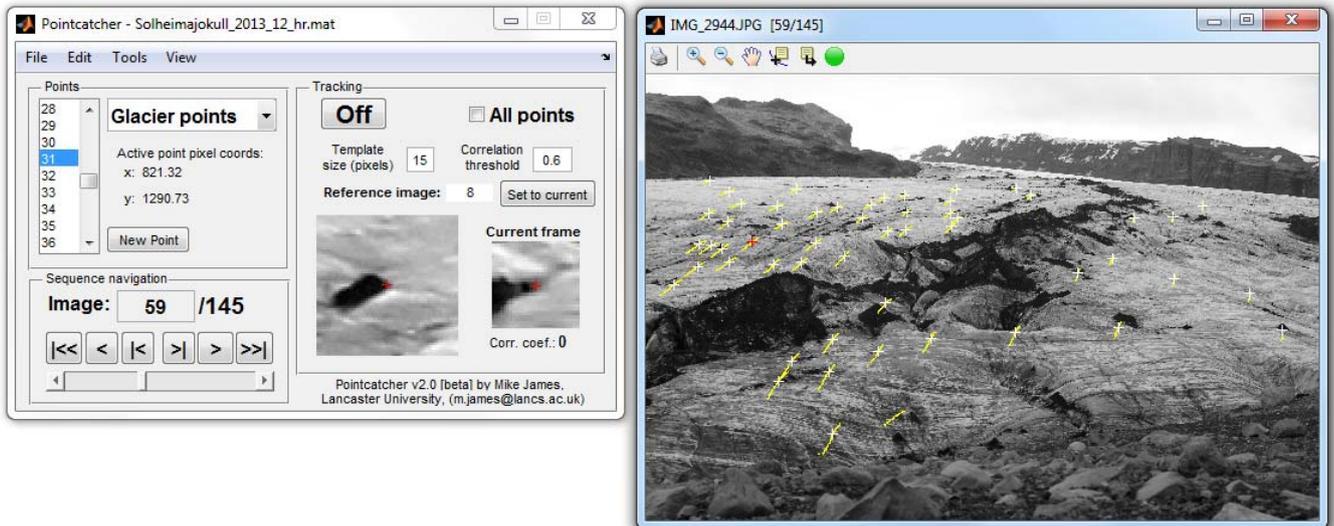


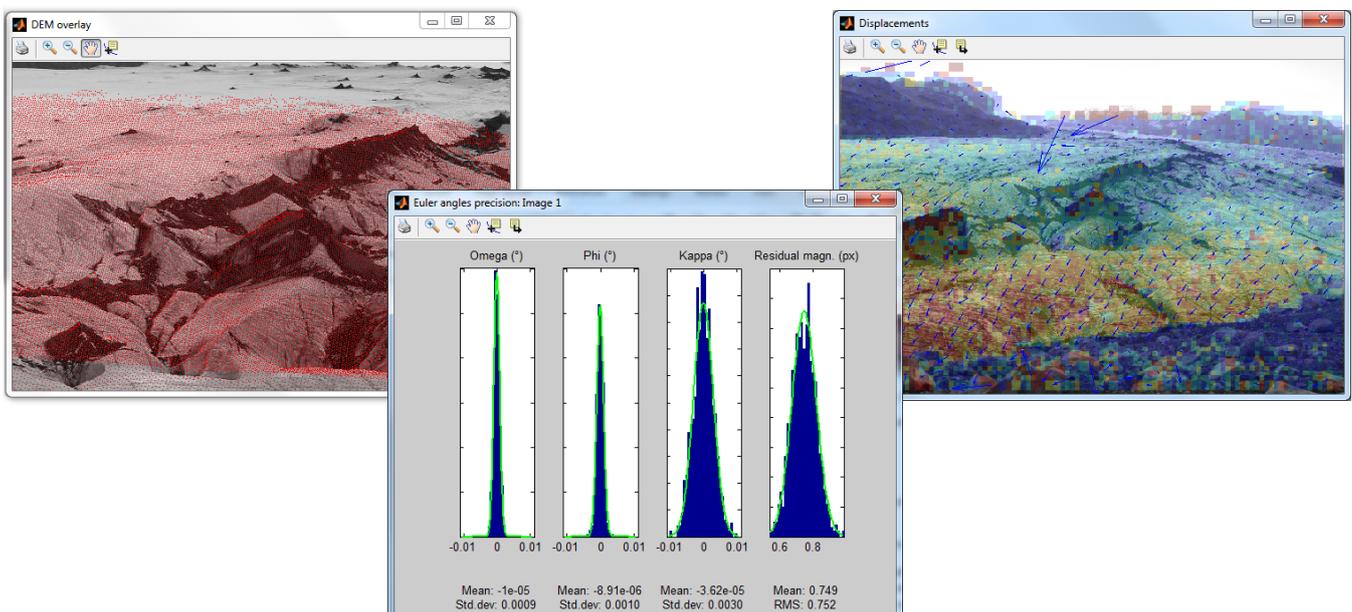
# Pointcatcher v.2.0

<http://www.lancaster.ac.uk/staff/jamesm/software/pointcatcher.htm>

Mike James, Lancaster University  
m.james@lancaster.ac.uk



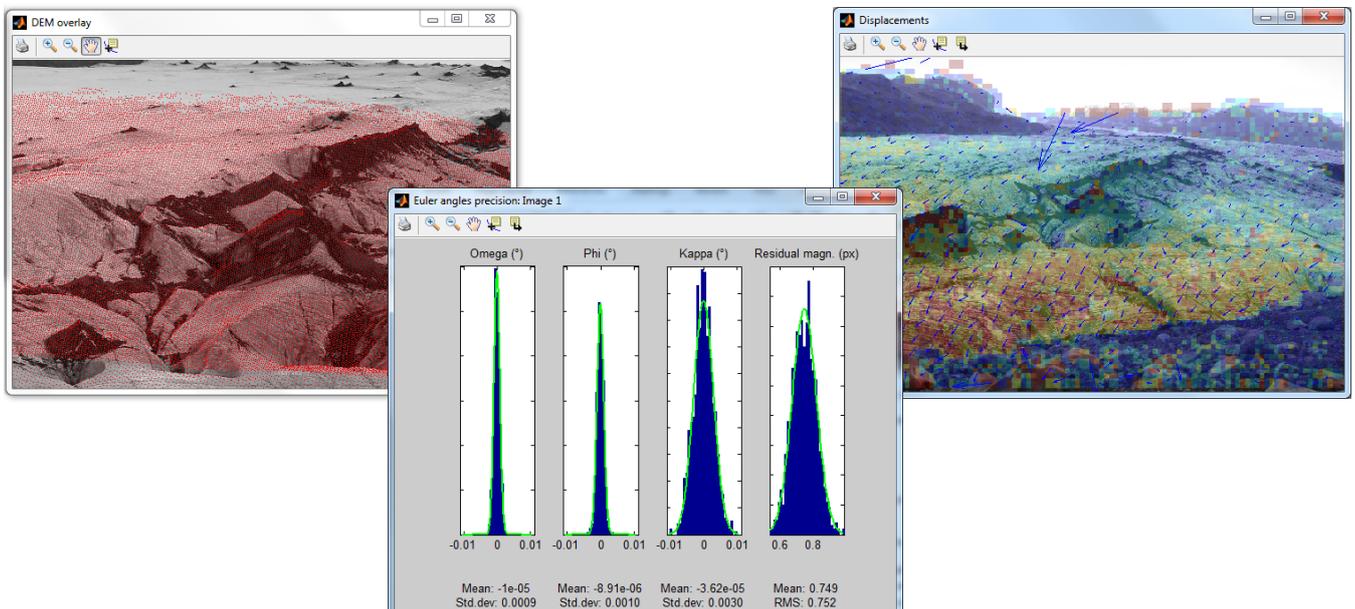
- Feature tracking through image and video time-lapse sequences:
  - Manual
  - Automated (normalised cross correlation)
- Camera and lens distortion models
- Image co-registration, with Monte Carlo-based precision estimates
- Geo-referencing and 3D point coordinates from re-projection to DEM surfaces
- Re-projection of Monte Carlo uncertainty onto DEMs



# Index

The following sections are roughly in order of increasingly complex use of Pointcatcher, so you can start at the beginning and work through.

1. Starting a Pointcatcher project .....	3
2. Image sequence navigation .....	3
3. Pointsets, Points and Observations .....	4
4. Visualising and exporting results .....	5
5. Automated point tracking.....	7
6. Point seeding.....	8
7. Image registration.....	9
8. Assessing uncertainties in image registration.....	11
9. Sequence geo-referencing .....	12
10. Re-project points to DEM surface .....	13
11. Estimate re-projection precision.....	13
12. Estimate displacement uncertainties.....	14
13. Calculate intermediate points.....	15
14. Bugs and requests .....	16



## 1. Starting a Pointcatcher project

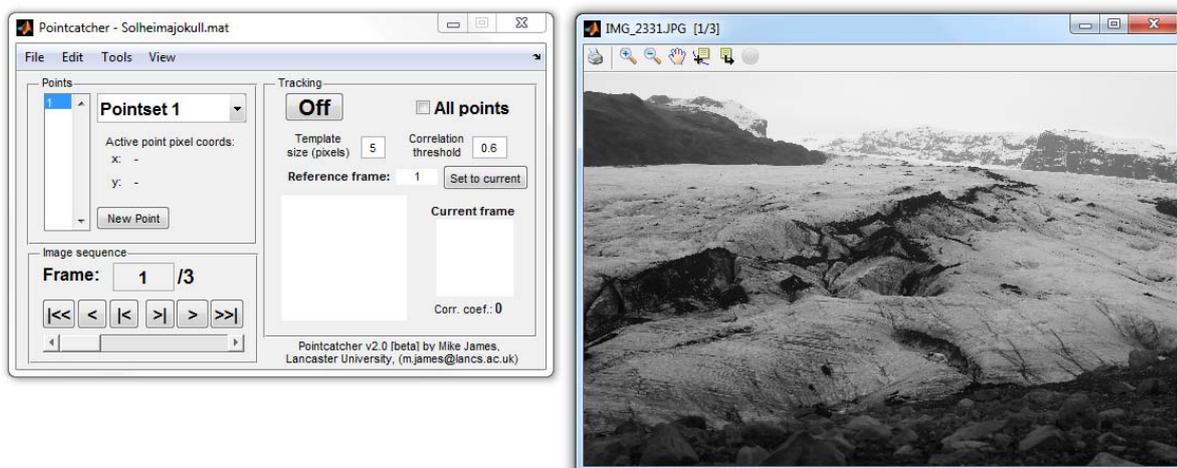
