The Pixel Farm®

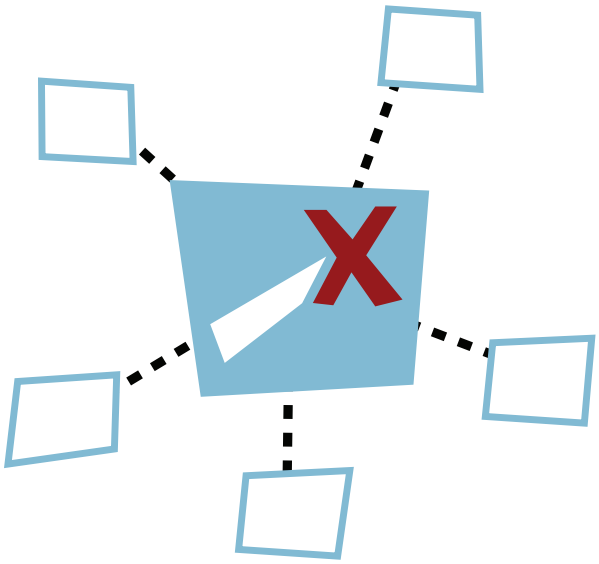

PFTrack™ 5.0

Version 5.0 Reference Manual

www.thepixelfarm.co.uk

PFTrack™ v5.0 - Reference Manual V 1.0

The Pixel Farm Ltd
67 Bexley Street,Windsor,Berks, SL4 5BX, UK

www.thepixelfarm.co.uk
support@thepixelfarm.co.uk

Contents

Introduction

Part 1 - The User Interface

Menus

Dialog Windows

Popup Menus

Part 2 - Camera And Footage Options

Part 3 - Tracking and Solving

Tracking

Scene Orientation

Object Tracking

Geometry Tracking

Stereo Cameras

Test Objects

Trouble Shooting

Exporting Data

Part 4 - Beyond Tracking

Optical Flow Analysis

Motion Capture

Z Depth Extraction

Image-based Modelling

Exporting Data

Part 5 - Keyboard Short Cuts

Part 6 - Scripting in PFTrack

Introduction

Using this Manual

Here you will find full details of all the capabilities and features of PFTrack. This manual has been designed so you can dip in and out as you need to. Issues discussed in the Getting Started manual are expanded in more detail with each tool and interface element within the software fully explained. Once you are familiar with the basics, it is here that you find how to get the most out of the software. The reference guide is organised in logical function groups to make it easier to find your way around.

If you run into problems, please consult this manual before contacting our support team. Our web site also features an FAQ section with additional information.

Conventions used in this Manual

As a rule most functions within PFTrack are identical on all operating systems. Where functions are specific to an operating system, the differences will be explained. Where screen shots are used they may appear slightly different on other operating systems. However, unless otherwise stated the options contained within the screen shot will be the same across all platforms.

Menu items are indicated as Italic text, as are options within Menus. The word Folder is used throughout the manual and refers to a folder in Mac OSX or Windows, or a directory in Linux.

All numerical edit boxes within PFTrack can be adjusted by clicking with the left or middle-mouse button and dragging left and right within the edit box to increase or decrease the value.

Part 1 - The User Interface

Finding your way around

As explained in the Quick Start manual the interface is split up into logical areas. The five main areas to note are:

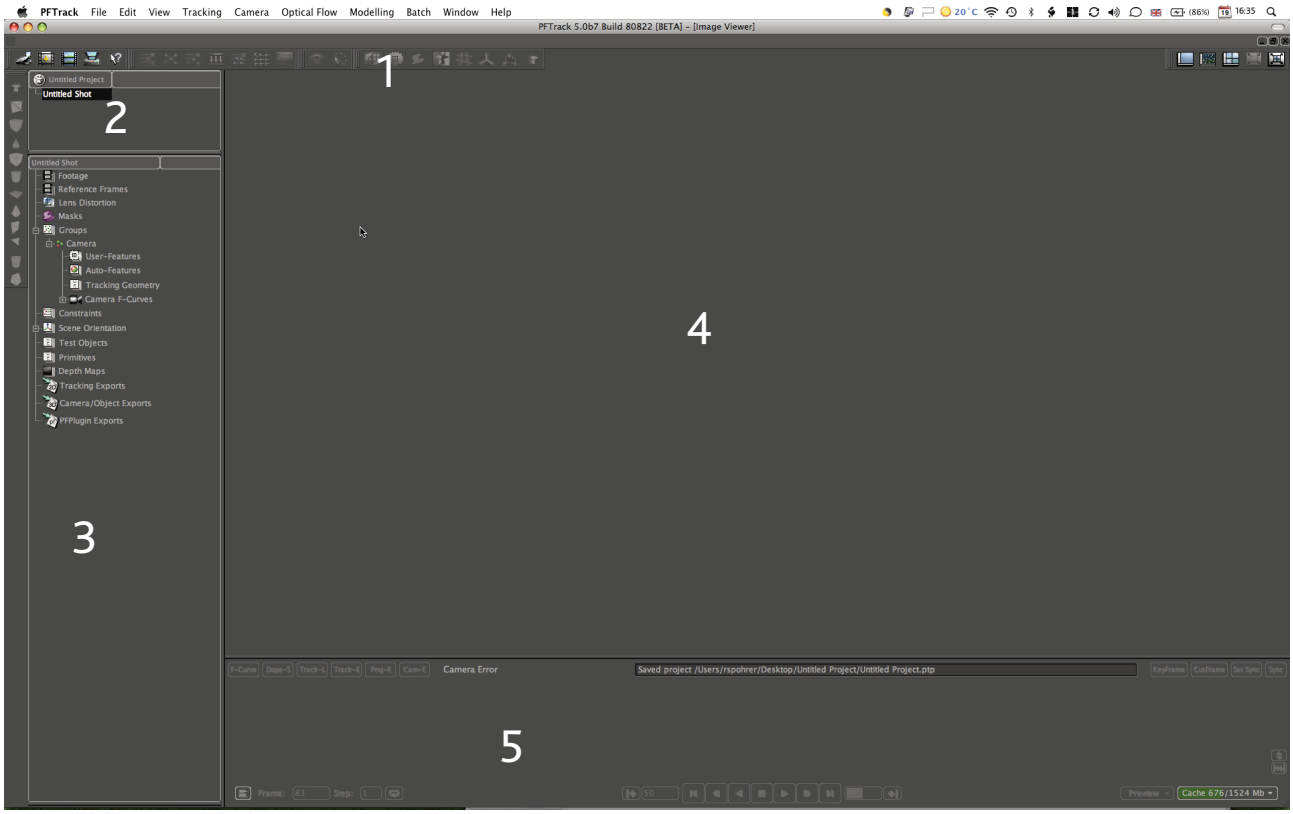

1. The Toolbar located at the top of the screen provides single click access to the most commonly used functions within the application. It also contains viewing options for the main windows along with options for window creation and layout. The tools are grouped into logical sets and each of these sets can be re-ordered or torn off the menu bar as floating palettes if required.

2. The Project Overview window lists all the shots within a project in the order in they were imported. Right clicking on a shot allows you to perform a number of actions including renaming, deleting and locking a shot. You can also re-order shots within the Project Overview using drag-and-drop.

3. The Shot Overview window lists all the elements and parameters relating to the current shot that is being worked on.

4. The main Image Viewing window appears grey when no footage is loaded, or shows the current footage when a shot is open. It can be split into a number of windows to provide orthographic and perspective views of the scene once it has been tracked.

5. The Timeline is located at the bottom of the screen and contains a number of elements: the timeline itself, playback controls, track data view buttons, sync controls, cache status and footage preview buttons. Also in the bottom left-hand corner is the log window button.

To quickly find out what something is you can use the 'What's This?' help system. This is activated by pressing the *'Shift'* and *'F1'* keys, and then clicking on the element or item you are interested in. A brief description of what the item is and does well be displayed. Hovering the mouse over an item will activate the contextual help system, if available, giving a short description of its function.

Project Overview

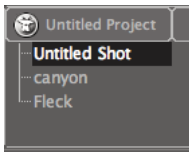

The Project Overview shows the name of the current project and a list of the different shots contained in the project. A shot is selected by clicking on its name, causing the shot to become active and information for that shot to be displayed in the Shot Overview (described below). The image and 3D viewer windows will also show the footage and scene for the current shot.

You can add as many shots into the project as are needed. For example, multiple versions of the same footage can be loaded into different shots to test different methods of tracking or camera solving. Right-clicking on a shot name will display a popup menu that allows a shot to be Locked (so not edits can be made), Renamed, Duplicated, Reset to its original state containing only the footage, Closed or Deleted (which will also delete the shot folder on disk).

To rename rename the project, use *File>Save Project As..* to specify a new name and save location.

Shot Overview

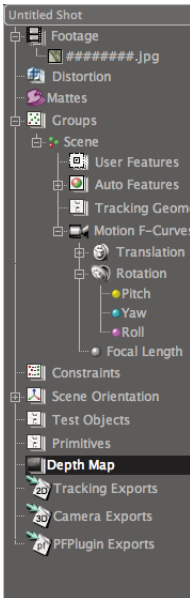

The lower window on the left hand side of the interface is the Shot Overview. The shot overview gives details of any footage, masks, user-features, tracking geometry etc.. that have been imported or created in the shot. Clicking on an individual mask or user-feature in the shot overview will select and highlight this item in the image window, and allow it to be manipulated or edited.

Right-clicking on many items or containers will display a popup menu containing related options. For example, to create a new tracking export, right-click on the 'Tracking Exports' container and choose New... from the popup menu. Right-clicking on the 'Auto-Features' container will display a popup menu allowing the list of auto-features to be sorted. It is also possible to delete many items from the shot by selecting the item with a click of the left mouse button and pressing the 'Delete' key, or by right-clicking and selecting Delete from the popup menu where available.

Many items can be renamed by double-clicking on them. The first click will select the item, and the second will allow you to edit the name. Branches of the shot hierarchy can be opened or closed by clicking on the small open/close indicators at each level.

Clicking items in the shot overview will place the PFTrack interface into an editing mode according to what is clicked. For example, clicking on the Scene Orientation item will perform the same task as pressing the Scene Orientation toolbar button.

Main Image Window

An Image Window is where the footage in the active shot is displayed. It is possible to have more than one image window open at the same time, as is often desired when working on a motion capture shot or with stereo cameras, but only one of the windows is ever active at one time. The currently active window is shown by a change in colour of the window title-bar, and can be changed by clicking in the title bar of a different image window. Image windows can be moved and resized if necessary, and can also be minimized or maximized by clicking on the standard title-bar buttons.

If you click and hold the middle or right mouse button inside an image window and move the mouse, the footage is zoomed or panned within the image window. These controls allow you to zoom and move the image to any position you wish. By clicking on the Fit to Window button the viewpoint will be reset back to the default position. The direction of zoom can be set in the Interface tab in the Preferences window (*File>Preferences...* or *PFTrack>Preferences...* on Mac OSX).

Clicking on the 'New Window' or 'New Viewer' buttons will open additional windows, and in this way any number of views can be displayed at the same time. When working with multiple pieces of footage, you can change the footage that is displayed in the active window by clicking on the footage name in the shot overview. The footage overview will also highlight the currently active footage when more than one piece of footage has been loaded.

Clicking the 'Tile Window Layout' button will arrange your windows within the main workspace. Repeated tilings will switch between alternative layouts. The 'Full Screen' button will toggle a full-screen interface. The 'Tile Window Layout' button can be clicked again after switching into full screen mode to reset the window layout, if desired, and the 'Fit to Window' can be used to reset the viewpoint within the active window. These options are also available from the Window menu.

