives slightly better results than fitting by least square methods, particularly when the photon counts are low. Maximum likelihood estimation requires correct setup of the camera parameters (photoelectrons per A/D count and the digitizer offset level). The recommended PSF model, which generally works well, is the integrated Gaussian. The initial size of sigma can be found by running ThunderSTORM on few images of the data sequence. A histogram of the fitted sizes of sigma (in pixels) can help to find the initial value. The size of the fitting radius should be an integer number close to 3*sigma

1.4. Visualisation: The proposed average shifted histogram method provides similar results as Gaussian rendering with a fixed sigma, but is orders of magnitude faster. Both methods result in a molecular density map. The Gaussian rendering approach can further show the calculated localization accuracy for each molecule.

1.5 QA of analysis: Its possible to check if your image analysis algorithm has done a good enough job by using simulated data. The Precision, Recall and Jaccard index and RMSerror can be calculated from this and can be used in final analysis. This is important for determining how good of a job the parameters are doing. If too many false positive / negative localisation are found in your data it will be invalid. There are no specific parameters here since every antibody, cell and tissue type is different. Reconstruction parameters will be specific to individuals experiments. See **section 7**

2 Image Filtering

This is used to remove any residual background / camera noise in the images and enhance the features (e.g. photoswitching events). There are several options available which are either low pass or convolution kernals:

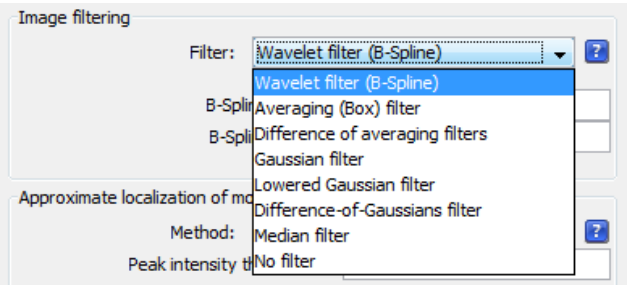

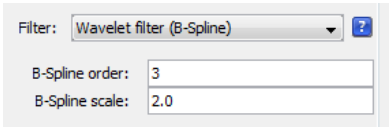

2.1 Wavelet filter (B-Spline): The wavelet transform is commonly used in modern signal-processing applications. This type of wavelet filter applies the *Ã trous* algorithm *[2, 1]* which is an undecimated scheme in which the filter responses are up-sampled, thereby inserting holes (*trous* in French) between the filter coefficients. The output of each filter level, therefore, contains the same number of samples as the input.

The 'Maths lite' explanation is that Wavelet filter denoises the image using a convolution kernel which is based on a polynomial function known as a Basis or B-Spline. The order in the box determines the order of the polynomial function and the scale is the scaling factor. These are the values suggested by Izzedin et al. *My suggestion is not to change these unless you are very convinced of what the mathematical implications will be.*

2.2 Averaging filter: This is the same as smoothing the image, it's a low pass filter. Using it results in an image in which the value of each pixel is equal to the average intensity of pixels from the input image in a given neighborhood. Filtering can be defined as a convolution of the image with a kernel given by a matrix containing equal entries that sum to one, thus

2.3 Difference of Averaging filters: Here the image is averaged using two filters – with differing and specified kernel sizes and the result is subtracted. This creates a type of band pass filter. Here the size of the kernel of the second filter needs to be larger than the first filter or the maths won't work

2.4 Gaussian filter: This is a straight forwards Gaussian Blur filter, which is already implemented in Image J and can be used for denoising images. The Sigma (i.e) error is set at 1.6 pixels but can be adjusted. It works on the assumption that the emission of a single molecule will be best represented by a Gaussian function. So if the PALM / STORM data are Gaussian functions this will be a standard method.

2.5 Lowered Gaussian Filter: This is the method implemented in DaoSTORM/ DaoPhot. It's a band pass filter. The *convolution kernel* is based on the *Gaussian kernel* which has been lowered*to have the sum of all its entries equal to zero,

The Maths lite explanation is that this method works by Gaussian filtering the image and then subtracting the mean value of all the elements from the result. The sigma value here is the same as in the Gaussian filter.

2.6 Difference of Gaussian filters: Here the image is averaged using two Gaussian filters – with differing and specified kernel sizes and the result is subtracted. This creates a type of band pass filter. Here the size of the kernel of the second filter needs to be larger than the first filter or the maths won't work

2.7 Median filter: A median filter is a non-linear image filter often used to remove impulse noise (i.e., salt and pepper noise). The value of each pixel in the resulting image is equal to the median of the intensity values from the original image in a given neighbourhood. If the neighbourhood mask has an odd number of entries, the median is the middle element of all numerically sorted entries. For an even number of entries, the median is calculated as the mean of the two entries in the middle. *Users need to input the kernel size and the type of neighbourhood mask*, which can be a cross (4-connected) or box (8-connected) pattern. This is the same as the one implemented in ImageJ.

2.8 No Filter: The option if the image does not require denoising

Now the image is denoised the molecules can be localised.

3. Approximate localisation of Molecules menu

This give a localisation 'per pixel' basically the software determines whether a photoswitching event which is work resolving to sub pixel event has occurred. Finding approximate positions of molecules

Finding approximate positions of the molecules in the input images follows the image filtering and feature enhancement step. ThunderSTORM offers three algorithms for this purpose: detection of local intensity maxima, non-maximum suppression, or calculation of the centroid of connected components of segmented objects. All three methods need a user-specified threshold defining an intensity below which possible molecular detections will be rejected.

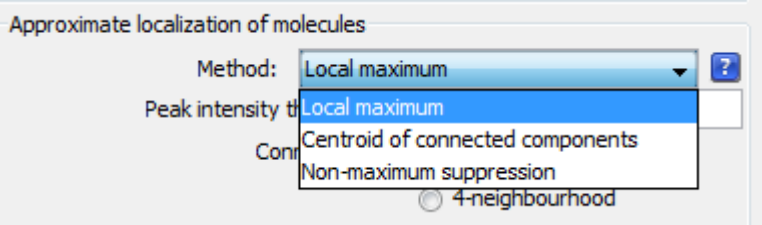

3.1 Setting the Threshold: (THIS IS IMPORTANT)

When finding the approximate position of molecules, choosing the right threshold value is important for obtaining good results, because the threshold influences the number of missed molecules (false negatives) and the number of erroneous detections of non-existing molecules (false positive detections). ThunderSTORM uses a single-valued intensity threshold which is updated for every raw input image **I** and applied to the filtered image **F.**

The threshold value can be specified by users as an expression combining mathematical functions and operators with variables based on the current raw or filtered image. This is a powerful option, because users can specify the threshold value systematically for unknown input images, in which the global intensity may slowly fluctuate over time.