Pointcatcher projects can be run in two different ways – ‘normal’ or ‘file-by-file’. ‘Normal’ projects load all the images simultaneously, so that moving through the sequence is quick and efficient. However, if you have too many images to hold in memory at once, you can use a memory-efficient ‘file-by-file’ mode in which only the active image is read into memory. Thus, very large projects can be used, but navigating through them is slower because each image has to be read from disc on advance. To start with, try a small number of images (e.g. 5) in ‘normal’ mode.

If you want to use the memory-efficient ‘file-by-file’ mode, ensure that the menu item File → File-by-file mode is checked before you start the project.

**Start a new project** by File → New Pointcatcher project and using shift-click to select all the image files that you want to analyse. Images will then be read in (this may take a while for many large images); when complete, the image window will appear.

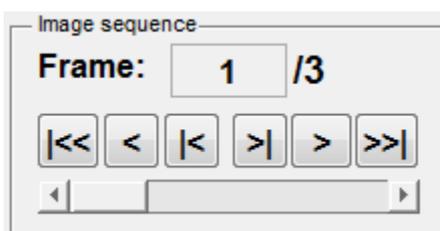
**Save the new Pointcatcher project** by File → Save Pointcatcher project as... (this does not save copies of your images, just the Pointcatcher data, which includes links to the original images).

**NOTE:** Pointcatcher project files save information about the image sequence, camera orientation, geo-referencing and point observations. 3D point data such as reprojected points, or imported DEMs are *not* saved. Any reprojected point coordinates should be exported for further analyses.



With a project started, the main Pointcatcher window (left) and an image window (right) will be displayed.