3D Perspective and Orthographic Views

In addition to showing footage, perspective and orthographic views can also be created within the main image window. Clicking the 'New Viewer' button will open a new perspective view. This view can be switched to a standard orthographic view by using the drop down menu in the bottom left of the window. Once you have tracked and solved your camera, these windows will display the camera path and 3D feature positions. In the picture below, the left image shows the footage whilst the right image is the 3D scene viewer.

The dots, which are colour coded to indicate quality, show the 3D position of features as calculated by PFTrack. The red pyramid and white line in the right hand picture is a representation of the camera's field of view and motion path. A low-resolution image is also displayed in the camera frustum. This can be disabled by switching off the *View>Frustum Image* menu option. The orientation of the viewpoint in the perspective window can be changed using the mouse: left mouse button rotates the view around the origin point, middle mouse button zooms and right mouse button pans the view.

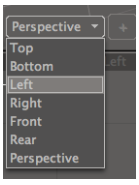

3D viewer windows will stay on top of image windows by default, although this behaviour can be changed from within the Interface tab of the preferences window. There are also options in the Display tab to control the size of the coordinate axes vectors and ground plane, as well as the size of the camera frustum and transparency of the image.

You can switch to a different orthographic viewpoint, or back to perspective at any time using the drop down menu in the bottom-left of the window. When in Orthographic mode, the '+' and '-' buttons allow you to alter the level of detail of the reference grid displayed in the window. The Near and Far edit boxes will be enabled when the *View>Clip* Planes menu option is set, and allow the near and far clip planes to be adjusted.

Icon Toolbar

The toolbar contains button for the more commonly used functions in PFTrack, making them easier to find and operate without having to access the main menus. By clicking on the vertical bars separating the different sets of toolbar buttons it is possible to move each block or detach them completely from the main interface. This allows you to place the buttons where you wish and customize your desktop interface. By default the Modelling Primitive toolbar is not displayed, but this can be switched on and off by selecting the Modelling>Primitive Toolbar menu option. Popup tool-tips can be activated by hovering the mouse over any button, providing information about its function.

Right-clicking in an empty part of the toolbar will display a popup menu allowing individual sets of buttons to be hidden or shown.

This is the icon for 'Open Project'. Clicking the icon enables you to open a previously saved project. If you are currently in a project that has been modified since it was last saved, you will first be asked if you want to save before it is closed and the new project opened.

This is the 'New Shot' icon. By clicking the icon PFTrack will add a new shot to the current project. You can have as many shots within a project as you wish.

This is the 'Import Footage' icon. If your shot does not yet have any movies or image sequences loaded then you can click this icon to load the footage you wish to work with. Once footage is loaded, the icon will be greyed out because each shot can only contain one piece of footage. The only exceptions to this rule are for motion capture and multi-image calibration, where multiple pieces of footage can be loaded into a shot at the same time.

This is the 'Save Project' icon. Clicking this icon will save the project to your hard disk.

This is the 'What's This?' icon. By clicking on this icon and then on any item in the PFTrack interface, a small description of what that item does will appear.

This is the 'Auto Track' icon. By clicking this icon, PFTrack will perform auto-feature tracking in the current shot using the parameters set in the *Tracking>Tracking Preferences*.. window.

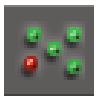

This is the 'Solve Motion' button and will only become available after user or auto-features have been tracked. Clicking this button will display the Solver Controls window and allow PFTrack to solve for camera or object motion using the feature tracks.

This is the 'Track and Solve' icon. By clicking this icon, PFTrack will perform an auto-feature track and then run a camera solve using the default Solver Control parameters.

This is the 'Improve Solution' icon. Clicking this button will open a window from which the current camera and object motion can can be improved after feature editing. This provides a way of fine-tuning an existing solve without starting from scratch.

This is the 'Track and Optical Flow' icon. By clicking this icon PFTrack will perform an auto feature track and then perform optical flow analysis on the current shot, using the auto-feature tracks as flow guides.

This is the 'Optical Flow' icon. By clicking this PFTrack will perform optical flow analysis on the current shot.

This is the 'Depth Map' icon. Clicking it will open the Depth Map parameters window and allow a per-pixel depth map to be generated. This is only available once the camera motion has been solved.

The next icons can be used to put the PFTrack interface into a particular mode of operation. PFTrack will automatically switch between modes in response to user selections. For example, if the user clicks a mask in the shot overview, then Mask Mode becomes active. Similarly, if a user-feature is selected in the shot overview, then User-Feature Mode becomes active.

This is the 'Navigation Mode' icon. Navigation mode allows the image in the currently active image window to be scaled and translated, or the 3D viewpoint in a viewer window to be adjusted. Navigation is the default mode. If you are in a different mode, you can temporarily put the interface into navigation mode by pressing and holding the 'n' key (or whatever shortcut has been assigned). Releasing the 'n' key will revert back to the previous mode.

This is the 'Selection Mode' icon. Selection mode allows tracked points and masks to be selected in the image window. When this mode is active, press and hold the left mouse button to draw a lasso around the tracked points or elements you wish to select. If you are in a different mode, you can temporarily put the interface in to selection mode by pressing and holding the 'l' key. Releasing the key will revert back to the previous mode.

This is the 'Lens Distortion Mode' button. This mode allows distortion lines to be drawn in the image window. Distortion lines are used to correct for geometric distortion introduced by the cameras lens. To draw a distortion line click and hold the left mouse button to place the initial point, drag the mouse and release the left button to place the end point. The distortion line can then be subdivided by clicking and dragging new points to fit a curve in the image. For more information on correcting distortion and exporting the distortion data, see the Lens Lens Distortion section of this manual.

This is the 'User-Feature Mode' button. This mode allows new user-features to be created and existing user-features to be edited. See the User-Feature Tracking section for more information on how user-features can be employed to help track shots. When this mode is active and no other user-feature is selected, a left click on the image will create a new user-feature at that location. User-features can be selected by hovering the mouse over them and clicking with the left mouse button. A selected user-feature can then be adjusted by clicking and dragging with the left mouse button. Holding the 'Shift' key whilst placing a user-feature will display a small popup window with close-up view of the image, providing a clearer view of where the user-feature is being placed.

This is the 'Mask Mode' icon. This mode allows new masks to be created and existing masks to be edited. See the Masks section of this manual for more information on how masks can be used. When this mode is selected, left clicks in the image will place points to define a mask boundary. Clicking the right mouse button will finalize the mask.

This is the 'Colour Key' icon. This mode allows you to quickly create a mask from a particular colour present in the footage, and is useful when tracking green/blue screen shots. For more information, see the Masks section.

This is the 'Coordinate-Axis Line Mode' icon. This mode allows coordinate axis lines to be created to help orient the scene and allow for focal length estimates to be made on still or moving images. Axis lines are drawn in the same manner as distortion lines, i.e. hold down the left button to place the initial point, drag the mouse and release the left button to place the end

point. Before drawing an axis line, select which axis you wish to draw in the shot overview by clicking the “X Axis”, “Y Axis” or “Z Axis” item in the Scene Orientation container. Alternatively, it is possible to change axis line assignment once a line is drawn by right clicking on one of the end points and selecting the appropriate direction from the popup menu.

This is the ‘Scene Orientation Mode’ icon. When selected it allows you to rotate, scale and translate the ground plane of your scene using the left, middle and right mouse buttons respectively. To constrain the rotation to be about the X, Y or Z axes, press and hold the ‘x’ ‘y’ or ‘z’ keys respectively while moving the mouse. When in this mode, buttons will appear along the bottom of the image or 3D viewer windows that allow translate, rotate and scale manipulators to be created as well as a “Fit” button to fit the ground-plane to selected feature points. This icon is only available after camera motion has been solved.

This is the ‘Camera Orientation’ icon. After solving for camera motion, this button performs a similar function to the 'Scene Orientation' mode, but allows the camera to be transformed using the left, middle and right mouse buttons rather than the ground-plane.

This is the ‘New Test Object’ icon. After solving for camera motion you can place a test object into the scene to check the quality of the camera path when playing through the sequence. Clicking this toolbar icon will display a window allowing you to add various 3D objects or load new Wavefront OBJ models to use as test objects. Once an object is created, pressing the ‘Trans’ 'Rotate' or 'Scale' buttons at the bottom of the image window will create manipulators to adjust the object's position. You can also place a test object at a specific feature location by right-clicking on a feature in an image window and selecting New Test Object from the popup menu. Existing objects can be snapped to a specific feature point by right-clicking on the feature and choosing Snap Selected Object.

The final set of icons are for window management:

This is the ‘New Window’ icon. When clicked, PFTrack opens a new image window with footage displayed in it, if any has been loaded. When working in Motion Capture mode this allows you to view several image sequences at once in order to synchronise them. You can change the footage displayed in the active window by clicking on the footage name in the shot overview.

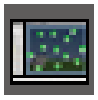

This is the ‘New Viewer’ icon. This opens a new window showing a perspective view of the scene. Once opened, a perspective window can be switched to an orthographic view using the drop down menu in the bottom left-hand corner of the window.

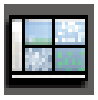

This is the ‘Tile Windows’ icon. When clicked, PFTrack re-arranges the open windows so that they can all be seen. Repeated clicking will switch between tiled, horizontal and vertical layouts.

This is the ‘Fit To Window’ icon. By pressing this icon the footage in the current image window will be re-sized so that it fills the window. If the View>Actual Size menu option is specified, the footage is displayed at its original size, where one pixel in the footage covers one pixel on the screen.

This is the ‘Full screen’ icon. By pressing this icon the interface toggles between windowed and full screen modes.

Timeline

The Timeline is located at the bottom of the PFTrack interface. The Timeline only becomes active once you have imported a movie or image sequence into your shot. You can use the standard playback control buttons to view your sequence or the left and right cursor keys to move forwards/backwards by the number of frames set in the 'Step' edit box (which defaults to 1). The up cursor key takes you to the out-point or last frame, and the down cursor key takes you to the in-point or the first frame. Clicking the left mouse button in the timeline caused the frame at that location to be displayed*. Holding the left mouse button down and dragging the mouse left and right allows you to scrub through the sequence, and pressing the space bar will start or stop playback. You can also scrub through the timeline by holding the appropriate keyboard shortcut (backslash, '\' by default) and moving the mouse pointer left or right.

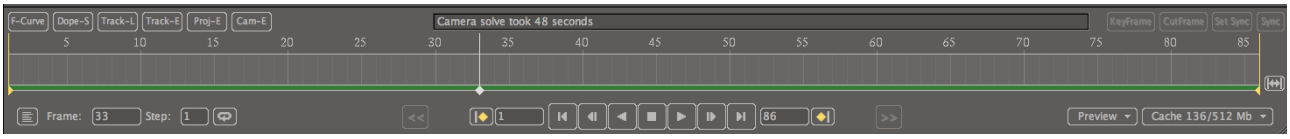

* Note that when one of the timeline mode buttons is active, you must click in the cache indicator underneath the timeline to change frame.

It is possible to set in and out points in the timeline to limit the number of frames PFTrack plays. This will also limit the frames that are considered for tracking and camera solving. To set an in-point (the frame PFTrack will start playback and tracking from), move to the desired frame and then click on the 'In-Point' button, or enter the frame number for the in-point into the edit box next to the In-Point button. Either of these will move the left yellow indicator bar to this frame. The same process is used to set an out-point, but this time you must use the 'Out Point' button or edit box. You can also change the in/out points by clicking on the yellow indicators with and dragging with the left mouse button.