3.2a Local Maximum

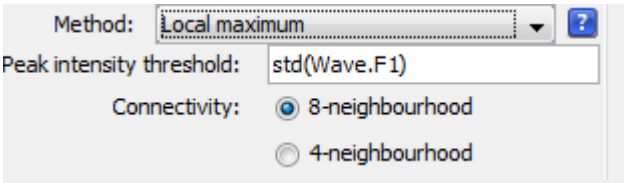

This method passes through each pixel of the filtered image **(F)** and determines if the pixel intensity F(x,y) is greater than a user-specified threshold and at the same time greater than or equal to all values within a 4- or 8-connected neighborhood. If this condition holds, the point is accepted as a candidate molecule for further processing.

Connectivity – This is the neighbourhood pixels which are searched when localising a pixel (x,y)

3.2b Centroid of connected components.

Here the position of each molecule is determined by computing the centroid of the relevant object using the formula below. It is possible to better separate the objects using watershed segmentation

if desired. The Watershed segmentation is based on the imageJ Find Maxima algorithm. This method is proposed by Izzedin et al (2012)

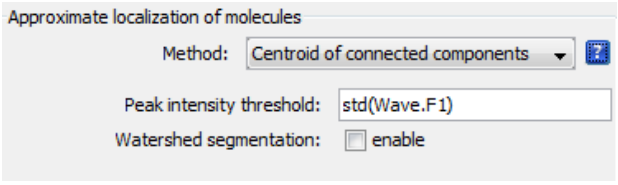

$$\tilde{x}_p = \frac{1}{c_p}\sum_{i=1}^{c_p} x_{i,p}\,, \quad \tilde{y}_p = \frac{1}{c_p}\sum_{i=1}^{c_p} y_{i,p}\,.$$

Here p indexes the objects, c_p is the number of elements in each object, and $x_{i,p}$ and $y_{i,p}$ are the integer pixel coordinates within each object. Segmented objects are determined as connected components by an algorithm based on a breadth-first search [2]. The watershed algorithm is based on the ImageJ ''Find Maxima'' function.

3.2c Non-Maximum suppression:

Essentially, a binary image is built with all non-maximum pixels set to zero and with all local maxima set to one. The maxima pixels are found by dilating the filtered image by the specified radius and finding maxima within that. See Sonka and Boyle (2007) for Mathematical proof.

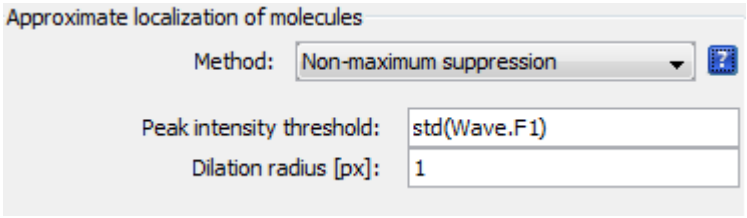

3.3 3D imaging

Approximate localization of molecules and threshold selection: As the signal to noise ratio is usually higher in the 3D calibration data with fluorescent beads, users should set the threshold, in the case of the wavelet filter, to 5 to 8 times the standard deviation of the 1st wavelet level, e.g., 6*std(Wave.F1). The rest of the settings are the same as in 2D data analysis. Use the Preview button to see detections of the calibration beads with the current settings.

4 Sub pixel localisation

Sub-pixel localization of single molecules with an accuracy below the diffraction limit is the basis of SMLM methods. ThunderSTORM supports the following localization methods: calculation of the centroid of the local neighborhood *[1]*, the radial symmetry method *[4]*, and fitting of point-spread function models by (weighted) least-squares methods or by maximum likelihood estimation *[2, 3]*. Users may also choose not to use any of the methods, thereby using the approximate positions of the molecules determined in the previous step. Note that the choice of method for sub-pixel localization of molecules is independent of the pre-processing methods described above.

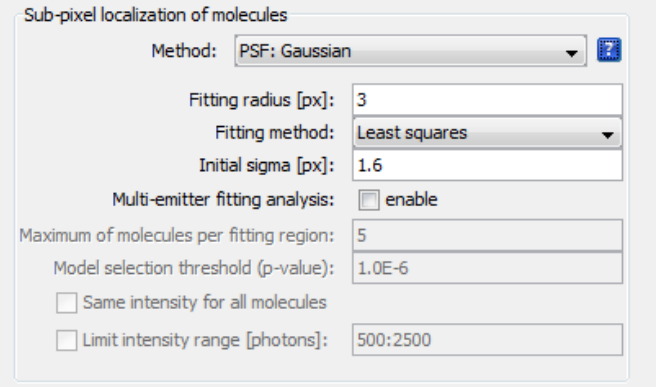

4.1 PSF: Integrated Gaussian (best one probably) The integrated form of a symmetric two-dimensional Gaussian function can be used to help to take into account the discrete nature of pixels present in digital cameras *[2, 1]*. Assuming a uniform distribution of pixels with unit size, a single molecule intensity profile can be expressed as:

$$\mathrm{PSF}_{\mathrm{IG}}\left(x, y \mid \boldsymbol{\theta}\right) = \theta_N E_x E_y + \theta_b ,$$

where $\mathrm{PSF}_{\mathrm{IG}}(x,y \mid \Theta)$ gives the expected photon count at the integer pixel position (x,y) for a vector of parameters Θ. The Vector parameter Theta accounts for the subpixel localisations (E_x+E_y), where E is the expected behaviour the camera pixel the Number of photons counted (N) and Background (b)

This method takes into account the discreet nature of pixels in scientific cameras the Gaussian function on its own doesn't do this

4.1aPSF: Gaussian Here an approximation of the the real PSF is fitted to each identified localisation . A 2D Gaussian function is used for this. This differs from the Gaussian fit as it doesn't

For both Gaussian methods The radius of fitting and method can be adjusted. (Fitting radius would be determined by expected Abbe limited PSF size and what would be mathematically reasonable).

The Gaussian fitting is as per this formula. The initial sigma should be 1.6 for our system as the initial – Abbe limited, molecule size would be a bit more than the 160nm pixel.

For fitting of point-spread function models (i.e. Gaussian methods) the Fitting Method needs to be specified. This + Multiple emitter fitting is discussed below.

4.2 PSF: Elliptical Gaussian (3D Astigmatism) 3D SMLM imaging can be performed by introducing a weak cylindrical lens into the imaging path to create slight astigmatism in the image [1]. This results in images of molecules with different ellipticity depending on their axial position. When a molecule is in focus, its image appears round. If the molecule is slightly above or below the focal plane, its image appears ellipsoidal. Calibration of the imaging system is needed to determine

the orientation of the imaged ellipsoid (the camera chip might not be aligned with cylindrical lens) and the relationships between the axial position and ellipticity of the imaged molecules.