## 2. Image sequence navigation

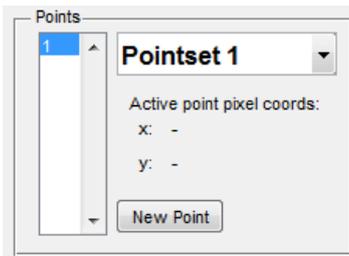


Navigate through the image sequence using the self-explanatory slider and buttons in the lower left panel of the Pointcatcher main window. The buttons allow play, step and advance to end in both forward and backward directions. If playing forwards, click the play button again to stop. You can go directly to any specific image by typing the number in the Image box.

You can also step forward/backward by using < or > key presses.

You can zoom in/out of the image window or pan by using the  tools, or through the mouse wheel (zoom) or right-click-drag (pan). To deselect a tool once active, click the tool's button again.

### 3. Pointsets, Points and Observations



A **'Point'** is an image feature that you want to track through the sequence; they are numbered sequentially from 1 and listed in the listbox in the Points panel of the main window, with the active point highlighted.

A **'Point observation'** is the location of that Point in any one image (not all Points may be observable in every image). Coordinates of the observation of the active point in the current image are given in the Points panel.

A **'Pointset'** is a group of Points (and their associated observations) that are collected together for convenience, or to enable a different type of analysis. The active Pointset is shown in the dropdown box in the Points panel. Typical use of Pointsets is to have one Pointset for static features in an image sequence, which can be used for registration, and another Pointset for features of interest (e.g. surface features on a moving glacier).

By default, new projects are initialised with one Point and one Pointset, but no Point observations.

#### **Points:**

The 'active' Point will be highlighted in the list box, and any observation of it will be shown as a red + on the image.

##### *Adding a new Point*

Click the 'New Point' button in the Points panel of the main window.

##### *Selecting a different Point*

A different Point can be made active by selecting its ID in the list box, or clicking on its observation in an image.

#### **Point observations:**

##### *Making a Point observation*

Make sure the Point you want is highlighted in the list box, and the cursor is crosshairs when over the image – then click on the image in the position you wish. (If the cursor is not crosshairs, then deselect the active navigation tool by clicking on the tool's button). A red + will appear, and coordinates will be shown in the Points panel. [Note: If you click too close to another Point observation, then this will just select the other point. Reselect the Point you wish to be active in the Point list, then zoom in on the image until the new position you need is separated from the existing observation of the other point).

##### *Changing a Point observation*

Ensure the correct Point is active, then click in the new position on the image.

##### *Deleting a Point observation*

Ensure the correct Point is active, and the image window is active, press the 'delete' key. For deleting multiple observations at the same time, the Edit menu provides a variety of options:

Edit → Point observations → Delete...

#### **Pointsets:**

##### *Adding a new Pointset*

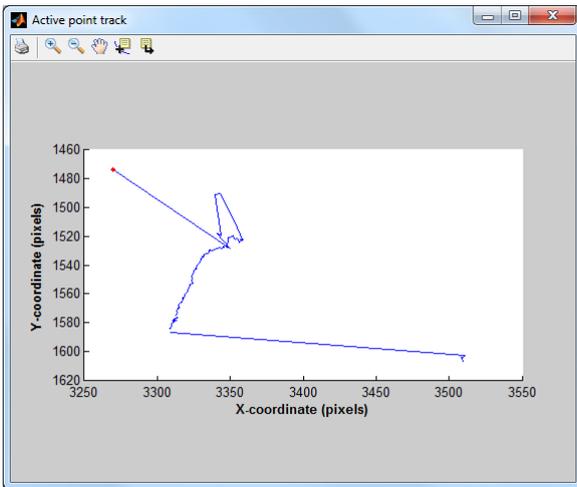
Go to the menu item Edit → Pointsets → New Pointset

##### *Renaming Pointsets*

Go to the menu item Edit → Pointsets → Edit pointset name

## 4. Visualising and exporting results

Pointcatcher provides a number of different ways to view the Point observation data collected in a sequence, which are accessed through the **View menu**. For most of these, data can be exported via the  toolbar button.

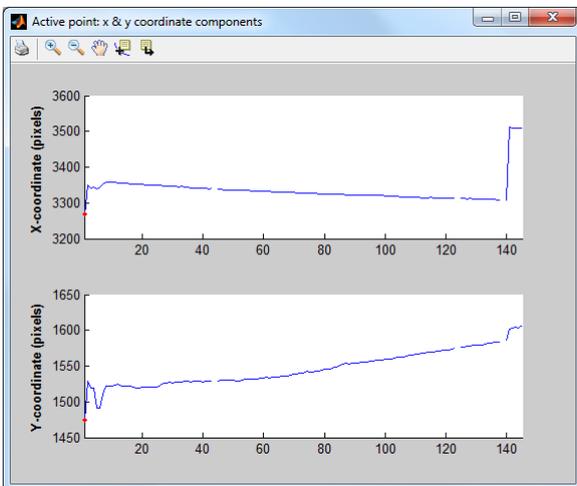


### Active point: x-y plot

Shows the pixel coordinates of the active point throughout the sequence.

The red symbol shows the point location in the currently active image.

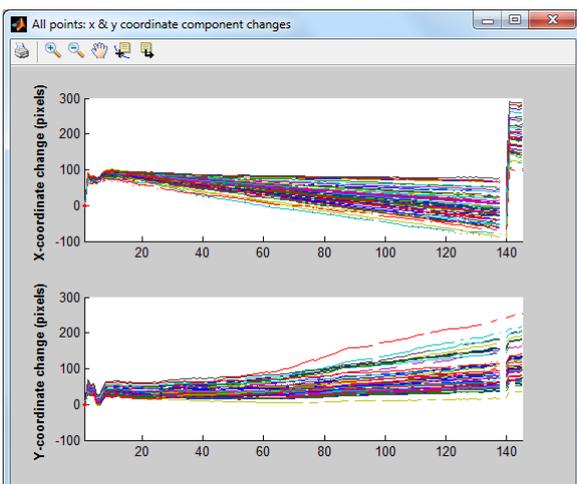
Mouse clicks in the graph change the current image to the one nearest the one represented by the clicked point coordinates.