If the sequence you load is very long, you may wish to zoom the timeline to concentrate on a particular area of interest. With the pointer over the timeline, press and hold the middle mouse button and move the mouse left and right to zoom in and out of the displayed frame range. Press and hold the right mouse button and move the mouse left and right to pan the displayed frame range.

Pressing the button shown here on the left switches between displaying the entire shot and the section between the in and out points.

Below the timeline are the playback controls. You can change the playback method by cycling the loop button to choose between Loop, Bounce, and Play Once.

This is Loop playback. With this option selected, the sequence will loop back to the in-point once playback reaches the out-point. With Bounce selected, playback will reverse direction once the out-point is reached, playing backwards towards the in-point. Play Once will halt playback once the out-point is reached.

The remaining buttons are as follows:

Go to the in-point of the sequence or, if already at the in-point, the start of the sequence.

Go backward by the number of frames set in the Step edit box.

Play backwards, from higher to lower frame numbers, in increments of 'Step' frames.

Stop playback.

Play forwards, from lower to higher frame numbers, in increments of 'Step' frames

Go forward by the number of frames set by 'Step'.

Go to the out-point of the image sequence or, if already at the out-point, the end of the sequence.

There are also some other icons and text boxes that comprise the timeline. These are, from left to right:

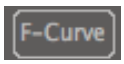

F-Curve displays the tracked camera or object motion parameters as an editable graph in the timeline. Clicking on camera F-Curves such as focal length, translation and rotation in the shot overview will open the F-Curve editor in the timeline and display the function curves (F-Curves). The graph window can be enlarged by dragging the horizontal bar between the timeline and the main image workspace. F-Curve editing is described in more detail later in the manual.

Dope-S displays the Dope Sheet in the timeline. The Dope Sheet is used to indicate where objects or parameters are keyframed, and the contents depends on the active mode and item that is selected.

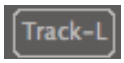

Track-L displays a graph indicating which frames are spanned by each feature. Frames are shown in red when they have too few features available to solve for the camera motion. When the Track-L graph is open, additional buttons appear on the right of the timeline allowing you to order the features by starting frame number ("O" button), or toggle between showing all features or only the selected features (the "A" button).

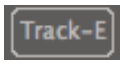

Track-E displays a graph showing the tracking errors for each feature. The tracking error measures the difference between the pixel window surrounding each feature from one frame to the next. A draggable threshold line is available that corresponds to the Tracking Threshold value in the *Tracking>Clean Tracks..* window. An additional button is available (the "A" button) to toggle between displaying all features or only the selected features. This graph can be used to identify features that track poorly from one frame to the next.

Proj-E displays a graph showing the projection errors for each feature, and is available after the camera solve. The projection error is the difference between the projection of the 3D feature point into the camera and the 2D track position. If the feature is solved perfectly in 3D, its projection error will be zero. A draggable outlier threshold line is available that corresponds to the Outlier Threshold value in the *Camera>Clean Features...* window. An additional button is available (the 'A' button) to toggle between displaying all features or only the selected features. This graph can be used to identify feature points that do not solve correctly in a particular frame.

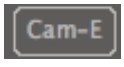

Cam-E displays graphs showing the overall camera error, which is a measure of the total projection error for features in each frame. Two graphs are displayed: one showing the total RMS (Root-Mean-Squared) error for all features in a frame, and the second showing the RMS error for the inlying features (those whose projection errors are below the Inlier Threshold value in the Calibration tab of the Preferences window). These graphs can be used to identify frames where badly solved feature points are affecting the overall camera solve.

When displaying one of the graphs described above, the timeline can be zoomed and panned vertically by holding the middle or right mouse buttons and moving the mouse up or down. To restrict zooming or panning to one direction only, hold either the 'X' or 'Y' keys whilst dragging the mouse. An additional button also appears on the right of the timeline to fit the vertical height of the graph to the current parameter.

The 'Log Window' button displays the text that PFTrack writes in the course of its calculations. Last entry in the log file is shown in the small log window above the timeline.

The current frame number

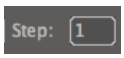

The Step value, which defines the frame increment during playback and jogging.

The current in-point.

The current out-point.

This is the Footage Preview button, which can be used to store a low-resolution, compressed version of the footage in RAM to enable fast playback of long sequences. Clicking this button will display a popup menu allowing the preview to be enabled or cleared, and the preview status to be viewed from the Performance tab in the Preferences window.

PFTrack can use a memory cache to store frames of the footage in RAM in order to speed up playback. A green mark under a frame in the timeline indicates that that a particular frame is stored in the memory cache. By default, the cache is not enabled, but it can be switched on in the Performance tab in the Preferences window, which is easily accessed by clicking on the cache status button at the bottom-right and selecting 'Properties' from the popup menu.

The Cache status shows the amount of cache used and total cache size in megabytes (Mb). The pull-down menu on this button allows for the memory cache to be emptied and for the cache properties to be set in the Performance tab of the Preferences window.

The 'KeyFrame' button is used to manually place keyframes for solving the camera path. This is described in more detail in the section on camera solving.

The 'CutFrame' button is used to place cut-frames at the points where the motion of the camera changes. This is used to segment the camera motion into different types, and is described in more detail in the section on camera solving.

The 'Set Sync' and 'Sync' buttons are used with motion capture data and are explained in the motion capture section of this manual.

Useful Keys

A full list of keyboard shortcuts is given at the end of this reference manual. These shortcuts can be changed from within the Keys tab of the Preferences window.

Left cursor key: Move backwards by the number of frames set in the Step edit box.
Right cursor key: Move forwards by the number of frames set in the Step edit box.
Up cursor key: The up arrow takes you to the out-point or, if already at this point, the end frame.
Down cursor key: The down arrow takes you to the in-point or, if already at this point, the start frame.
Escape key: Whilst tracking user-features or geometry hitting the escape key at any time will terminate the track. The escape key can also be used as a shortcut to the 'Abort' button whenever a progress window is displayed.
Space bar: Start/stop movie playback.
N key: Hold and release to toggle in/out of Navigation mode
L key: Hold and release to toggle in/out of Selection mode
Backslash key ('\'): Hold to enabled scrubbing in the timeline using left/right mouse motion. Sensitivity can be controlled from the Interface tab in the Preferences window.

By clicking and holding the middle or middle mouse buttons in any numerical text entry box and then dragging the mouse left and right, the value in the text entry box can be adjusted without having to use the keyboard.

Menus

Most tools and functions within PFTrack can be found either in one of the menus across the top of your screen, or in popup menus displayed with a right mouse click on the item of interest. The most common options within these menus are also found in the toolbar. Items are grouped in logical families, so Camera related functions will be found in the *Camera* menu.

File menu

New Project: Opens a newproject.
Open Project: (Ctrl+O): Opens an existing project.
Save Project: (Ctrl+S): Saves the current project.
Save Project As...: Saves the current project under a different name.

New Shot: (Ctrl+N): Creates a new shot.
Open Shot: Opens an existing new shot.
Duplicate Shot: Makes a copy of this shot in this project.
Reset Shot: Removes everything from this shot except the imported footage.
Close Shot: Close the current shot without deleting it from the project.
Delete Shot: Delete the current shot, including the shot folder on disk.

Import Footage...: (Ctrl+I): Loads an image or movie into the current shot.
Import Reference Frame...: Loads a reference frame into the current shot.
Import 2D Tracks...: Load 2D tracking positions as user or auto-features. The format of the tracking file must be the same as the files generated by the 2D tracking export plugin:

"feature name"
<number of frames>
<frame number> <x> <y> <residual>

Import Mask Image...: Loads an image sequence or movie to use as a mask.
Import Tracking Geometry...: Loads a Wavefront OBJ file for use as geometry tracking in the selected group.

***Preferences...:** Displays the preferences window.
Recent Files: Displays a list of the most recently loaded files.
***Quit:** Exit out of PFTrack.

*These items are found under the PFTrack menu on Mac OS X

Edit Menu

Undo: (Ctrl+Z): Undo the last change.
Cut: (Ctrl+X): Cuts the selected item to the clipboard.
Copy: (Ctrl+C): Copies the selected item to the clipboard.
Paste: (Ctrl+V): Pastes the item in the clipboard to the selected area.
Delete: (Del): Deletes the selected item.

View Menu

Darken: (Alt+K): Reduces the display intensity of the background image.
Actual Size: (Alt+1): Displays the image sequence at its original size.
Unit Pixel Aspect Ratio: (Alt+S): Displays the image using square pixels.
Frustum Image: (Alt+X) **:** Toggles the display of the footage in the camera frustum of the 3D scene viewer.
Frame Boundary: (Alt+9): Toggles display of the frame clipping boundary.

Axis Lines: (Alt+L): Toggles the display of coordinate-frame axis lines.
Lens Distortion Lines: (Alt+I): Toggles the display of the distortion lines.
Masks: (Alt+M): Toggles the display of masks.
Colour Key: (Alt+H): Toggles the display of the colour key mask.

User Features: (Alt+U): Toggles drawing of user-features.
Auto Features: (Alt+A): Toggles drawing of auto-features.
Centre On Selected Feature : (Alt+N): Keeps the image window positioned so the currently selected user-feature is at the centre.
Feature Tracks: (Alt+C): Toggles drawing of feature tracks.
Feature Labels : (Alt+B): Toggles the display of feature labels.
Search Windows : (Alt+R): Toggles the display of user-feature search areas.
Group Assignments: (Alt+O): Toggles the display of feature group assignment colours.
Constraint Assignments: (Alt+P): Toggles the display of feature constraint assignment colours.
Tracking Geometry: (Alt+2): Toggles the display of tracking geometry.

Optical Flow: (Alt+Q): Toggles the display of optical flow data.

Solved Features : (Alt+Y): Toggles the display of solved feature positions after a camera solve.
Missing Features: (Ctrl+M): Toggles the display of projection points for features that aren't tracked in the current frame.
Projection Errors: (Alt+E): Toggles the display of projection errors after a camera solve.
Ground Plane: (Alt+G): Toggles drawing the ground-plane after the camera solve.
Horizon Line: (Alt+5): Toggles the display of the virtual horizon line.

Group Bounds : (Alt+8): Toggles display of a bounding box around each tracking group.
F-Curve Targets : (Alt+7): Displays the camera corresponding to any F-Curve targets that the user has drawn. This gives you a preview of the camera F-Curves before improving the camera solve or copying motion to the main camera.

Modelling Primitives: (Alt+4): Toggles the display of modelling primitives.
Vertex Guides: (Alt+0): Toggles display of the guidelines that help you determine where modelling primitive vertices should lie in the Image Window.

Stereo Convergence: (Alt+[ ): Toggle displace of the stereo convergence point in 3D viewer windows.
Test Objects: (Alt+3): Toggles the display of 3D tracking test objects.
Clip Planes: (Alt+D): Toggles the display of the current clip planes in 3D viewer windows.
Depth Cue: (Alt+Z): Change the colour of solved feature points according to their distance from the camera.
Depth Map: (Alt+6): Toggles the display of depth map.
Depth Mesh: (Alt+J): Toggles the display of the depth mesh in 3D viewer windows.

Tracking Menu

Tracking Parameters: (Ctrl+F2): Displays the tracking parameters window.

New User Feature: (Ctrl+F3): Creates a new user-feature.
New Mask: (Ctrl+F4): Creates a new mask to assist auto-feature tracking or optical flow.
Colour Key...: Opens the colour key window.