This method was authored by Bo Huang

4.3 Radial Symmetry: This algorithm finds the sub-pixel position of a molecule by determining the point with maximal radial symmetry in the data as described in *[1]*. The general idea is to find the origin of radial symmetry (i.e., the center of a molecule) as the point with the minimum distance to gradient-oriented lines passing through all data points. The calculation of each molecular position is very fast due to an analytical solution, but the algorithm *does not estimate the intensity or imaged size of a molecule*. Radial symmetry is a robust feature in SMLM data, making the algorithm resistant to noise. (This method was developed in 2012 by Parthasarathy)

4.4 Centroid of local neighborhood: Calculation of the centroid in a local image neighborhood is a *very fast method* for sub-pixel molecule localization and is used in *[1]*. The *algorithm does not estimate the intensity or imaged size of molecules, and is sensitive to noise*. The main idea is simply to calculate the mean pixel positions weighted by the intensity of the image data.

If you just want a quick a dirty look at your data this is a good method. But it doesn't give the best localisations. This method is also known as QuickPALM

4.5 No Estimator: This doesn't bother doing subpixel localisation, only the list of approximate positions is returned

Method of subpixel localisation for Gaussian Fits

4.6.1 Least squares fitting: To approximate the data with a point-spread function, least-squares methods [4, 1, 5] are employed to minimize the sum of (weighted) squared residuals.

The residual value for a molecule (x,y) is the difference between the observed or recorded Photon Intensity **I** (x,y) and the value expecte from the approximated Gaussian PSF ($PSF_{IG}(x,y|\Theta)$)

$$\chi^2(\boldsymbol{\theta} \mid \mathcal{D}) = \sum_{x,y \in \mathcal{D}} w\left(\tilde{I}(x,y) - \mathrm{PSF}(x,y \mid \boldsymbol{\theta})\right)^2.$$

4.6.2. Weighted Least Squares fitting. As above but here the residuals are weighted by w = 1 / I (x,y) taking into account the uncertinaty in the number of detected photons (This is a better metric for fitting for EMCCD cameras)

FOR BOTH OF THESE FITTING METHODS THE USER HAS TO CHOOSE THE STARTING POINT AS THE APPROXIMATE MOLECULAR WIDTH OF AN FWHM ABBE GAUSSIAN FUNCTION (probably can measure from image)

4.6.3 Maximum likelihood fitting: Requires correct setup of the camera parameters

This approach assumes that the number of photons collected by a single camera pixel follows the Poisson distribution. The equation supposes that samples are drawn independently from the Poisson distribution, with the expected photon count given by the gaussian point-spread function model, and the observed photon count given by the image intensity expressed in photons (I (x,y). It then iterates through the parameters. It starts optimising by looking at the difference between the maximum and minimum intensity values for the molecules localised in the Approximate Localisation step and using the **user defined value** for the resolution (which is the Abbe resolution) but expressed in pixels. It eventually computers the subpixel fitting which has the maximum statistical likelihood of being correct.

4.7 Constraining parameters of PSF models – assumptions made

The Levenberg-Marquardt algorithm and the Nelder-Mead method used above search for values of the Gaussian parameters over an infinite interval. The optimization process can therefore converge to a solution with negative values which is impossible for variables corresponding to image intensity or to the standard deviation of a Gaussian PSF. We therefore limit the interval of possible values by transforming the relevant parameters so that the modulus is displayed. The optimization process is still unconstrained but will result in positive PSF parameters. Such a transformation also improves the stability of the fit.

4.8 Multiple Emitter Fitting

High spatial densities of activated molecules can result in what Mathematicians refer to as a crowded field problem, in which single molecules are not adequately resolved (and what Biologists refer to as a normal situation).

Since super-resolution analysis methods were invented by Physicists and Mathematicians, with scant knowledge of biology, their methods were devised to fit only single molecules. This is fine if the structures you work on are punctate. But a problem if they are abundant or filamentous.

To help solve this problem, ThunderSTORM uses a multiple-emitter fitting analysis (MFA) approach similar to the algorithm described by Bo Huang.

4.8.1 Guidelines for the choice of parameters

Multiple emitter fitting analysis (MFA) can be used in high-density data to estimate the number of molecules detected as a single blob. We recommend setting 3 to 5 molecules per fitting region. Unfortunately, multiple emitter analysis is a computationally costly method, and may run quite slowly depending on the user-specified maximum number of molecules within the fitting area. See 4.4 for why.

The stability of the algorithm is improved if the molecular brightness is limited to a realistic range (perhaps 300 to 5000 photons), depending on the sample. This range can be estimated by ThunderSTORM by processing a few frames of the data and evaluating the number of detected photons from each molecule using the histogram function (this relies on correct entry of the camera parameters).

The algorithm can also be more stable by forcing all molecules to have the same intensity. Which works at the expense of losing intensity data from the output. Note that the size of the fitting radius might need to be increased slightly to accommodate larger blobs.

4.8.2 Mathematics underpinning the multiple emitter fitting

The multiple-emitter fitting analysis approach uses a PSF defined by a model

$$\mathrm{PSF}_N\left(x, y \mid \boldsymbol{\Theta}\right) = \sum_{i=1}^{N} \mathrm{PSF}\left(x, y \mid \boldsymbol{\theta}_i\right)$$

where a finite number of molecules are allowed in the fitting region (N), and the equation contains parameters describing position and shape of the imaged molecules modelled by the PSF (Θ).

The fitting of multiple-emitter models to the raw data proceeds according to the following algorithm. First, the algorithm fits (a single molecule model) with an initial molecular position. The fitted PSF is subtracted from the raw data and the position of the maximum intensity value in the residual image is taken as an approximate position of a second molecule. The fitting is now repeated on the raw data with (a model containing two molecules) and with the initial positions estimated in the previous steps. The result of the fit is subtracted from the raw data to find an approximate position of a third molecule in the residual image. This routine is repeated until the maximum number of molecules allowed in the fitting region is reached. Note that initial positions of the molecules are further adjusted, during the multiple-emitter fitting analysis, by a˜Push&Pull" process *[2]*. To find the optimal number of molecules, statistical tests are required, see Section Model selection.

Users can specify the size of the fitting region, the maximum number of molecules allowed in one fitting region, the type of PSF, and the fitting method (least-squares methods or maximum likelihood estimation). Optionally, users can constrain the multiple-emitter fitting algorithm such that all fitted molecules have the same intensity or an intensity in a given range. The background offset is constrained to the same intensity for all fitted molecules.