### Active point: x & y coordinate components plots

As above, this shows the pixel coordinates of the active point throughout the sequence, but now split by x and y coordinate components. Unlabelled abscissa axis is image number.

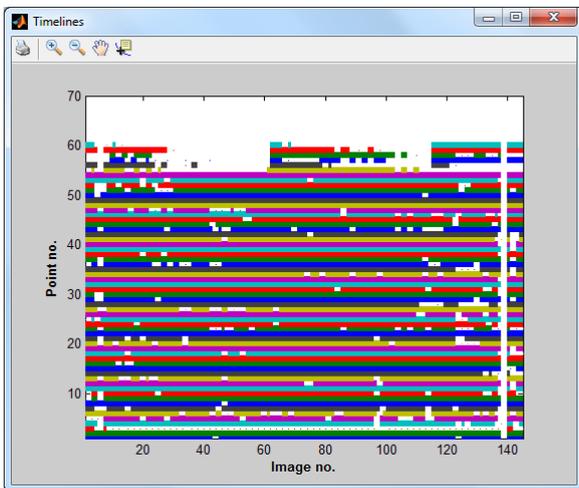
Mouse clicks in the graph change the current image to the one nearest the one represented by the clicked point coordinate.



### All points: x & y coordinate components plots

Shows the change in pixel coordinates of all points throughout the sequence, split by x and y coordinate components. Unlabelled abscissa axis is image number.

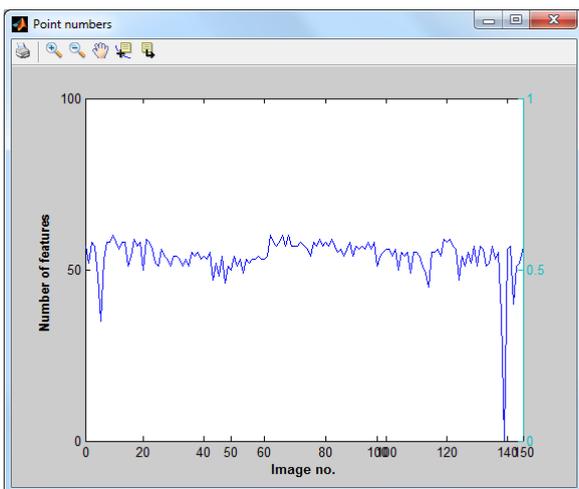
Mouse clicks in the graph change the current image to the one nearest the one represented by the clicked point coordinate.



### All points: Timelines

Shows the continuity of the tracks for all points. The bars indicate the image numbers for which there are observations of the individual points.

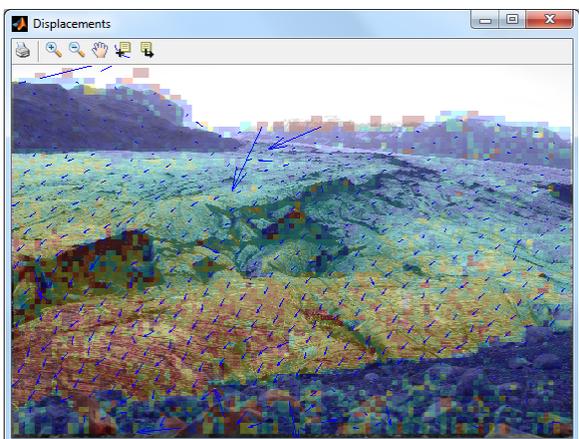
Mouse clicks in the graph change the current image to the one nearest the one represented by the clicked point coordinate.



### Point number statistics

Indicates the number of point observations per image. If the active Pointset is the one that is used for image registration, then the number of inlier/outlier points are indicated.

Mouse clicks in the graph change the current image to the one nearest the one represented by the clicked point coordinate.



### Displacement vectors (image pair)

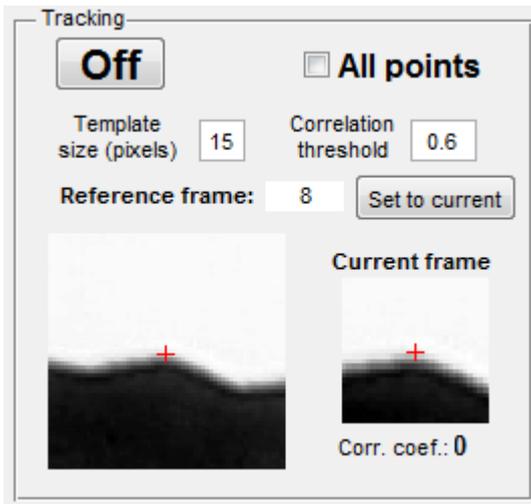
This is specifically meant for observing dense arrays of points in image pairs – e.g. such as provided by seeding using a regular grid.

The first image of the pair is used as the background image, overlain with derived displacements indicated by vectors and a semi-transparent coloured overlay.

The associated control window provides enables changing vector and raster overlay parameters.

Note: With large colour images, updating is slow.

## 5. Automated point tracking



The procedures above describe how to click on an image to manually allocate a Point observation, and this effectively represents a fully interactive point tracking procedure. Pointcatcher implements normalised cross-correlation-based tracking to facilitate automated measurements. Each point is associated with an image template which is shown on the left hand side in the Tracking panel of the main window. For automated tracking, this template is used to try and find a match in the current image. Successful matches are determined by whether the maximum correlation coefficient exceeds a threshold. The size of the image patch (the template size), the image from which it is taken, and the correlation threshold can be changed at any point.

*To track a new point:*

- 1) Make an observation of that point in an image, then click the main Tracking button to **ON**.
- 2) Use the sequence navigation buttons to step or play through the sequence.
- 3) Where correlation coefficients exceed the threshold, the point observation is stored.

Tracking works by calculating a normalised cross-correlation over a search area in the current image to look for a match. The centre position of the search area is determined by a set of rules.

- 1) If an observation of this point already exists in this image, then this is the centre for the search.
- 2) If there is no observation, the central position is forecast depending on the strategy selected in Tools → Point tracking → Position forecasting
  - Constant position - the point is assumed not to have moved since the last observation, this is the default option.
  - Constant velocity – if a point velocity can be calculated (e.g. two or more previous observations), then this will be used to define an approximate offset from the last observation
  - Locked to reference – the point is assumed to always be close to where it is in the reference image.
- 3) Finally, if there is image registration information available to describe image movement along the sequence (see next section), then this is used by default within the position forecast. This can be de-selected, in which case camera motion will not be accounted for.