Auto Track: (Ctrl+F5): Automatically select and tracks a number of auto features throughout the shot.
Track and Solve: (Ctrl+F6): Performs auto-feature tracking and camera solving with default parameters in a single step.
Clean Auto-Feature Tracks: Displays the Clean Auto-Feature Tracks window for removing badly tracked features.

Survey Data...: Opens the survey data window.
Groups...: Opens the motion groups window.
Constraints...: Opens the feature constraints window.

Clear User Features: Clears all user-features.
Clear Auto Features: Clears all auto-features

Camera Menu

Camera Parameters: (Ctrl+P): Displays the camera parameter window.
Segment Motion Types...: (Ctrl+Q): Displays the segment motion types window, to allow different motion types to be set for different segments of the camera move.
Lens Distortion...: (Ctrl+L): Displays the lens distortion window.

Estimate Focal Length: (Ctrl+E): Estimates the focal length at the current frame, using manually drawn coordinate axis lines. If no coordinate axis lines are drawn, a focal length estimation tool is created.
Solve Motion...: (Ctrl+F7): Solves camera motion in the current shot, displaying a window containing the Solver Controls.
Statistics...: Opens the statistics window after the camera has been solved, showing numerical per-frame and per-feature data.

Clean Auto-Features...: Displays the Clean Auto-Features window, allowing you to remove poorly solved features.
Improve Solution... (Ctrl+F8): Opens the Improve Solution window, where an existing solve can be improved after feature editing.
Depth Map... : Displays the depth map parameter window.

Orient Scene: (Ctrl+F9): Orient the coordinate system using either user-drawn coordinate axis lines, or so the ground plane fits selected user/auto features.
Scale Scene...: Opens scene scale window, where the distance between any two features can be set to a known value.
Transform Camera...: Opens the camera transformation window.

Un-Solve Frames: Open the Un-solve Frames window, allowing specific frames to be removed from the solution..
Delete Solution: Deletes any camera and object motion data.

Optical Flow Menu

Parameters: Opens the Optical Flow parameter window.
Track and Flow: Performs auto-feature tracking and optical flow analysis in a single step.
Calculate Flow...: (Ctrl+F10): Starts optical flow analysis for the current shot.

Edit...: (Ctrl+F11): Opens the optical flow clean up and editing tool window.

Modelling Menu

Primitive Toolbar: Toggle display of the primitive modelling toolbar in the interface.
Build Mesh...: Builds a geometric mesh by triangulating selected features from the current camera viewpoint.
New Primitive: Open the primitive popup menu allowing the selection of a new primitive modelling object.
Clear Primitives: Deletes all modelling primitives from the current shot.

Extract Textures From: Toggles between texture extraction from the 'current frame' or the 'best frame' within the shot.
Texture Fill Method: Specifies how PFTrack will fill holes in extracted texture maps.
Extract Textures: Extracts textures and maps them to the modelling geometry within the shot.
Clear Textures: Deletes all textures within the shot.

Batch Menu

Open Script: Opens an existing script.
Edit Script...: Edits the current script.
Execute Script: Runs the currently open script.
Clear Script: Removes the current script.

Launch Batch Manager: Opens the Batch Manager window.
Track off-line: Sends the current shot to the batch manager to be tracked.
Track and Solve off-line: Sends the current shot to the batch manager to be tracked and solved.
Track and Optical flow off-line: Sends the current shot to the batch manager for tracking and optical flow analysis.
Solve off-line: Sends the current tracked shot to the batch manager for solving.
Optical Flow off-line: Sends the current shot to the batch manager for optical flow analysis.

Window Menu

New Window: (Alt+W): Open a new Image Window.
New Viewer: (Alt+V): Open a new 3D perspective or orthographic viewer window.

Layout: Displays the popup layout menu, allowing you to select a pre-defined window layout, or create and edit current layouts.
Layout>Store: Store the existing window layout for future use.
Layout>Rename: Rename the active user-created window layout.
Layout>Delete: Delete the active user-created window layout.

Tile Windows: Window (Alt+F): Adjusts the window viewpoint to fit the contents. This will work in Image or 3D viewer windows.
Fit Image To Window: Window (Alt+F): Adjusts the window viewpoint to fit the contents. This will work in Image or 3D viewer windows.

Help Menu

Log Output: Toggles the writing of information to the log file.
Show Log: Shows the current log file. This can also be done by clicking the "Show Log Window" button at the bottom-left of the timeline.
What's This?: (Shift+F1): Put the system into 'What's This?' mode for context-sensitive help. To get context-sensitive for popup menu items, display the popup menu using the right-mouse button, move the mouse over the menu item you are interested in and then press Shift+F1.

Dialog Windows

Import Footage Window

This window is used to load footage into a shot, and is displayed by clicking the import footage button or by selecting *File>Import Footage...* from the menu.

Current folder: Display the current folder you are browsing. You can browse through different folders by clicking on the menu button next to the folder name.
Add Favourite: Add the current folder to the favourite list, displayed underneath the current folder. Clicking on a Favourite entry in the list will move directly to that folder.
Up: Move up one folder from your current location.
Thumbs/List/Detail: Change the displayed information from image thumbnails to a file listing, or file size/creation/owner details.
Preview: Play a preview of the currently selected footage in the window.
Selected file: displays the name of the currently selected footage.
Single Frames: Specifies whether individual files in a folder should be displayed separately even if their frame numbers run in sequence.
From frame: Specifies the first frame to load.
To frame: Specifies the last frame to load.
Load: Loads the currently selected footage.
Cancel: Close the import window without loading any footage.

Note: when importing reference frames that are numbered numerically, such as ref.1.tif ref.2.tif etc.. the Single Frames box will need to be ticked to ensure the individual images are displayed.

Tracking Parameters

This window contains the parameters that control auto and user-feature tracking, and is opened from the *Tracking>Tracking Parameters...* menu.

General:
Window Width/Height: The size of the window that represents each feature (units: pixels).
Similarity: The algorithm used to compare features to decide if they are the same. 'NCC Score' uses the Normalized Cross-Correlation measure. 'RMS Error' uses the root-mean-squared pixel difference.
Group: Defines the group into which auto-features are stored. This can also be changed by selecting a group in the shot overview before pressing the Auto-Track button.

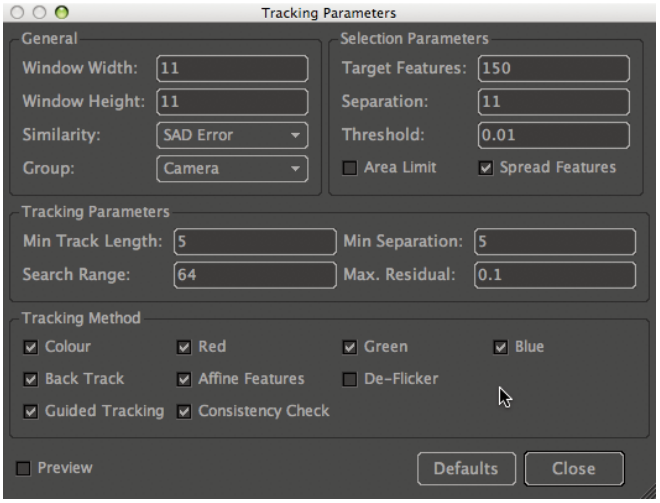

Selection Parameters:
Target Features: The target number of auto-features that should be present in each frame (this is only a rough guideline, and the number of features may not match this exactly).
Separation: The minimum allowed distance between auto-features (units: pixels).
Threshold: The threshold used to determine if an individual pixel contains an 'interesting' feature. Increasing this value will pick more poorer-quality features.
Area Limit: Limits the number of auto features created in each frame to take into account the fraction of the image covered by masks.

Spread Features: Tries to ensure that auto-features are spread uniformly across the image.

Tracking Parameters:
Min Track Length: The minimum number of frames that an auto-feature must track for to be kept.
Search Range: The maximum distance (in pixels) a feature can travel between frames of the sequence.
Min Separation: The minimum distance (in pixels) allowed between two tracked features before they are removed.
Max Residual: The residual level above which a feature is considered to have tracked incorrectly. Increasing this value will allow more poorly tracked features to remain.

Tracking Method:
Colour: Tracking takes account of image colour, rather than just luminance.
Red: Tracking accounts for the colour red.
Green: Tracking accounts for the colour green.
Blue: Tracking accounts for the colour Blue.
Back-Track: Features are tracked backwards during a second tracking pass. This will take extra time, but will also result in a better set of auto-feature tracks.
Affine Features: Tracking uses an affine transformation to allow a small amount of rotation inside each feature window.
De-Flicker: Tracking accounts for global contrast/brightness changes between frames.
Guided Tracking: Estimates of camera motion are used to guide feature tracking. When tracking features for optical flow, if your footage contains moving objects you may want to switch this off to ensure you get auto-feature tracks on all moving objects in the scene.
Consistency Check: When checked, PFTrack uses a slower but more accurate process to identify which features are moving consistently with the scene.

Preview: This uses the current tracking parameters to display a preview of where new auto-features will be placed when tracked. It may help to switch on the View>Darken menu option to see the features more easily.
Defaults: Set all parameters back to their default values.
Close: Close the tracking parameters window.

Note: When tracking features for optical flow, if your footage contains moving objects you may want to switch off Guided Tracking and Consistency Check to ensure you get auto-feature tracks on all moving objects in the scene.

User-Feature Tracking Parameters

This window contains the parameters that control tracking of a specific user-feature. It is opened by right-clicking on a user-feature and selecting the *Edit Parameters...* option from the popup menu.

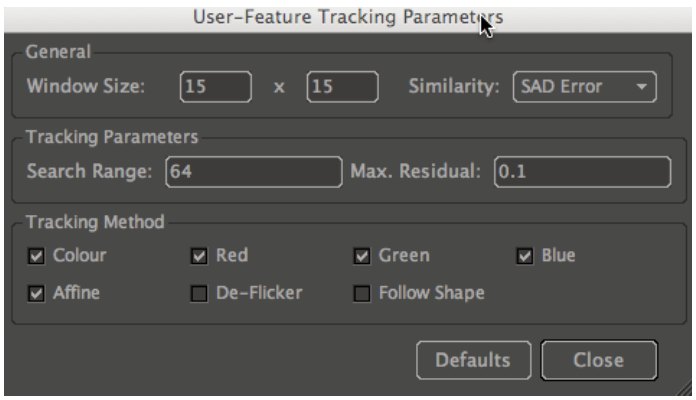

General:
Window Size: The size of the window that represents each feature (units: pixels).
Similarity: The algorithm used to compare features to decide if they are the same. 'NCC Score' uses the Normalized Cross-Correlation measure. 'RMS Error' uses the root-mean-squared pixel difference.

Tracking Parameters:
Search Range: The maximum distance (in pixels) a feature can travel between frames of the sequence.
Max Residual: The residual level above which a feature is considered to have tracked incorrectly. Increasing this value will allow more poorly tracked features to remain.

Tracking Method:
Colour: Tracking takes account of image colour, rather than just luminance.
Red: Tracking accounts for the colour red.
Green: Tracking accounts for the colour green.
Blue: Tracking accounts for the colour Blue.

Affine: Tracking uses an affine transformation to allow a small amount of rotation inside each feature window.
De-Flicker: Tracking accounts for global contrast/brightness changes between frames.
Follow Shape: When selected, the contents of the feature window will be updated after each frame according to the new feature position. This can help track features for longer before the tracking fails, but may also lead to an accumulation of error towards the end of the track.