4.9 Model selection threshold

Because a model with more parameters will always be able to fit the data at least as well as a model with fewer parameters *[3, 1]*, statistical tests are required to determine whether the more complex model provides a significantly better fit of the underlying data. Statistical tests are usually based on pair-wise model comparison. Here a fit by $PSF_{(1\ mol)}$ is compared with a fit by $PSF_{(2\ mols)}$, the better of the two is compared with a fit by $PSF_{(3\ mols)}$, etc. Pair-wise comparisons are based on an F-test *[3, 1]* or on a log-likelihood ratio test *[3, 2]* as described below.

F-test for least squares multi emitter fitting: An F-test *[3, 1]* arises in the case of fitting by least squares methods, when we need to compare significance of the fit between two models, where one model (the null model) is a special case of the other (the alternative model) for some choice of parameters.

The null hypothesis assumes that the alternative model does not provide a significantly better fit than the null model. The F-test statistics computed in Equation (2) has an F-distribution with n

degrees of freedom. The null hypothesis is rejected if the value of F computed from the data is greater than the critical value of the F distribution for a user-specified -value.

Log-likelihood ratio test for maximum likelihood multi emitter fitting:To compare between the fits of two models, in the case of fitting by the maximum likelihood estimation method, we use a model selection criteria based on a log-likelihood ratio test *[3, 2]*. This Assumes that one model (the null model) is a special case of the other (the alternative model) for some choice of parameters.

The probability distribution of the log-likelihood ratio computed in Equation (3), assuming the null hypothesis that the alternative model does not provide a significantly better fit than the null model, can be approximated, This approximation is usually valid even for small sample sizes *[3]*. The null hypothesis is rejected if the log-likelihood ratio computed from the data is greater than the critical value of the distribution for some value specified by the users.

5. Visualisation of Results

Visualization (rendering) of SMLM data involves creation of a new super-resolution image based on the coordinates of the localized molecules. ThunderSTORM offers several methods, all of which support both two-dimensional visualization and slice-by-slice three-dimensional visualization. In the three-dimensional case, each image slice contains a visualization of molecules with axial positions in the user-specified range. The desired magnification ratio of the new super-resolution image to the original image is user-specified. The visualisation of results occurs on the fly, the magnification of the visualisation and the update frequency can be changed. Reducing update frequency can speed processing

5.1 Averages shifted histograms: This visualization algorithm uses a density estimation approach based on averaged shifted histograms *[1]*. It works quickly.

In the one-dimensional case, this method works by averaging histograms with the same bin width , but with the origin of each histogram shifted by a given quotient (= width of the bin / number of histograms) from the previous histogram.

In our implementation, the width of the histogram bin is determined the number of user specified shifts multiplied by the super-resolution pixel size in the final image. The number of shifts in the lateral and axial directions can be specified independently.

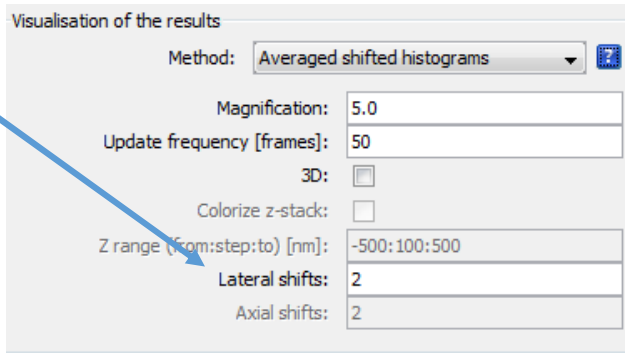

For 3D visualisation check the 3D box, and specify the Z range acquired. A colourised Z stack can help with depth interpretation.

5.2 Scatter plot: Scatter plot visualization [1] is the simplest method, and does not usually provide high quality results. A simple binary image is created with pixel intensity values set to one at locations corresponding to molecular positions. All other pixel intensity values are set to zero. This method is fast but does not reflect the density of molecules.

5.3 Gaussian rendering This method draws a normalized symmetric 2D or 3D Gaussian function integrated over the voxel volume for every localized molecule, with a standard deviation equal to the computed, or user-specified localization uncertainty. The visualized molecules are added

sequentially to the final super-resolution images. The contribution of one molecule to the voxel intensity at the integer position is related to the intensity, the axial and lateral localisation uncertainty. It is possible to force all molecules to have the same lateral (and Axial)

5.4 Histogram: Histograms are used to estimate the density of data by counting the number of observations that fall into each of the bins. In our case, a two-dimensional histogram of molecular positions is created with the bin size corresponding to the pixel size of the final super-resolution image *[1]*. Thus, for every localized molecule, the bin value (i.e., the image brightness) at the corresponding molecular positions is incremented by one.

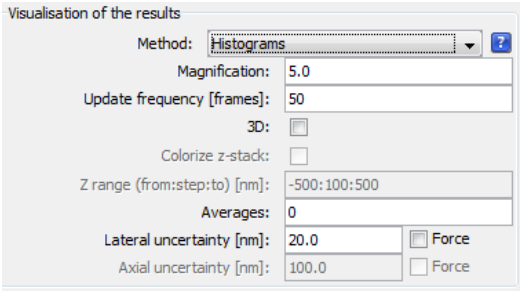

The histogram visualization optionally supports jittering*[2]*. When enabled, aÂ random number drawn from the normal distribution, with a standard deviation equal to the computed (or user-specified) localization uncertainty, is added to the coordinates of every molecular position before creating the histogram. This step is applied multiple times and all generated histograms are averaged together. As the number of jitters increases, the final image approaches the result of the Gaussian rendering. For a small number of jitters, the histogram visualization is much faster than the Gaussian rendering but the resulting images may appear noisy.

6. Post-processing

The order of post-processing steps is user-specified. However, we recommend the following order: Remove duplicates (if MFA was enabled), filtering (to remove outliers), Z-stage offset (if applicable), drift correction, and finally merging.

The first pass reconstructed image should (unless you are lucky) show some poorly localised molecules (Its recommended to include a bit of unlabelled area in the image if possible to assist with this, e.g. an area outside the cell / tissue

6.1 Results table The results table shows how many photoswitches has been fitted in the first pass.

IMPORTANT. You need to report how many localisations you found in your results

- The actual results of STORM are the localisations (either x,y or x,y,z) which have been subpixel fitted.
- Sigma represents the standard deviation of the fitting distance.
- Intensity is the recordedintensity of the detected ‘blink’ / photoswitch either in photons or in intensity counts (the accuracy of assigning photons depends on what is entered into camera parameters)
- Offset is used determine what the background camera offset would be.
- BKGstd Is the background
- **Uncertainty** = . It is determined according to Equations (31) or (30), where we use $\sigma^2 = \theta_{\sigma 1}, \Theta_{\sigma 2}$. In the axial direction, we use a constant, user-specifed value.