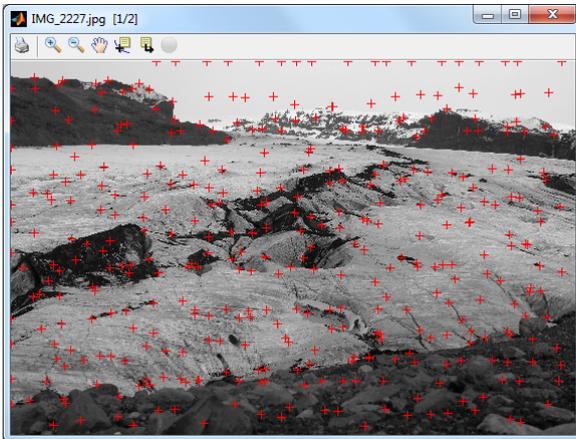
Tracking also needs to know how large the search area should be, and this is set as a multiple of the template size, via Tools → Point tracking → Set search radius; the default is 2.

*All points*

By default, only the active point will be tracked during movement through the sequence. To track all points simultaneously, check the 'All points' box.

## 6. Point seeding

For applications in which tracked points are required across an image, initial Point observations can be automatically determined using Point seeding. Seeding works by using an 'interest operator' (Harris) across the image to identify good features for tracking, or by defining a regular grid.

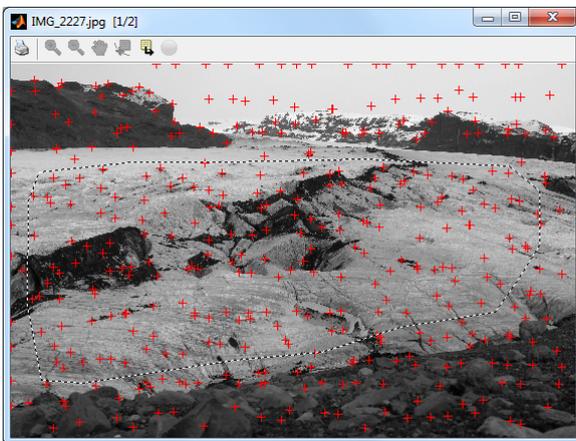


To use **Harris features** (as shown left):

Tools → Point seeding → Harris then give an approximate number of seeds to locate across the image, (default is 400).

For a **regular grid**:

Tools → Point seeding → Regular grid then give the grid interval.



With the seed points detected, you can remove areas that you don't want by using

Tools → Point seeding → Select features

Click in the image to select the region of seeds that you wish to retain. Right-click to finish the selection; seeds *outside* the selected region will be removed.



Finally, the remaining seed points need to be converted into new Points and Point observations via

Tools → Point seeding → Convert features to trackable points

for ongoing use in Pointcatcher.

They will then be added as new points and point observations in the active image and Pointset.

## 7. Image registration

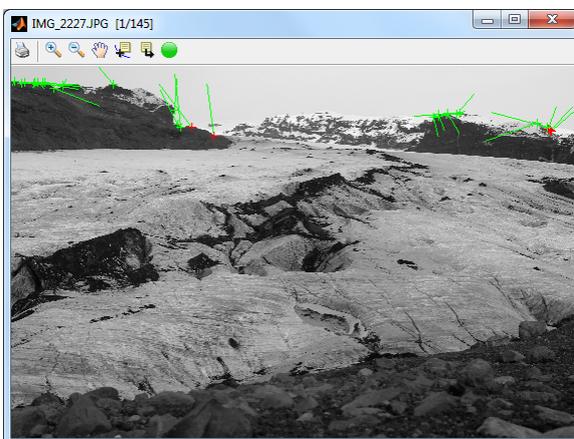
Image registration enables movement (rotations) of the camera during the sequence to be accounted for. To register images, you need observations of 3 (preferably many more, e.g. 20) points in different images. If you want to register images, then dedicate one Pointset to collating the points you will use.

- 1) Create a new Pointset to contain the registration points: Edit → Pointsets → New Pointset
- 2) Rename the Pointset if you wish, so that it is clear what it is for: Edit → Pointsets → Edit pointset name
- 3) Add Points to represent static features in the scene, and collect observations of these features throughout the sequence. See Section 3 for adding Points, and Section 4 for automated tracking through the sequence. If camera motion is minimal, then automated tracking can be successful throughout the sequence. If there are sudden offsets, then automated tracking can fail. In this case, make a few point observations manually, calculate an image transform (see below), then retry the automated tracking. The transform should guide the tracker by enabling better estimates for the search positions.
- 4) To calculate image registration transforms, select an image to act as the reference image: Tools → Image registration → Set Parameters... → Reference image, and enter the image number. The reference image is the image that other images will be registered to and it should be the one with the most point observations in it. Default reference image is the first image.
- 5) Perform an initial registration of the images: Tools → Image registration → Register all images

As the registration proceeds, reports on how many features were accepted as inliers in each image, and the RMS of the residual error magnitudes are given. If many points are rejected, you may want to increase the threshold for outlier detection (Tools → Image registration → Set Parameters... → Max distance for inliers (default is 1.5 pixels). The status of observations in an image can be visualised: View → Points... → Colour by out/inlier

White = not used in the calculation (i.e. no observation of this point in the reference image)  
Green = inlier  
Red = outlier

Residual vectors can also be shown: View → Registration residuals, and their magnification (default = 100) varied: View → Set residuals multiplier. If the 'residuals multiplier' = 1, then a residual of magnitude 1 pixel is indicated as a green line of length 1 pixel.