Defaults: Set all parameters back to their default values.
Close: Close the tracking parameters window.

Colour Key

This window contains the controls for generating a colour key to use as a mask for tracking. This can be useful when tracking green/blue screen shots. The window is opened using the *Tracking>Colour Key...* menu option.

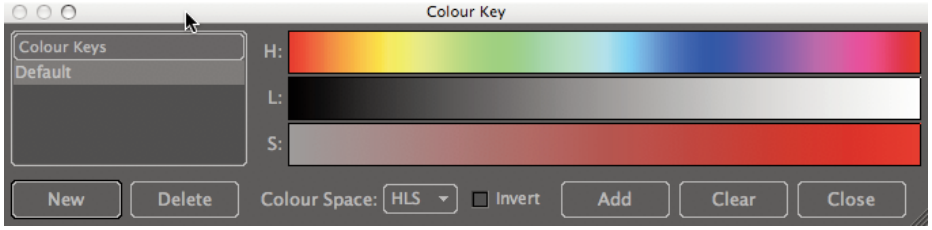

The Colour Keys list on the left shows the available keys. Multiple keys can be used to create a single mask.
New: Create a new key in the colour key list.
Delete: Delete the current colour key from the list.
Colour Space: Change the colour space used to build the key. Choices are HLS (the default), YUV, YIQ and RGB.
Invert: Invert the mask generated from the colour key
Add: Clicking this button will let you sample colours from an image window by clicking and dragging with the left mouse button. As colours are added, the colour bars will update to show the range of colours that are selected.
Clear: Clear the selected colours from the key.
Close: Close the colour key window.

Once colours are added, the vertical lines in each colour bar can be dragged left and right to increase or decrease the range of selected colours.

Clean Auto-Feature Tracks

This window is used to remove poorly tracked auto-features, and is opened from the *Tracking>Clean Auto-Feature Tracks..* menu. The window also provides tools to remove features that exhibit very sharp accelerations.

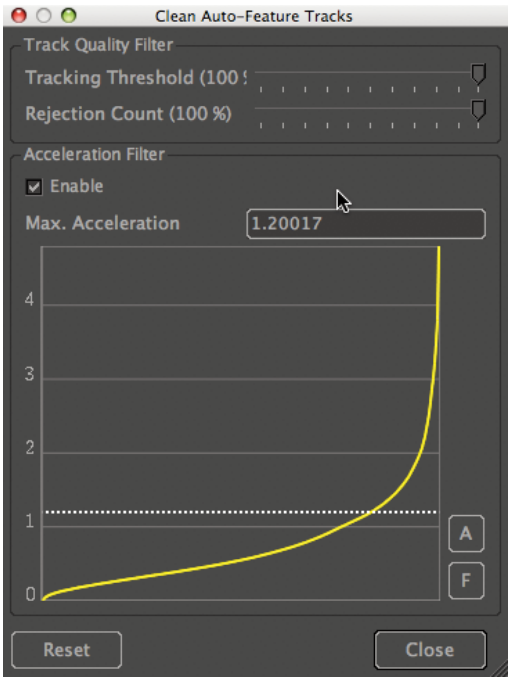

Tracking Threshold: This is the threshold used to decide whether to remove part of an auto-feature track (1 to 100%). A threshold of 100% corresponds to the 'Max Residual' value in the Tracking Parameters window. If the tracking error for a feature in one frame is larger than this threshold, the auto-feature will be removed from that frame.
Rejection Count (1 to 100%): This controls the removal of entire auto-feature tracks. An auto-feature will be removed entirely if the percentage of frames where the Tracking Threshold is exceeded is larger than this value. For example, if this is set to 50% then any auto-features that exceed the Tracking Threshold in more than half of their frames will be removed entirely.

Acceleration Filter: Enabling this filter will display a graph showing the range of feature track accelerations throughout the entire shot. Acceleration is the change in velocity from frame-to-frame, so features that zig-zag incorrectly during tracking can be quickly identified and removed using this tool. Features are plotted in increasing order of acceleration, and the Max. Acceleration edit box specifies the maximum allowable acceleration before the feature is removed entirely. The Max. Acceleration value can also be changed by clicking and dragging in the graph with the left mouse button, and the graph can be zoomed using the middle mouse button.

'A' button: clicking this will automatically estimate a suitable value for Max. Acceleration. Note that the accuracy of this value will depend on the type of motion that the features are undergoing, and the length of the shot.
'F' button: Fit the acceleration graph to the window size.

Reset: Reset the filtering to their default values.
Close: Close the window.

Survey Data

The groups window is used to create additional motion groups for tracking objects separately from the camera. Each separate object (whether it is tracked using feature-points or tracking geometry) must be placed in a different motion group. The groups window is opened from the *Tracking>Groups...* menu, or by right-clicking on the group name in the shot overview and selecting *Edit...* from the popup menu..

Available Features: This lists the features available to use as survey points. If features were selected before the survey window was opened, only those features will be listed here, otherwise all features will be listed.
Add/Remove buttons: These buttons let you add or remove features from the survey list.

Survey table: This table shows survey data, including the feature name, the survey position (Survey X, Survey Y and Survey Z) the current feature position (X, Y and Z) and the survey point type. In the Type column Survey points can be specified as either Exact or Approximate. Exact points are assumed to be measured exactly. Approximate survey points have an associated "uncertainty", so for example, you can specify that the point is approximately at position (1,2,3) with an uncertainty of 0.01. The final column indicates whether a survey point is active or not. Inactive survey points will be ignored during the camera solve.
Clear: Clear all existing survey data from the table.
Close: Close the survey window.

Groups

The groups window is used to create separate motion groups for tracking objects separately from the camera. Each separate object (whether it is tracked using feature-points or tracking geometry) must be placed in a different motion group. The groups window is opened from the *Tracking>Groups...* menu.

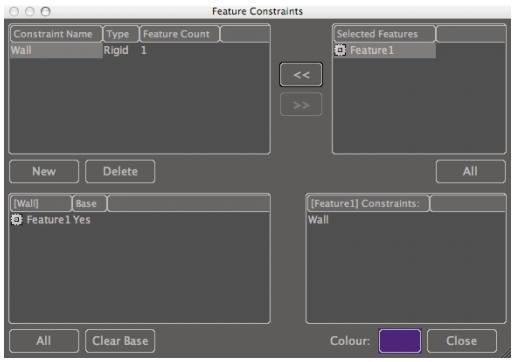

Number of Motions: This allows you to set the number of distinct motions within a shot. For example, if your scene contains a moving camera and a moving object and you want to track both, you would enter '2' in this box.

Group Name: Clicking on a group name will select the group and let you add features to it. Clicking again allows you to rename the group (hit return to confirm any change made).
Motion: Allows different motions to be assigned to a group. To change the motion assigned, right-click in the Motion column and choose an option from the popup menu.
New: Create a new group.
Delete: Delete the currently selected group.
Grouped Features: This lists the features contained in the selected group.
Available Features: Shows all the features available for adding to the group. If features were selected before the group window was opened, only those features will be listed here, otherwise all features will be listed.
All/None: Select/de-select everything in the Grouped Features list or the Available Features list.
Colour: Shows the colour assigned to the current group. The colour can be changed by clicking on the button and choosing a new colour from the colour chooser window.
Close: Close the group window.

Feature Constraints

The Feature Constraints window is used to create constraints that can limit 3D feature positions to a line or a plane, or specify the distance between two features. Feature constraints can also be created for Motion Capture shots to ensure that some of the feature points move as a rigid object. The Feature Constraints window is opened from the *Tracking>Constraints...* menu.

Constraint Name: The name of the current constraint. Constraints can be selected by clicking on them. Clicking again will allow the constraint to be renamed.

Type: The type of constraint. Options are 'None for no constraint, 'Linear' to constrain features in a straight line, 'Planar' to constraint features in a plane, or 'Distance' to constrain the relative dis-

tance between features. For motion capture shots, 'Rigid' constraints are available to constrain features that move rigidly with respect to each other.
Feature Count: The number of features in the constraint.
New: Create a new constraint.
Delete: Delete the selected constraint.
Available Features: Shows all the features available for adding to the constraint. If features were selected before the constraint window was opened, only those features will be listed here, otherwise all features will be listed.
Constraint Features: This list (at the bottom-left of the window) shows all the features assigned to the selected constraint. The 'Basis' column specifies whether a feature will be used to find the support for the constraint. For example, if three features in a planar constraint are marked as "basis" features then these form the plane that all other features will be constrained to. If more than three basis features are selected, the plane will be chosen as the best-fit between all features. To change the basis state of a feature, right-click and choose 'Yes' or 'No' from the popup menu.
Constraints: This list (at the bottom-right of the window) shows the constraints that are applied to the feature selected from the Available Features list.
All/None: Select/deselect all features in the Available Features list or the Constraint Features list.
All Basis/No Basis: Mark all/none of the features in the Constraint Features list as basis features.

Colour: Shows the colour assigned to the current constraint. The colour can be changed by clicking on the button to open a colour chooser window. Constraint colour assignments can be shown in a viewer window using the *View>Constraint Assignments* menu option.
Close: Close the Feature Constraints window.

Camera Parameters

The Camera Parameters window contains the settings that control the behaviour of the camera in the shot. It is opened from the *Camera>Camera Parameters..* menu. On the left-hand side of the window are the various controls and menu options. On the right is a diagram illustrating the current parameters that have been specified. The diagram can be hidden or shown by clicking the Less/More button at the top-right of the window.

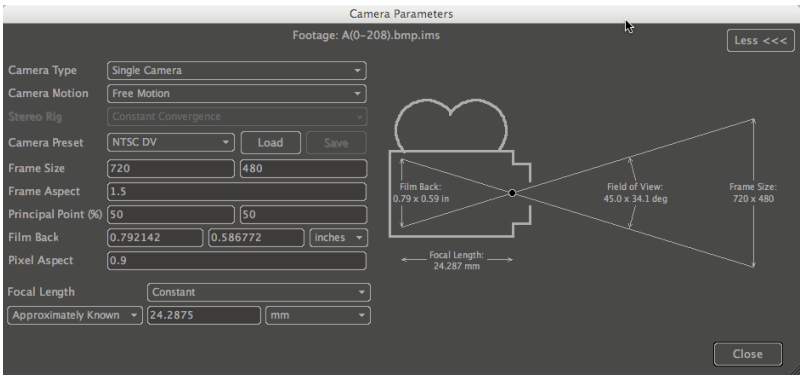

Camera Type: The type of camera used in the shot: Single Camera, Stereo Camera or Motion Capture.
Camera Motion: The motion that the camera undergoes. Options are Free Motion, Translation Only, Rotation Only, Linear Translation + Rotation, Planar Translation + Rotation or Stationary.
Stereo Rig: For stereo shots, this specifies the convergence behaviour during the shot: Constant Convergence, Variable Convergence or Parallel.

Camera Preset: Choose a camera preset for the footage (for example, PAL D1, PAL 16:9 etc.. or custom).
Load: Load a previously saved preset into the list.
Save: Saves the current camera settings as a custom preset. For your saved preset to be picked up automatically when you reload PFTrack, it must be placed in the "presets" folder where PFTrack is installed.
Frame Size: The size of the image in pixels. The first box contains the image width and the second the height.