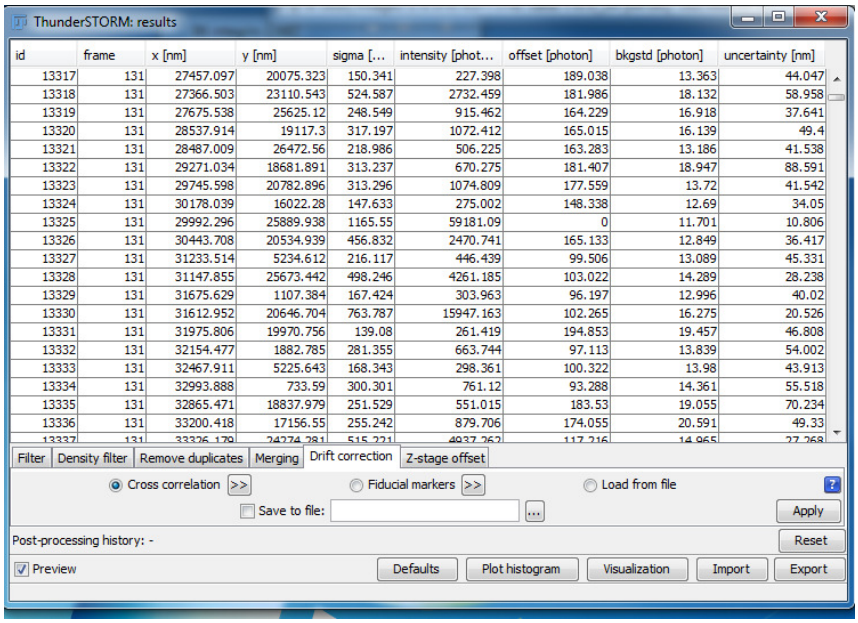

6.2 Remove duplicates

Repeated localizations of single molecules in one frame may occur due to overlapping fitting sub-regions when using multiple-emitter analysis approach. ThunderSTORM allows users to specify a radius below which molecules in close proximity are grouped together. Only the molecule with the smallest localization uncertainty in the group is kept, other molecules are removed. The radius can be specified as a mathematical expression which yields a scalar (same radius for all molecules) or a vector (different radius for every molecule).

Examples of expressions

50, 5*mean(uncertainty), 3*uncertainty

6.3 Filtering:

Outliers can be removed by using a histogram feature, where a region of interest can be selected and applied to the filter rules. Users may need to use a manual range and binning, if outliers are too far from the rest of the data (e.g., one very large value which distorts the automatic histogram function). Check histograms especially for fields: sigma, intensity, offset, and uncertainty. More complex filtering rules for minimizing detection errors can be found in [4].

The columns in the table are treated as unitless (vectors) and filters can be applied on a true or false (Boolean) basis. E.g. Intensity <300 or Uncertainty is >50. Or frame<1000 & Sigma >200. This is useful if the first few frames have too many molecules to sensibly fit.
An example of removing bad fits (see nasty grid artefact) is here:

Example of creating a filter expression using a histogram

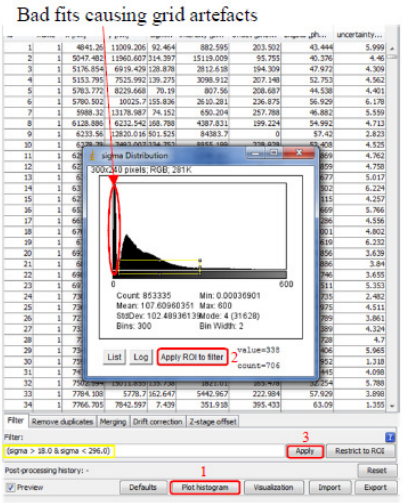

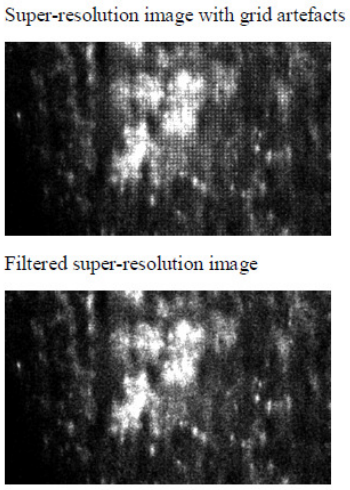

First to work out where the error is coming from plot a histogram (Illustrated in the diagram)
The histogram of sigma (error of fitting) shows a high number of background localisation

To resolve this issue one can draw an ROI around the datapoints and then check Apply ROI to filter, this then removes the obvious outliers.

6.4 Drift correction:

Even with the sample chamber heated to a few degrees above room temperature there is still chance of a small amount of system drift. Even 50nm perturb the final results so this must be corrected for. Either by cross correlation or by use of fiducial markers.

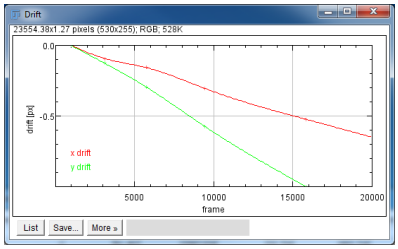

6.4.1 Drift correction by cross-correlation: The parameter "Number of bins" controls the time resolution of the drift trajectory by splitting the image sequence into an appropriate number of bins. Molecular localizations from each bin are used to reconstruct one super-resolution image. "Magnification" controls the lateral resolution of the drift trajectory through the magnification of the reconstructed images. A small number of localized molecules requires a smaller number of bins so that there will be enough data in each sub-sequence. This decreases the time resolution of the drift estimation. A smaller magnification setting can also help to obtain resolvable peaks in cross-correlation images created from images with less data or with unclear structures. Cross-correlation images with detected peaks can be viewed by checking the "Show cross-correlations" checkbox.

6.4.2 Drift correction by tracking fiducial markers: Fiducial markers are automatically detected as molecules that stay in the "on" state at one position for a substantial amount of time. The lateral tolerance for identification of a marker is controlled by the setting "Max distance". The parameter "Min marker visibility ratio" controls the fraction of frames where the molecule must be detected to be considered as a fiducial marker. The ratio should be set higher than the longest "on" state for a regular blinking molecule. Values higher than 0.5 might not work due to possibility of missed detections. "Trajectory smoothing factor" controls smoothness of the drift trajectory and ranges from 0 (no smoothing) to 1 (highest smoothing).

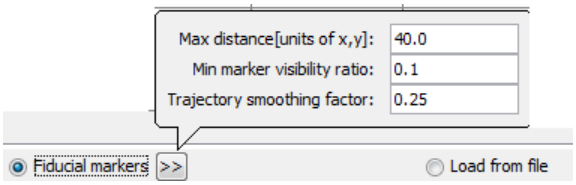

Analyzing samples with fiducial markers yields localizations of both the blinking fluorophores and the fiducial markers. This may slow down the merging algorithm. For faster marker identification, the merging process can be limited to regions containing only the fiducial markers. The drift trajectory can then be saved to a file and applied later to the whole dataset.