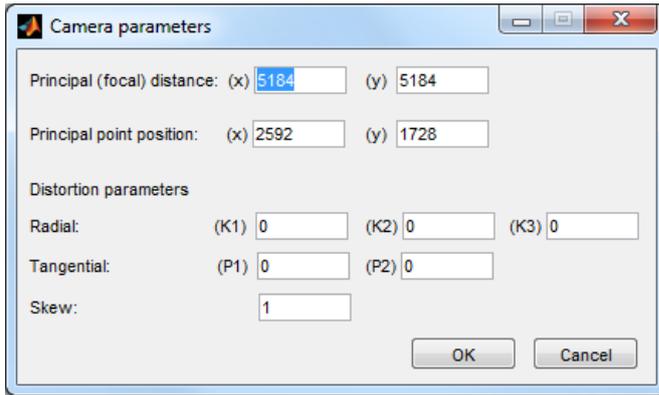


Example of a Pointset used for image registration, plotted as coloured by out/inlier, and with overlain residual vectors (magnified by 500).

Note that initially registered images will show a circular yellow symbol in their toolbar. A symbol with an 'R' denotes the reference image.

Registration information can be reset or updated by re-calculating the registration following edits to their points or their observations.

- 6) To refine the image registration by representing it in terms of rotations of a physical camera, a camera model is required.



Use Edit → Camera parameters to define the parameters via the dialog box. The default values are estimated, with the parameters defined as described here:

[http://www.vision.caltech.edu/bouguetj/calib\\_doc/htmls/parameters.html](http://www.vision.caltech.edu/bouguetj/calib_doc/htmls/parameters.html)

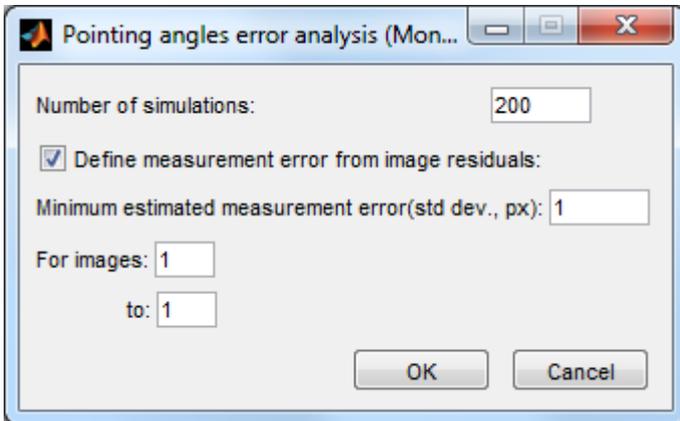
- 7) Registration refinement can now be carried out by checking the Tools → Image registration → Camera model option, then re-running Tools → Image registration → Register all images. Images for which this is successful will show a green circular symbol in the toolbar. Note that this refinement uses only the point observations that were deemed inliers in the initial registration. Inlier/outlier designations are not updated.

Note that images with refined registrations will show a circular green symbol in their toolbar.

## 8. Assessing uncertainties in image registration

For an image that has been registered using a camera model, the uncertainty in the recovered camera angles can be estimated using a Monte Carlo approach. Random offsets are repeatedly added to the point observations used in the image registration, and the camera angles are estimated for each combination. Overall precision of the recovered camera angles is then given by the distribution of angles.

Tools → Error analyses → Estimate precision of camera angles (Monte Carlo)...



The dialog box has options for:

Number of simulations (this is the number of randomised runs per image, default is 200)

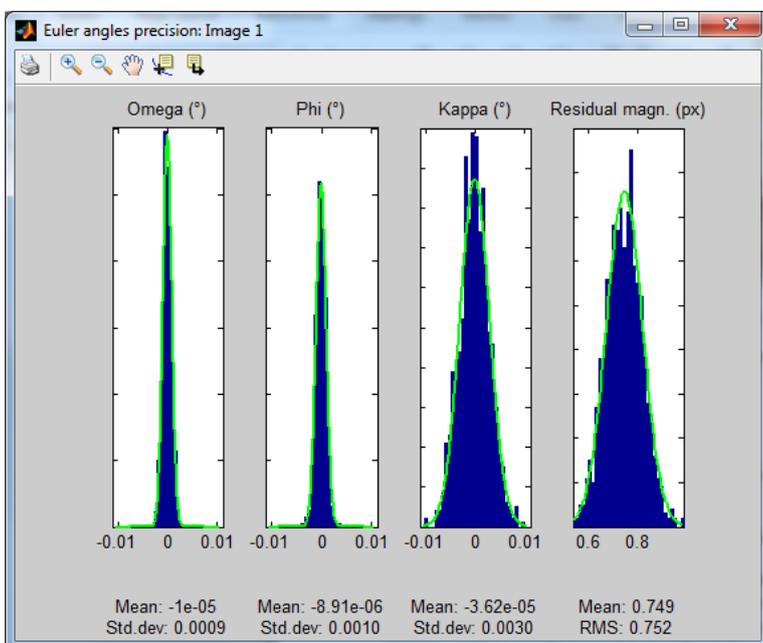
Define measurement error from image residuals: check this box to take the RMS of the offset applied to the point coordinates from the RMS error on the control points. If it is unchecked, then the value will be taken from the box below.

Minimum estimated measurement error (std., px)

This is a minimum RMS value that will be used (i.e. even if the RMS on the control points is smaller for that image). NOTE: this value is the error magnitude, and will be proportioned into random x and y components. If you want to guarantee that the value(s) used are from the error on the control points, set it to zero.

For images / to : Select the range of images to carry out the analysis on.

When you hit 'OK', the analysis will run (it can be slow for large numbers of simulations) and results will be shown sequentially for each image:



Histograms show the distributions of changes in recovered camera angles, with associated Gaussian curve fits. If a distribution passes a Chi-squared test for normality, the curve fit is shown in green, fails are given in red.