Frame Aspect: The aspect ratio of the image frame.
Principal Point: Specifies the position of the camera principal point as a percentage of image width/height.
Film Back: Sets the size of the film back area (width and height) in units of either inches or millimetres.
Pixel Aspect: Specifies the pixel aspect ratio for the footage.
Focal Length: Specifies a variable or constant focal length. You can also specify whether the focal length is known, unknown or approximately known, and enter the focal length value and units. If you want to enter a focal length in millimetres, make sure that the camera film back size is specified correctly because otherwise your focal length measurement will not correspond to the correct field-of-view.
Close: Closes the camera parameters window.

Segment Motion Types

This window is displayed using the *Camera>Segment Motion Types...* menu option, and contains motion information for each segment of the camera path. A camera path can be segmented by clicking the 'CutFrame' button at the top-right of the timeline whilst the DopeSheet is open. The table contains one row for each segment.

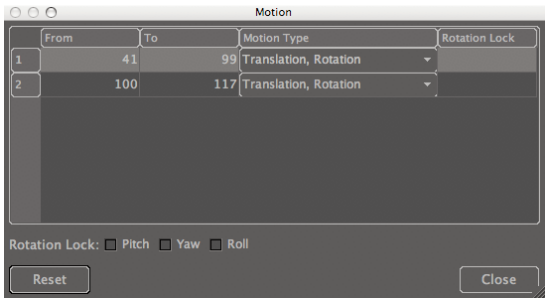

From/To Columns: Shows the start and end frame of the segment.
Motion Type Column: Drop down list for selecting one of the seven camera motion types available within PFTrack. The options are:

Translation + Rotation: this is the same as the main "Free Motion" type;
Translation Only: Camera translation with no rotation;
Linear Translation, Rotation: Straight line camera motion, a dolly shot could produce this motion;
Planar Translation, Rotation: Camera is moving in a flat plane;
Rotation Only: Camera can rotate, but not translate;
Rotation + Small Translation: Useful for shots where the camera is not moving far, relative to the distance from the set;
Interpolate: This will initially ignore the segment during the camera solve, and then interpolate the camera motion after the solve has completed.

Rotation Lock Column: Display the rotation lock state for the segment. Locking rotation for any segment ensures that the camera does not rotate around the appropriate axis.
Rotation Lock (Pitch, Yaw, Roll): Check boxes to set the rotation lock parameter for the current segment.
Reset: Resets the motion type and rotation locks for the all segments to the default settings.
Close: Closes the segment motion window.

Lens Distortion

The Lens Distortion window contains the controls for calculating or specifying the type of lens distortion in the footage. It is opened from the *Camera>Lens Distortion...* menu.

Distortion Model: Choose the distortion model for the shot. The main options are 'None' for no distortion, or 'User-Specified'. Also available are distortion models for Cooke Optics lenses, including the Cooke CXX Zoom lens.

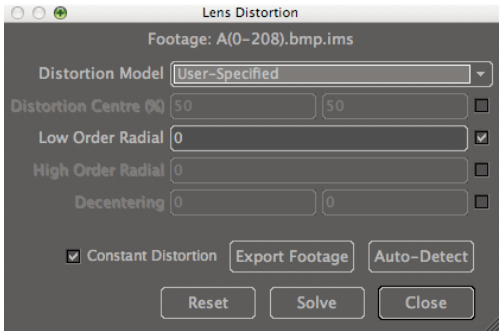

For the 'User Specified' distortion model, the following parameters are available:

Distortion Centre: Specifies the centre of distortion, as a percentage of image size. For most lenses, you can keep this at the centre of the image (50%,50%), but if your image has been cropped then the centre of the lens and hence the distortion may be offset.
Low Order Radial: Specifies the low order coefficient of radial lens distortion.
High Order Radial: Specifies the high order coefficient of radial lens distortion
Decentering: Specifies low and high order de-centering coefficients.
Enable: The enable check-boxes will enable/disable each parameter as required.

For the Cooke Optics distortion models, the following parameters are available:

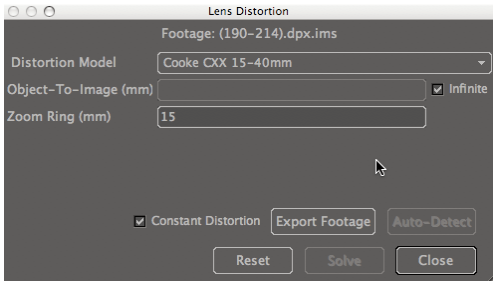

Object-To-Image (mm): This specifies the object-to-image distance for the shot (in millimeters). The 'Infinite' check-box can be used to specify an infinite object-to-image distance.
Zoom Ring (mm): Specifies the Zoom Ring setting on the camera, and is only available for Cooke CXX Zoom lens distortion models.

Constant Distortion: Specifies whether the distortion parameters are constant throughout the shot. Un-check this box if you wish to use different distortion parameters for different frames.
Export Footage: Export an undistorted version of the footage.
Auto-Detect: Automatically detect distortion lines in the image.
Reset: Reset all distortion parameters back to their default values.
Solve: Solve for the distortion parameters using manual or automatic distortion lines.
Close: Closes the lens distortion window.

Solver Controls

The Solver Controls window is displayed when solving either camera or object motion, and can be used to control the way in which the camera/object solve is performed. Further details of these options are available elsewhere (in particular, the Solving Camera Motion and Stereo Camera sections of this manual).

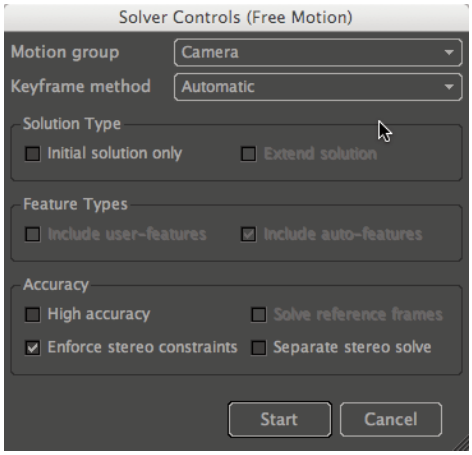

Motion group: Select the motion group to solve. This will always be Camera unless you have created additional motion groups. Additional groups cannot be solved before the main camera motion.
Keyframe method: Select Automatic or Manual placement of keyframes to create an initial solution. Manual placement of keyframes can greatly assist the solver in obtaining a good quality solution.

Initial solution only: Choosing this option will produce only the initial part of the solution between the initial keyframes.
Extend solution: Selecting this option will allow an existing solution to be extended, instead of solved from scratch.

Include user-features: Select this to include user-features in the solve.
Include auto-features: Select this to include auto-features in the solve.

High Accuracy: This enables a slower but more accurate solve mode, where additional frames will be added to the initial solution one at a time. This will often give increased accuracy but will also increase the time it takes to complete the solution.
Solve reference frames: When enabled, reference frames will be solved in addition to the main camera motion.
Enforce stereo constraints: Enabling this option will try to enforce constraints on the position of the secondary camera during the camera solve.
Separate stereo solve: This option can be used to solve motion for the secondary camera as a separate step after the primary camera has been solved, rather than having them both contribute to the solution.

Start: Start the camera/object solve.
Cancel: Closes the solver controls widow.

Statistics

The Statistics window displays tables containing numerical error values for the camera or object solve. It is opened by choosing *Camera->Statistics...* from the menu. The data can be sorted by clicking in any of the column headers. Clicking in a row of the table will either select that feature (when in Per-Feature mode), or move to the frame (when in Per-Frame mode).

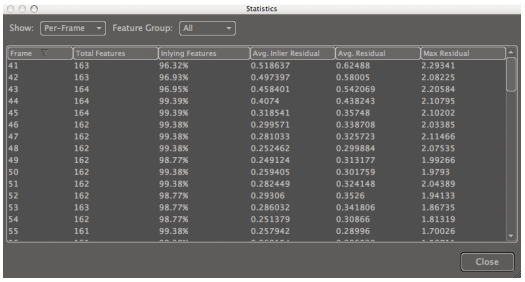

Show: Change the table to display statistics for each frame, or for each feature.
Feature Group: Displays statistics for all groups or only for a specific group.
Close: Closes the statistics window.

Per-Frame statistics:
Frame: The frame number. For stereo or motion capture cameras, the frame number will be marked as either Primary or Secondary.
Total Features: The total number of features in the frame.
Inlying Features: The percentage of features whose projection errors are below the Inlier Threshold set in the Calibration tab of the Preferences window (these are the features that are displayed in the image window as green dots).
Avg. Inlier Error: The average RMS error for this frame (measured in pixels) for all inlying features.
Avg. Error: The average RMS error for all features (inlying and outlying).
Max Error: The largest RMS error in the frame.

Per-Feature statistics:
Feature: The feature name.
Group: The group that contains the feature.
Total Frames: The total number of frames that the feature was tracked for.
Inlying Frames: The percentage of frames where the feature projection error is below the Inlier Threshold set in the Calibration tab of the Preferences window (these are the features that are displayed in the image window as green dots).
Avg. Inlier Error: The average RMS error for this feature (measured in pixels) over all inlying frames.
Avg. Error: The average RMS error over all frames (inlying and outlying).
Max Error: The largest error over all frames.

Clean Auto Features

This window is used to remove poorly solved auto-features, and is opened from the *Camera>Clean Features..* menu.

Outlier Threshold (pixels): This is the threshold used to decide whether to remove an auto-feature from a frame. By default, the threshold is set to the largest projection error for a feature. The threshold corresponds to the horizontal "Outlier Threshold" line that is drawn in the Proj-E timeline graph. Reducing this threshold will remove features from frames in which their projection error is too high.
Rejection Count (1 to 100%): This controls the removal of entire auto-features from the solution. An auto-feature will be removed entirely if the percentage of frames where the Outlier Threshold is exceeded is larger than this value. For example, if this is set to 50%, then any auto-features that exceed the Outlier Threshold in more than half of their frames will be removed entirely.

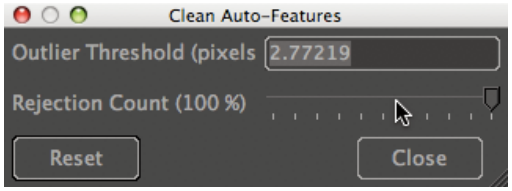

Reset: Reset the Outlier Threshold and Rejection Count to their default values.
Close: Close the Clean Features window.

Improve Solution

This window contains the controls that can be used to improve an existing camera/object solve. This is necessary after editing feature tracks or adding new tracks after the camera is solved. The window can be opened from the *Camera>Improve Solution...* menu or by clicking the Improve Solution button in the toolbar.

When trying to improve the solution, the cameras and feature points will be adjusted to try to reduce the projection error for each feature. The various parameter weights are expressed relative to the feature projection error, so specifying a Focal Length weight of 2.0 will mean that twice as much weight is given to ensuring a constant focal length compared to the feature projection error. This in turn will mean that the projection error for features may increase slightly, in order to better match the focal length.