6.5 Merging

A single photo-activated molecule may appear in several sequential images, then disappear for several frames, appear again, and finally bleach completely. Such molecules can be combined into one single molecule based on a user-specified distance within which molecules are merged together in the subsequent frames, and the allowed number of frames in which the molecule can disappear. Users can also specify a maximum number of consecutive frames such that a repeating event is still considered a single molecule, see the timing diagram. After merging, a new column called detections appears in the table of results, showing the number of merged molecules. This new column can also be used for filtering.

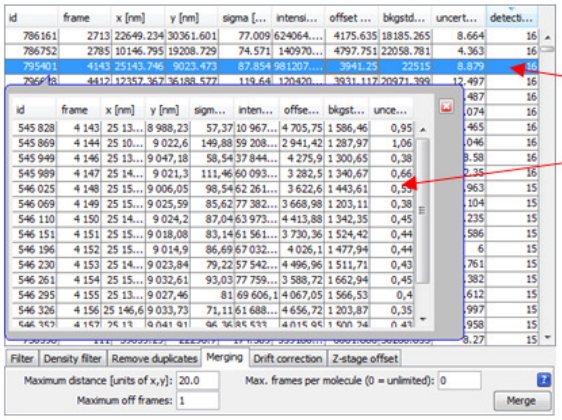

Right-click on a particular row in the table of results and select Show list of merged molecules to see the list of molecules merged in that row.

6.6 Z-stage offset.

This corrects for any offset between where the stage position really is and what is recorded. A Z stack of 100nm beads is required for this.

7 Performance evaluation

7.1 Introduction The performance of the parameters which have been selected by the user: e.g. Filtering, Approximate localisation Threshold, Subpixel localisation method, initial sigma, can be evaluated using artificial datasets.

A 'ground truth' simulated dataset can be generated by Thunderstorm with the precise positions of molecules determined.

The parameters used in the STORM analysis of a dataset can then be compared for efficacy (e.g. how precisely do the parameters localise the images, how good of a job do they do are there false positives and negatives). Several useful metrics then give an indication of how good these parameters are.

A basic protocol would involve:

- Analysing the 'real, biological' data with the suggested parameters.
- Generating simulated data based on the output of this
- Analysing the simulated data with the parameters used with the 'real biological data' and seeing how good they are.
- Adjusting the parameters, testing on the simulated data
- Reanalysing the 'real biological' data with the recommended parameters

7.2 Generation of simulated data

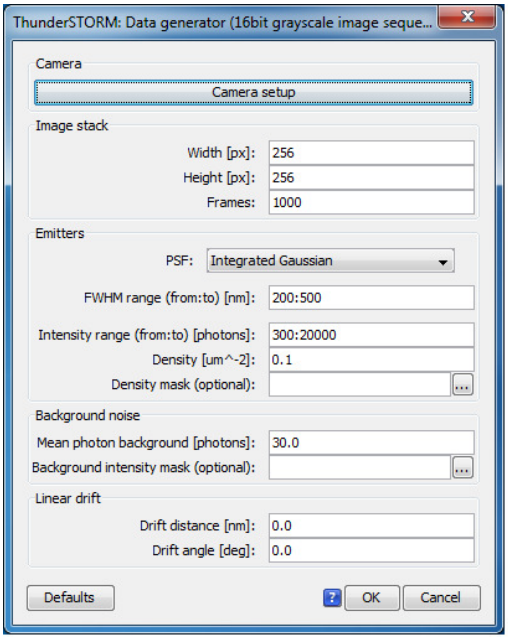

ThunderSTORM is capable of generating a sequence of SMLM-like images in which the ground-truth positions of the molecules are known. This allows users to perform Monte Carlo simulations *[4, 3]* and to quantitatively evaluate the performance of applied localization algorithms by calculating, e.g., the Jaccard index or F1 score using the Performance evaluation plugin. In addition to the image size and sequence length, users can specify the intensity, imaged size, and spatial density of the generated molecules. The resulting images can be subjected to sample drift. Noise in the generated images simulates the behavior of CCD or EMCCD cameras.

7.2.1 How does Thunderstorm generate Ground Truth Data

For each frame, an ideal, noise free, SMLM-like image is created simulating the expected number of photons detected in each camera pixel. Image formation starts by creating a list of molecules with FWHM and intensity chosen randomly in user-specified ranges, and with random positions of molecules given by a user-specified spatial density (see below). Users can also specify any of the implemented PSF models, including 3D models, to create the simulated images of molecules. The modelling includes Poisson distribution of intensities to better simulate biological data. The generated molecules are added sequentially to the final image similarly as in the Gaussian rendering

method. A user-specified offset is added to the generated image sequence to simulate photon background.

Alternatively, a gray-scale image, in which each pixel value is normalized to the interval [0,1], can be used as a weighting factor of the offset level in different parts of the generated images to simulate an irregular background as might be encountered in real samples. This is added using the Density Mask [optional] diaglogue box and requires an image.

Finally, the signal in the camera register is digitized by converting the photons to digital counts. The CCD sensitivity (in photons per A/D count), and the camera digitizer offset (in A/D counts) are user-specified in the camera setup, as well as the camera pixel size (in nanometers) as projected to the sample plane.

7.2.2 Fixed or spatially varying density of molecules: Users can specify a fixed or spatially varying density of simulated molecules in the generated images. A Value of greater than 2 is recommended otherwise the image looks like white noise.

Note that the frame size should be at least the same size as the desired super-resolution image in pixels order to preserve high resolution in the final reconstruction. The coordinates of the molecular centers are down-scaled appropriately.

7.2.3 Additional sample drift: The generated molecular positions in the image sequence can be subjected to a lateral sample drift. Users need to specify the speed and direction of the drift, which is constant throughout the image sequence. This can be modelled based on the experiment itself from the drift data see 6.4.1 (Drift correction)

7.3 Evaluating data. The analysis of the ground truth molecule positions is checked against a set of parameters used to analyse a 'real biological' experiment. (Plugins-Thunderstorm-Run analysis). This is quick as the Ground Truth data will only have 1000 frames.