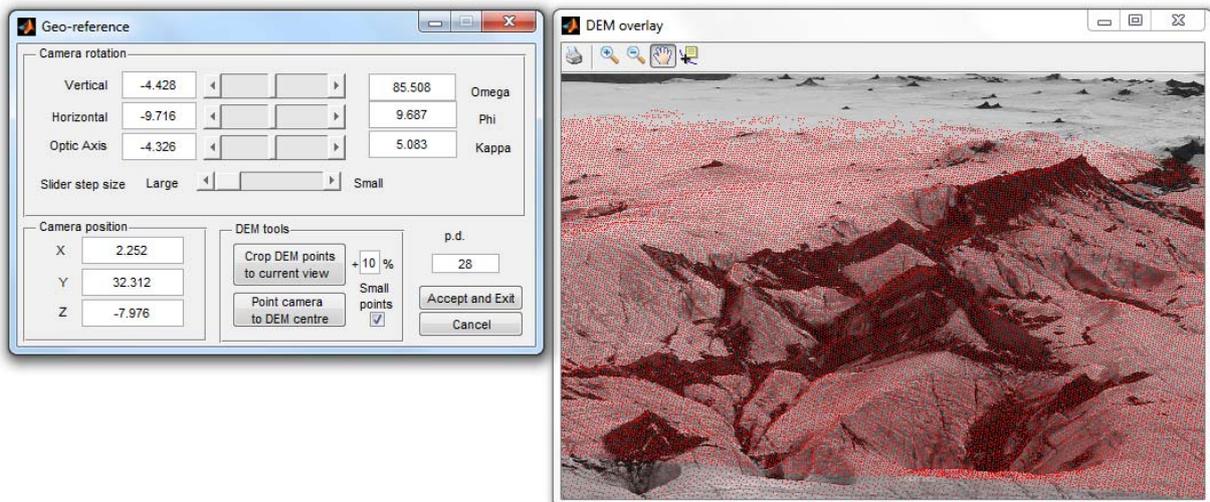
Mean values (except the residual magnitude) should be small, indicating that the distributions are centred on the actual orientation values.

The precision estimates, as given by the standard deviations, are stored in saved Pointcatcher files, and can be exported from the figure via the  toolbar button.

## 9. Sequence geo-referencing

Sequence geo-referencing can be carried out if a DEM is available and the camera position coordinates are known. Geo-referencing is achieved by visually aligning the image to a projection of the DEM points. The only currently supported format for DEMs is as an ASCII file with 3 columns, of x, y, z point coordinates.

- 1) Make sure that the image you want to use for geo-referencing is being shown.
- 2) Tools → Sequence geo-referencing → via DEM
- 3) If a DEM has not already been loaded for this image, then a dialog box will come up for selecting the DEM file to use.
- 4) If DEM import is successful, two new windows will appear, one with controls and the other showing an overlay of the DEM points (in red) over the image.
- 5) Enter the camera position coordinates in the lower left box. If no DEM points are shown then the camera is looking in the wrong direction – the default direction is horizontal and due North. If nothing can be observed, the 'Point camera to DEM centre' button can help, otherwise, pan the camera round by moving the 'Horizontal' slider until red DEM points appear. They will be there somewhere....
- 6) Refine the camera alignment to the DEM through changing the camera angles, and using smaller slider step sizes to finalize adjustments. If the points are too large, tick the 'Small points' button.



- 7) When you are happy with the overall orientation, click Accept and Exit

Pointcatcher will have now stored a geo-referencing transform for that image. This transform is linked to other images via their registrations.

## 10. Re-project points to DEM surface

Following geo-referencing, point observations can be re-projected onto the DEM surface, to enable conversion into 3D geographic coordinates. A DEM will be associated with the image that is current when it is loaded.

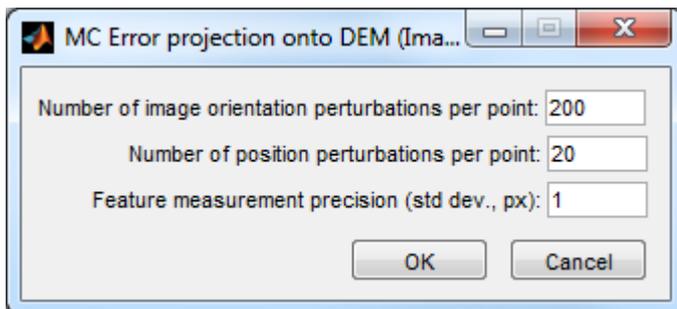
Tools → Re-project pointset to DEM Tools, then select the option you need.

Re-projection is implemented using a graphical approach, aimed at providing accurate coordinates from a triangulated DEM surface, using a generalised solution that fully accounts for occlusions and local horizons. The returned 3D point coordinates are derived by the intersection of virtual rays from the camera with the triangles of the surface. The graphical approach is used for speed, the results are then verified analytically before being returned. For large numbers of points on very large DEMs, the process can be slow.

## 11. Estimate re-projection precision

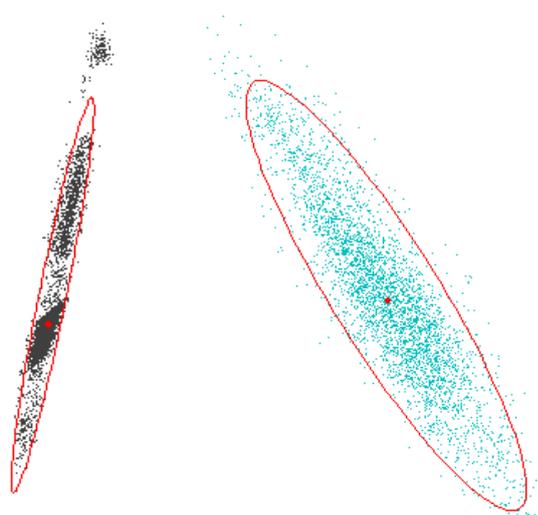
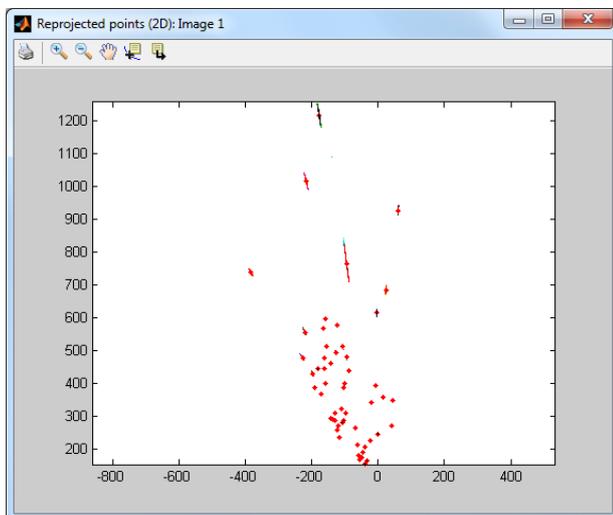
The precision of the geographic coordinates for a particular point (i.e. following re-projection) can be estimated by combining the camera orientation precision with the precision of the point observation, within a Monte Carlo analysis. Note that process currently this does NOT encompass errors in geo-referencing, the DEM itself or the camera model, and uses points from the active Pointset.