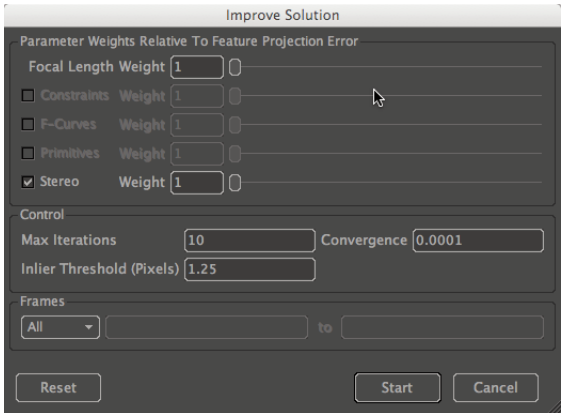

Focal Length Weight: The relative weight applied to ensure the shot has a constant focal length. This is not available for shots with a variable focal length.
Constraints: When selected, feature constraints will be accounted for when trying to improve the solution.
Constraints: The relative weight applied to any feature constraints.
F-Curves: When selected, F-Curve edits will be accounted for when trying to improve the solution.
F-Curves Weight: The relative weight applied to the any edited F-Curves.
Primitives: When selected, modelling primitive positions will be accounted for when trying to improve the solution.
Primitives Weight: The relative weight applied to any modelling primitive positions.
Stereo: When active, stereo constraints will be accounted for when improving the solution

Control:
Max Iterations: The maximum number of iterations that the improvement algorithm will run for.
Convergence: The change in error that signifies when the solution cannot be improved any more.
Inlier Threshold: This specifies the distance (in pixels) that will be used to classify a feature

point as being an inlier. Only inlying feature points will contribute to the solution improvement, so this value can be increased in order to "pull" outliers towards their track positions.

Frames:
All: Improve the solution for all frames between the in and out points.
Current: Improve the current frame only (and features that are tracked in the current frame)
From/To: Improve the solution for a subset of frames between the in and out points.
Start: Start the improvement process. When completed, if the results are not as expected you can undo the change using Edit>Undo (or by pressing Ctrl+Z).
Cancel: Closes the window without improving the solution.

Depth Map

The Depth Map window is opened from the *Camera>Depth Map...* menu, or by clicking the Depth Map icon in the toolbar. It contains the controls that are used to calculate per-pixel depth maps after camera motion has been solved.

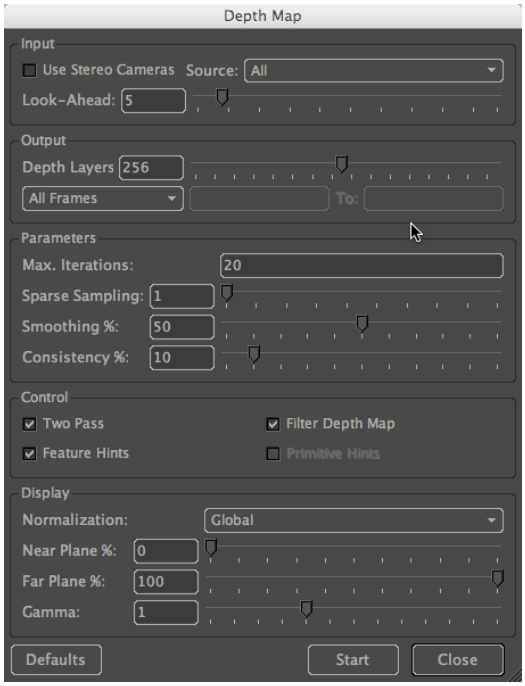

Input
Use Stereo Cameras: When stereo cameras have been tracked, this option is available and will allow an independent depth map to be created for each frame, using image data both the primary and secondary cameras.
Source: For stereo shots, when Use Stereo Cameras is not specified, this menu allows you to specify which footage will be used to generate the depth map. Setting this to All will calculate a depth map for both the primary and secondary cameras.
Look-Ahead: This specifies the separation between frames that will be used to estimate depth. PFTrack requires that there is parallax and camera translation between pairs of frames to estimate depth reliably. In the case of a shot with minimal camera motion, increasing this parameter will produce better depth estimates. Similarly, for fast moving cameras, decreasing this may also improve the quality of the depth estimates.

Output
Depth Layers: Specifies the number of depth layers in the depth map. Increasing this value will give more accurate depth maps, but will also increase computation time. If near and far camera clip planes have been specified in a 3D Viewer Window, these depth layers will be positioned between those planes.
Frame Menu: Specifies which frames will have depth maps calculated (All Frames, Current Frame or From Frame).

Parameters
Max Iterations: Specifies the maximum number of iterations used to build the overall depth map for each frame. Decreasing this value will speed up the depth map calculation, but will also decrease the accuracy of the final result.
Sparse Sampling: Specifies the density of pixels used to measure depth. Increasing this value will speed up the depth map calculation but may also give less accurate results. A value of one indicates that every pixel will be used in the calculation. A value of 5 would mean that only every 5th pixel was used.
Smoothing: Specifies the amount of smoothing within one frame. This controls attempts to

assign similar depth values to nearby regions in the image that have a similar appearance.
Consistency: Controls the amount of consistency between depth maps from different frames.

Control
Two Pass: When building depth maps for multiple frames, the two-pass algorithm will try to make the depth estimates consistent from frame-to-frame. This will give more accurate results but will also increase computation time. This option is not available when using stereo cameras.
Filter Depth Map: Apply a smoothing filter to a depth map after it has been calculated.
Feature Hints: Use inlying 3D feature positions as hints to the depth solver
Primitive Hints: Use modelling primitive positions as hints to the depth solver

Display
Normalization: Specifies how depth values are converted in a grey-map with values between 0 and 1. 'Local' normalization means the near/far depths are mapped to white/black values independently for each frame. Global normalization means the mapping is consistent between all frames, with white corresponding the the nearest distance over all frames, and black the furthest.
Near Plane: Adjusts the near clipping plane when mapping depth values to a grey-map. The range 0 to 1 represents the distance between the smallest and largest depth values.
Far Plane: Adjusts the far clipping plane when mapping depth values to a grey-map. The range 0 to 1 represents the distance between the smallest and largest depth values.
Gamma: Adjust the linearity of the grey-map. When this value is 1, depth values between the near and far plane will be mapped linearly to grey scales. Other gamma values control the mapping in exactly the same way as standard image gamma adjustments.

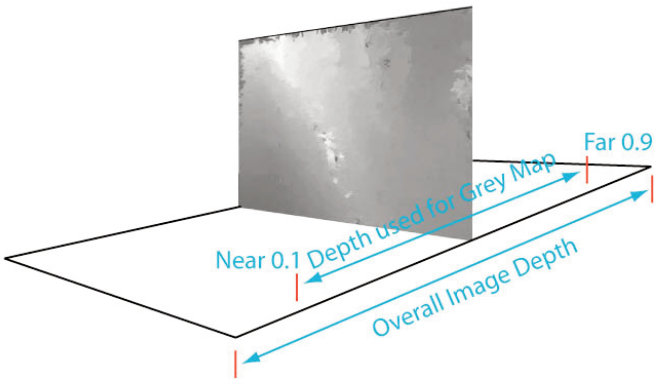

Defaults: Resets all the parameters to their default values.
Start: Starts the depth map calculation process.
Close: Closes the depth map window.

Scale Scene

This window can be used to scale the camera and feature positions so that the distance between two 3D feature points is a specific value. This can be useful if you wish to specify an exact scale for your scene, for example, the distance between features at the bottom and top of a door-frame is exactly 190cm. It also provides options to scale the scene to set a desired inter-ocular distance between stereo cameras. The window is opened using the *Camera>Scale Scene...* menu option.

The table lists the features that are available for adjusting scale. If features were selected when the scale window was opened, only these will be listed. Otherwise, all features will be listed. Pairs of features can be selected by clicking with the left mouse button. Next to the

name of each feature are its X, Y and Z coordinates. To set the distance between two features, click on them in the table with the left mouse button and enter a distance in the 'Distance between features' edit box.

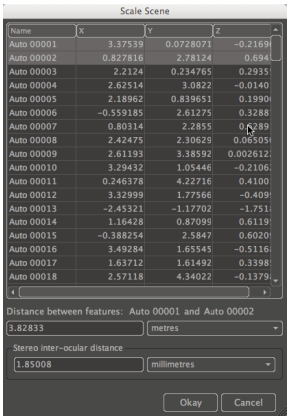

Stereo inter-ocular distance: Editing this value will scale the scene so that the average distance between the primary and secondary stereo cameras is the desired value. For example, if the main scene units are in metres, you can adjust the inter-ocular distance to be exactly 65mm, relative to those scene units.

Okay: Closes the scale window and accept any scale changes.
Cancel: Cancels any changes and close the window.

Transform Camera

This window is used to transform the entire scene relative to one camera, after the camera motion has been solved. For example, you can use this to place the camera for one frame in a specific position, rotate the camera in one individual frame, or transform one segment of the camera motion in a specific way. It can also be used to transform the secondary stereo camera into a fixed position relative to the primary camera. The existing footage playback and timeline controls can be used to change the current frame at any time, and the transformations can be undone at any time using the *Edit>Undo* menu or pressing Ctrl+Z.

The window is opened from the *Camera>Transform Camera...* menu option. It may help to have a 3D viewer window open when transforming the camera, so you can see the effect the transforms have more easily.

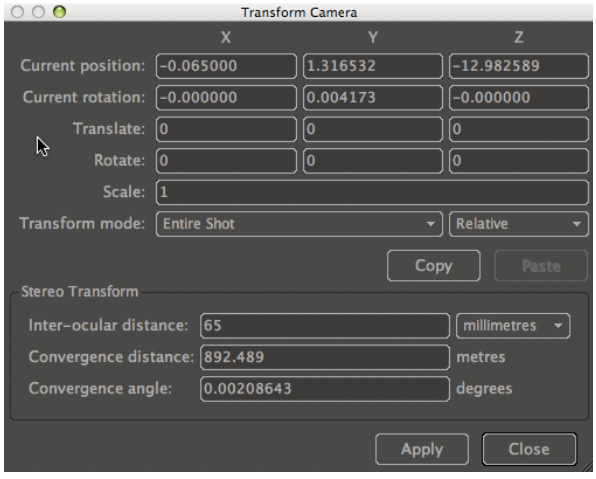

Current position: Shows the X, Y and Z camera position for the current frame.
Current rotation: Shows the camera rotation (in degrees) for the current frame in X, Y and Z.

Translate: Enter translation adjustments in X, Y and Z. These are either relative or absolute adjustments, according to the Transform mode.
Rotate: Enter relative or absolute rotation adjustments (in degrees).
Scale: Enter a relative or absolute scale adjustment.
Transform mode: Specify the type of adjustment that will be made. Entire Shot will transform the camera in all frames between the in/out point; Current Frame Only will transform the camera for the current frame only; and Current Segment Only will transform cameras in the current segment (the current segment is the segment containing the current frame).
Relative/Absolute: Changing the transform mode to Relative will apply transformations relative to the current camera. Specifying Absolute will enable absolute transformations.

Copy: Take a copy of the current camera position and rotation.
Paste: Paste a copied camera position and rotation into the current frame.

Stereo Transform

These controls can be used to adjust the secondary stereo camera position. They will not affect the position of the primary camera. Note that applying these transforms to stereo camera rigs that do not have the secondary camera placed directly next the primary may increase the number of outlying feature points.

Inter-ocular distance: Changing this value will position the secondary stereo camera at a fixed distance from the primary camera.
Convergence distance: This edit box is used to specify a fixed convergence distance for the primary and secondary cameras. For parallel cameras, this value will be infinite.
Convergence angle: This contains the stereo convergence angle (in degrees) for the primary and secondary cameras.

Apply: Applies the transform to the camera without closing the window.
Close: Closes the transform camera window.

Un-Solve Frames

This window can be used to remove certain frames from the camera solve. This can be useful to remove spikes in the camera solve by first un-solving the incorrect frame(s), editing feature tracks if necessary and then extending the existing solution into those frame(s). The window is opened from the *Camera>Un-Solve Frames* menu.