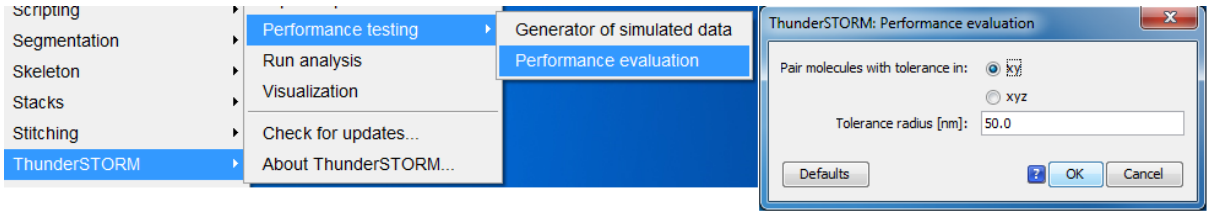

The performance evaluation plugin pairs the localized molecules from user specified analysis parameters with the closest molecule in the ground-truth data.

The user specified the tolerance radius which they expect these molecules to fit in in the Performance evaluation dialogue box e.g. 50nm as shown above (this could be considered the actual resolution. The software then identifies True Positive localisation (TP). This is where a localisation from the analysis parameters falls within the given tolerance radius of the ground truth data. A False Positive (FP) is when the distance from the analysis parameters is greater than or equal to the tolerance radius from the ground truth data. Ground-truth molecules which were not associated with the localized molecules are counted as FNs.

With a growing density of molecules it becomes more important how the algorithm performs the matching. To solve the problem of finding the correct matching between localized molecules and the ground-truth data, the Gale-Shapley algorithm *[1]* is used. KD-trees *[2]* are employed for an effective implementation.

7.4 Output of the performance evaluation.

Parameters	Distance radius [nm]	# of TP	# of FP	# of FN	Jaccard index	Precision	Recall	F1-measure	RMSE x [nm]	RMSE y [nm]	RMSE lateral[nm]	RMSE total [nm]	Approx Localisation threshold		Sub Pixel Localis'n	Initial Sigma 4 fitting
1	50	133072	13288	34435	0.73604	0.90921	0.794427	0.847952	3.7906	3.783	5.93753137	5.9375314	1*Std(Wave.F1)			1.6
2	50	126794	12856	40713	0.70299	0.907941	0.756947	0.825597	3.6443	3.6299	5.70016491	5.7001649	2*Std(Wave.F1)			1.6
3	50	0	18304	167507	0	0	0	NaN	NaN	NaN	NaN	NaN	0.5*Std(Wave.F1)			1.6
4	50	132732	9545	34775	0.74968	0.932913	0.792397	0.856933	3.7809	3.7743	5.92286152	5.9228615	Non-Max Suppression	Dilation 3	Max Likli	1.6
6	50	132770	6781	34737	0.76179	0.951408	0.792624	0.864788	4.8957	4.8952	7.66406844	7.6640684	Non-Max Suppression	Dilation 4	Least Sq	1.8

The output metrics are shown (light grey columns) The parameters used to 'analyse' the ground truth data are in the dark grey columns

- **Distance Radius:** Tolerance of fitting for the comparison
- **#TP / #FP:** Number of True and False positives

Precision and Recall: Statistical measures related to the number of correctly or incorrectly detected molecules, or missed molecules, are the recall (also called sensitivity) and the precision (also called positive predictive value) *[4, 5, 3]*. Their definitions are given by

$$r = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FN}},$$
$$p = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FP}}.$$

TP = True positives, FN = False Negatives. FP = False positives

Recall measures the fraction of correctly identified molecules, and precision measures the portion of correctly identified molecules in the set of all localizations. The theoretical optimum is achieved for values of recall and precision both equal to 1.0.

F1 score: For purposes of comparison between multiple algorithms, it is convenient to combine precision and recall into a single measure of performance with some trade-off between both values. A traditional method for this applies the F1 score *[4, 3]* defined by

$$F_1 = \frac{2pr}{p+r}.$$

Values of the **F1** score close to zero indicate both bad recall and precision while values approaching 1.0 signify a good ratio between recall and precision.

Jaccard index: Another measure suitable for comparing similarity and diversity of sets of samples is the Jaccard index *[4]* defined by the formula

$$J = \frac{|\mathcal{A} \cap \mathcal{B}|}{|\mathcal{A} \cup \mathcal{B}|}.$$

Here $\mathcal{A}$ is the set of ground-truth molecular positions, $\mathcal{B}$ is the set of all molecular positions localized by processing the data, intersection $|\mathcal{A} \cap \mathcal{B}| = \mathrm{TP}$ gives the number of true positive detections, union $|\mathcal{A} \cup \mathcal{B}| = \mathrm{TP} + \mathrm{FP} + \mathrm{FN}$, and $|\cdot|$ denotes the size of the set. The Jaccard index ranges from zero to one and a theoretical optimum is achieved for values of the Jaccard index equal to 1.0.

RMS distance: For all molecules identified as true positives, we also calculate the root-mean square distance between the ground-truth positions of the molecules and their localizations. This again is a measure of the localisation accuracy and can be included in figure legends etc.

8. References

Filtering references

- **(B-Spline method)** I. Izeddin, J. Boulanger, V. Racine, C. G. Specht, A. Kechkar, D. Nair, A. Triller, D. Choquet, M. Dahan and J. B. Sibarita(2012) Wavelet analysis for single molecule localization microscopy, Optics Express 20 (3), pp.Â 2081â€“95. External Links: Document. Cited by: Wavelet filter, Wavelet filter, Wavelet filter, Wavelet filter.
- J.-L. Starck and F. Murtagh(2002) Astronomical Image and Data Analysis, Springer-Verlag. Cited by: Wavelet filter, Wavelet filter, Wavelet filter.
- **(Diff averages filter)** F. Huang, S. L. Schwartz, J. M. Byars and K. A. Lidke(2011) Simultaneous multiple-emitter fitting for single molecule super-resolution imaging, Biomedical Optics Express 2 (5), pp.Â 1377â€“93. External Links: Document. Cited by: Difference of averaging filters.
- [1] S. J. Holden, S. Uphoff and A. N. Kapanidis(2011) DAOSTORM: an algorithm for high-density super-resolution microscopy, Nature Methods 8 (4), pp.Â 279â€“80. External Links: Document. Cited by: Lowered Gaussian filter.
- [2] P. B. Stetson(1987) DAOPHOT - A computer program for crowded-field stellar photometry, Publications of the Astronomical Society of the Pacific 99, pp.Â 191. External Links: Document. Cited by: Lowered Gaussian filter.
- **Centroid of connected components**
- The algorithm which identifies connected components. D. E. Knuth(1997) The Art Of Computer Programming, 3rd edition, Vol. 1, Addison-Wesley, Boston. Cited by: Centroid of connected components.
- Watershed Algorithm: M. Å onka, V. HlavÃ¡Ä and R. Boyle(2007) Image Processing, Analysis, and Machine Vision, 3rd edition edition, Cengage Learning. This is based on the ImageJ Find Maxima function.