Tools → Error analyses → Estimate precision of re-projected points



The re-projection dialog will be pre-populated with values if you have just run uncertainty estimations. In the example here, the position of each point will be estimated for 200 different camera orientations, with 20 different random offsets (of up to 1 pixel) for each orientation – i.e. 4000 re-projections per point.

Once analysis is complete, the results will be displayed in a 2D map form. Large red symbols indicate the position of the 'error-free' points, surrounded by smaller symbols that give the distribution of the point position under the Monte Carlo analysis. These results are fitted by an ellipse to illustrate the error. Ellipse fits represent 95% confidence on the mean for a normally distributed dataset – but note the typically non-normal distribution of points within the ellipses. Ellipses are in green for points whose distribution passes a chi-squared test for normality, and red if they fail.

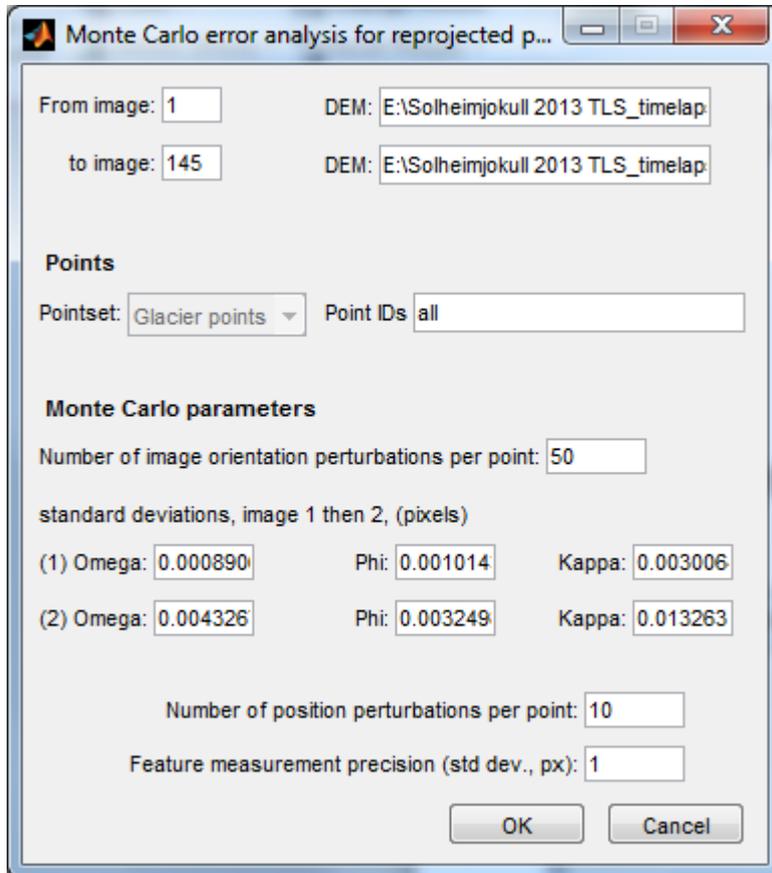


An example output is given above, with two close-up illustrations of point distributions on the right.

## 12. Estimate displacement uncertainties

In order to examine the effect of re-projection uncertainties on the calculation of displacements between two images, Monte Carlo re-projection analyses for both images can be combined.

Tools → Error analyses → Estimate precision of 3D point displacements



Enter the image IDs that you want to run the analyses between, and click in the empty boxes to link to DEMs as required.

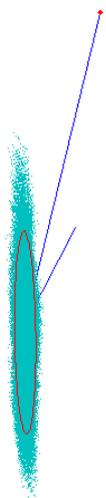
Enter the point numbers you want to analyse. 'all' for all points, or a range:

- 1 - Point 1 only
- 1, 4 - Point 1s and 4
- 1:5, 9 - Points 1 to 5 and 9
- ~1:5, 9 - All points except 1 to 5 and 9

Enter the number of different image orientations to use, and the standard deviations of the distributions they will be taken from.

Enter the number of different point positions, and the standard deviation of the distribution they come from (i.e. the estimated point measurement precision).

NOTE: To determine displacement uncertainty, all combinations of position uncertainty between the two images are combined, leading to a large number of estimates. For example, in the dialog box above, 50 image orientation perturbations are used with 10 point position perturbations. Thus, each image gives 500 positions per point, and therefore 250,000 different displacements between the images. Putting larger numbers in the dialog box can lead to failure of the analysis due to memory (RAM) constraints.



Results are presented in 2D view as displacement vectors from a starting point. The arrow gives the 'no error' displacement vector, and the pointcloud represents the result of the Monte Carlo analyses. **Note that the pointcloud spread is magnified by a factor of 10** for visibility. Pointclouds are fitted by ellipses in an attempt to describe the error but they are seldom normally distributed. Distributions that fail the chi-squared test for normality are given as red ellipses, if they pass, the ellipse is green.

### **13. Calculate intermediate points**

In applications such as the glacial example on the cover page, two DEMs may exist, but 3D point positions may be required for all images in the sequence. These positions can be calculated if 3D points are first derived for the times of the two DEMs, and if the points can be assumed to move in planimetric straight lines (i.e. they travel within fixed vertical planes).

- 1) Make the image contemporaneous with the first DEM the active image, then re-project the point observations from that image onto the DEM (see section 10; the DEM will be imported as part of the process).
- 2) Navigate to the image contemporaneous with the second DEM, and repeat the process for this image and DEM.
- 3) To calculate 3D point coordinates for all other image observations Tools → Reproject pointset to DEM → Derive intermediate 3D points
- 4) The results will be displayed as a 3D plot, and can be exported from the figure.

## **14. Bugs and requests**

Please report any bugs ( [m.james@lancaster.ac.uk](mailto:m.james@lancaster.ac.uk) ) and I will do my best to resolve them. Suggestions and requests also welcome.