References sub pixel localisation

Sub pixel localisation References

- [1] F. Huang, S. L. Schwartz, J. M. Byars and K. A. Lidke(2011) Simultaneous multiple-emitter fitting for single molecule super-resolution imaging, Biomedical Optics Express 2 (5), pp.Â 1377â€“93. External Links: Document. Cited by: PSF model: Integrated form of a symmetric 2D Gaussian function.
- [2] C. S. Smith, N. Joseph, B. Rieger and K. A. Lidke(2010) Fast, single-molecule localization that achieves theoretically minimum uncertainty, Nature Methods 7 (5), pp.Â 373â€“5. External Links: Document, ISSN 1548-7105. Cited by: PSF model: Integrated form of a symmetric 2D Gaussian function.
- [1] R. Henriques, M. Lelek, E. F. Fornasiero, F. Valtorta, C. Zimmer and M. M. Mhlanga(2010) QuickPALM: 3D real-time photoactivation nanoscopy image processing in ImageJ, Nature Methods 7 (5), pp.Â 339â€“340. External Links: Document. Cited by: Centroid of local neighborhood.
- B. Huang, W. Wang, M. Bates and X. Zhuang(2008) Three-dimensional super-resolution imaging by stochastic optical reconstruction microscopy, Science 319 (5864), pp.Â 810â€“3.
- R. Parthasarathy (2012) Rapid, accurate particle tracking by calculation of radial symmetry centers, Nature Methods 9 (7), pp.Â 724â€“6. External Links: Document.

- [1] P. R. Bevington and D. K. Robinson(2003) Data reduction and error analysis for the physical science, McGraw-Hill Higher Education, McGraw-Hill. External Links: ISBN 9780072472271. Cited by: Least-squares methods.
- [2] Commons-Math(2013-04) The Apache Commons Mathematics Library; version 3.2, External Links: Link. Cited by: Least-squares methods.
- [3] F. Huang, S. L. Schwartz, J. M. Byars and K. A. Lidke(2011) Simultaneous multiple-emitter fitting for single molecule super-resolution imaging, Biomedical Optics Express 2 (5), pp.Â 1377â€“93. External Links: Document. Cited by: Maximum-likelihood estimation.
- [4] M. Kendall and A. Stuart(1979) The Advanced Theory of Statistics, London: Charles Griffin. Cited by: Least-squares methods, Maximum-likelihood estimation.
- [5] K. I. Mortensen, L. S. Churchman, J. A. Spudich and H. Flyvbjerg(2010) Optimized localization analysis for single-molecule tracking and super-resolution microscopy, Nature Methods 7 (5), pp.Â 377â€“381. External Links: Document. Cited by: Least-squares methods, Maximum-likelihood estimation.
- [6] R. Oâ€™Neill(1971) Algorithm AS 47â€“function minimization using a simplex procedure, Applied Statistics 20, pp.Â 338â€“45. External Links: Link. Cited by: Maximum-likelihood estimation.
- [7] C. S. Smith, N. Joseph, B. Rieger and K. A. Lidke(2010) Fast, single-molecule localization that achieves theoretically minimum uncertainty, Nature Methods 7 (5), pp.Â 373â€“5. External Links: Document, ISSN 1548-7105. Cited by: Maximum-likelihood estimation.

Visualisation References

- [1] D. Baddeley, M. B. Cannell and C. Soeller(2010) Visualization of localization microscopy data, Microscopy and Microanalysis 16 (1), pp.Â 64â€“72

[2] P. Kek, I. Raļka and G. M. Hagen(2011) Minimizing detection errors in single molecule localization microscopy, Optics Express 19 (4), pp.Â 3226â€“35.

Postprocessing Refs

1. L. J. Van Vliet, F. R. Boddeke, D. Sudar, I. T. Young. Image Detectors for Digital Image Microscopy, in Digital Image Analysis of Microbes, Modern Microbiological Methods, M. H. F. Wilkinson and F. Schut, Eds. Chichester, United Kingdom: John Wiley & Sons, pp. 37-64 (1998).
2. J. Janesick, K. Klaasen, T. Elliott. CCD charge collection efficiency and the photon transfer technique, Proc. SPIE, Solid State Imaging Arrays, vol. 570, pp. 7-19 (1985).
3. I. Izeddin, J. Boulanger, V. Racine, C. G. Specht, A. Kechkar, D. Nair, A. Triller, D. Choquet, M. Dahan, J. B. Sibarita. Wavelet analysis for single molecule localization microscopy. Optics Express 20(3), pp. 2081-95 (2012).
4. P. Křížek, I. Raška, G. M. Hagen. Minimizing detection errors in single molecule localization microscopy. Optics Express 19(4), pp. 3226-35 (2011).

Evaluation of results References

- [1] D. Gale and L. S. Shapley(1962) College admissions and the stability of marriage, The American Mathematical Monthly 69 (1), pp.Â 9â€“15. External Links: Document. Cited by: Counting localized and missed molecules.

- [2] D. E. Knuth(1997) The Art Of Computer Programming, 3rd edition, Vol. 1, Addison-Wesley, Boston. Cited by: Counting localized and missed molecules.
- [3] P. KÅ™ÃÅ¾ek, I. RaÅ¡ka and G. M. Hagen(2011) Minimizing detection errors in single molecule localization microscopy, Optics Express 19 (4), pp.Â 3226â€“35. External Links: Document. Cited by: Precision and recall, F1 score.
- [4] P.-N. Tan, M. Steinbach and V. Kumar(2005) Introduction to Data Mining, Addison-Wesley Longman Publishing Co.. External Links: Link. Cited by: Precision and recall, F1 score, Jaccard index.

- [5] S. Wolter, M. SchÃ¼ttpelz, M. Tscherepanow, S. van de Linde, M. Heilemann and M. Sauer(2010) Real-time computation of subdiffraction-resolution fluorescence images, Journal of Microscopy 237 (1), pp.Â 12â€“22. External Links: Document. Cited by: Precision and recall
- [1] M. Hirsch, R. J. Wareham, M. L. Martin-Fernandez, M. P. Hobson and D. J. Rolfe(2013) A stochastic model for electron multiplication charge-coupled devicesâ€“from theory to practice., PLoS ONE 8 (1), pp.Â e53671. External Links: Document. Cited by: Image formation.
- [2] M. Kendall and A. Stuart(1979) The Advanced Theory of Statistics, London: Charles Griffin. Cited by: Fixed or spatially varying density of molecules.
- [3] P. KÅ™ÃÅ¾ek, I. RaÅ¡ka and G. M. Hagen(2011) Minimizing detection errors in single molecule localization microscopy, Optics Express 19 (4), pp.Â 3226â€“35. External Links: Document. Cited by: Generator of simulated data.