Group: Specifies the motion that will be un-solved. This will always be Camera, unless you have tracked multiple motion groups.
From frame: the lowest frame number to un-solve.
To frame: the highest frame number to un-solve
Set: Clicking a Set button will put the current frame number into either the From frame or To frame edit boxes. The current frame can be changed at any time from within the timeline.

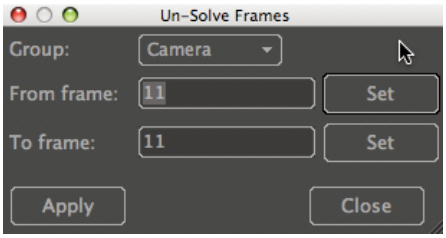

Apply: Clicking the apply button will un-solve the specified range of frames (inclusive).
Close: Close the un-solve frames window.

Optical Flow

This window contains the controls for generating optical flow. It is opened by choosing the *Optical Flow>Parameters...* menu option, or by clicking the Optical Flow toolbar button.

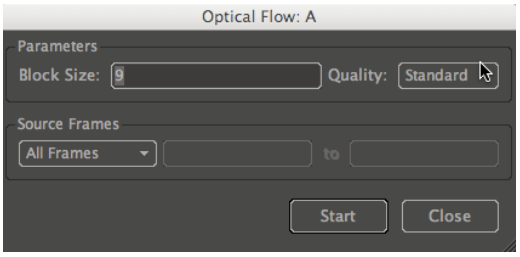

Block Size: The size of the block in pixels used to evaluate optical flow. The optical flow algorithm will try to match image data by comparing blocks of pixels in each frame. This parameter controls the size of the block that is compared. Decreasing this value may give better optical flow results in areas of complex motion, but will also increase the overall running time.
Quality: Select the quality parameter to be used when calculating optical flow. 'High Quality' mode will perform additional calculations that improve the per-pixel flow vectors after they have been calculated. This will increase running time but give a higher quality flow field.
Source Frames: Sets the range of frames to be used for the optical flow calculation. Options are 'All Frames', 'Current Frame' or 'From/To Frame'.

Calculate Flow: Start the optical flow calculation.
Close: Closes the optical flow window.

Optical Flow Edit
The optical flow edit window is used to visualize and edit an optical flow field after it has been calculated. The edit window is opened by selecting the *Optical Flow>Edit...* menu option, or by right-clicking on the Optical Flow entry in the shot overview and selecting Edit... from the popup menu.

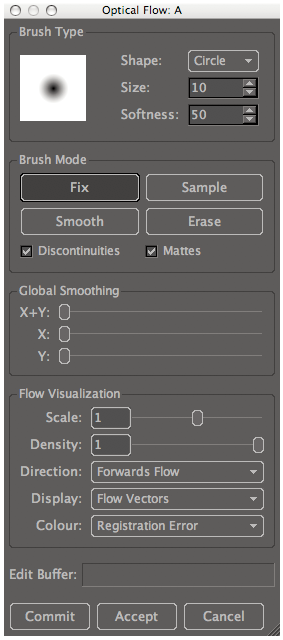

Brush Type:
These options define the shape and type of brush that is used to paint in the optical flow field. The brush shape (circle or square) can be changed, along with its size and softness. To the left of these options is a brush preview window, showing the size and shape of the current brush.

Brush Mode:
These buttons control the brush behaviour when painting into the flow field. There are four brush modes available:

Fix: Painting the flow field with this brush will adjust the flow vectors so they point in the direction of the brush stroke. Brush vectors will not be adjusted if doing so would increase the error.
Sample: Hold the shift key and use the left mouse button to sample the flow field and calculate an average flow vector. Painting without the shift key held will paste this average vector into the flow field.
Smooth: The brush will smooth the flow field in painted areas.
Erase: Paints a discontinuity into the flow field, effectively removing motion vectors for pixels that are occluded in the next/last frame.
Discontinuities: When active, the paint tools will respect discontinuities so that they remain unaffected.
Masks: When active, the paint tools will respect mask boundaries so that vectors within a mask boundary are unaffected by the smooth brush.

Global Smoothing:
These sliders controls the level of flow overall field smoothing. One slider for each of the X and Y and a third for X and Y combined.

Flow Visualization:
These controls affect the way the flow field is displayed on screen.

Scale: Slider control to set the scale of vectors displayed over the frame, from zero to maximum of two (twice as long as normal). Note that this only affect how the vectors are displayed and will not scale the flow vectors stored on disk.
Density: Slider control to set the density of vectors displayed over the frame, from zero to a maximum of one (full density). Note that this only affect how the vectors are displayed and will not adjust the density of flow vectors stored on disk.
Direction: Displays either the forwards or backwards flow field. For one pair of frames, the forward flow will be in the direction X->X+1, and the backward flow is reversed, from X+1->X.
Display: This Menu offers choices of how the flow field is displayed. 'Flow Vectors' displays the field as small arrows, 'Image Overlay' displays the field as coloured pixels according to current colour choice (see the Colour option below), and 'Morphed Frame' will display an in-

between frame, rendered at the mid-point between the current frame and the next.
Colour: Specifies the colour of the vectors or image overlay. Choices are 'None' for no colour, 'Registration Error' to display the error for each pixel (colour coded in the standard green/yellow/red scheme), and 'Direction/Magnitude' to colour pixels according to the direction and/or distance they move.
Edit Buffer: This is a progress bar that shows how full the 100Mb RAM-based edit buffer is. Each brush stroke will be stored in RAM until the edits are accepted and the flow data written back to disk. Once the edit buffer is full, PFTrack will start to commit edits to the flow field to disk to free up more space.

Commit: Commit edits for the current frame to disk.
Accept: Commit edits for all frames back to disk and close the edit window.
Cancel: Cancel all edits and close the edit window.

Script Editor

The script editor window can be used to run pre-defined scripts, or send scripts to the Batch Manager for processing elsewhere. It is displayed after opening a script using the *Batch>Open Script...* menu option or by selecting *Batch>Edit Script...* after a script is opened. At the top of the script window is a text editor where the script can be viewed and edited. Below the editor is an information window containing useful messages about the script.

The buttons below the information window are used to access and run scripts:

Open: Open an existing script for editing.
Save: Save the current script to disk.
Verify: Verify the script contents to make sure there are no syntactic errors. Errors are highlighted in red in the editor window and further information displayed below.
Execute: Execute the current script.
Batch: Send the current script for processing by the Batch Manager.
Clear: Clear the current script.
Close: Close the script edit window.

Batch Manager

The Batch Manager window is used to control how different tasks executed by PFTrack are sent out for processing on other machines. More details of the Batch Manager and how to use it are given later in this manual. The Batch Manager window is opened by selecting *Batch>Launch Batch Manager...* from the menu.

The text line at the top of the window shows the connection state between PFTrack and the Batch Manager. The two top window panes show a summary of the batch processing network. At the top-left is the list of clusters, machines and their loads, and at the top-right is a summary of the status of each submitted job. Selecting an item from either of these lists by clicking with the left mouse button displays further information about that item in the windows at the bottom of the window. Right-clicking on a job in the top-right window displays a popup menu allowing the job to be killed, paused/resumed or re-submitted.

The buttons at the bottom allows you to filter what information is displayed. The machine button toggles display of all machines, those you can submit jobs to, and those which have spare capacity. The job button toggles between displaying all jobs, or jobs submitted by the current user.

Footage Format

The Footage Format window can be used to change the characteristics of a piece of footage, such as the frame-rate or de-interlacing mode. The window is opened by right-clicking on a footage name in the shot overview and selecting Edit Format from the popup menu.

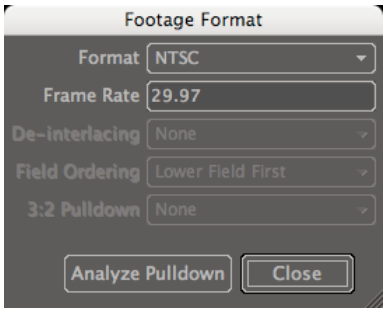

Format: Sets the format of the footage to be NTSC, PAL or Film. This will adjust the frame rate accordingly.
Frame Rate: Specify the frame rate of the footage. If you are importing Cooke S4i dataLink files, it is important that the frame-rate is set correctly in order to allow timecode calculations to be performed correctly.
De-interlacing: Specify the method of de-interlacing the footage. Choices are: 'None' to ignore interlacing, 'Field Averaging' to average the two fields together to create the frame, 'Interpolate Lower Field' to keep the upper field and interpolate the lower field, 'Interpolate Upper Field' to keep the lower field and interpolate the upper field, and 'Field Separation' to treat each field as a separate frame. Choosing 'Field Separation' will double the length of the shot and also double the frame rate.
Field Ordering: Specifies the field order. Active only when a de-interlace mode is selected. Options are 'Lower Field First' and 'Upper Field First'.
3:2 Pulldown: Specifies the method for removing 3:2 pulldown and revert back to 24fps footage (Only available for NTSC sequences). Options are defined by 'W' for a whole frame, and 'S' for a split frame.
Analyze Pulldown: Automatically analyse the NTSC footage to determine the best method for removing 3:2 pulldown.
Close: closes the footage format window.

Proxy Management

The proxy management window is used to control the generation of proxy images for faster feature-tracking. It is opened by right-clicking on a footage name in the shot overview and selecting *Proxy>Manage Proxies...* from the popup menu. By default, all proxies are scaled automatically from the original clip. However, if you already have proxy representations stored on disk you can specify their locations here. The proxy table displays the following information:

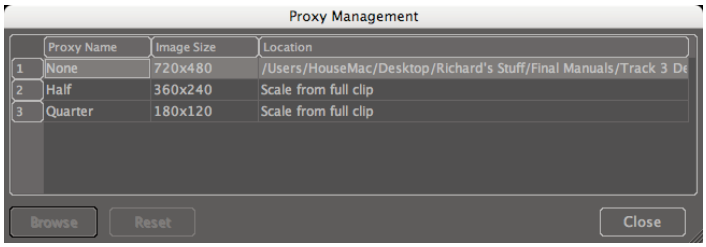

Proxy Name: The name of the proxy ('None', 'Half' or 'Quarter').
Image Size: The size of the proxy image.
Location: The location of the proxy footage. When no disk-based proxy is used, this will display 'Scale from full clip'.

When a proxy is selected by clicking the left mouse button in the proxy table, the following buttons are available*:

Browse: Browse to the location of the disk-based proxy sequence.
Reset: Reset the proxy to 'Scale from full clip'.

* Note that you cannot change the full resolution clip from this window. Also, you cannot change the currently active proxy.

Close: Closes the proxy management window.

Cooke S4i Data Sync

This window is used to synchronise data from a Cooke S4i data file with the timecode present in the footage. This is necessary to ensure that each frame of the footage contains the correct data from the Cooke S4i dataLink device. The window will be displayed when a Cooke S4i data file is imported by right-clicking on the footage name in the shot overview and selecting *Import Cooke S4i Data File...* from the popup menu.

Once the data file is imported, PFTrack will attempt to match the timecode in the footage with timecodes stored in the data file. If this is not possible, a warning message will be displayed.

The top part of the window contains information about the timecode that has been read for the current frame. Currently, PFTrack only supports reading the timecode from DPX footage. For all other types of footage, the timecode will be calculated automatically according to the frame number and the frame-rate of the footage. Before importing a Cooke data file, ensure that the frame-rate for the footage is correct, because this will affect the results of the timecode calculations.

The lower part of the window shows matching Cooke S4i data, if any is found. When no data can be matched, this will display 'No data found'. In this case, it becomes necessary to manually edit the timecode associated with one frame of the footage. This is achieved by clicking the 'Edit' button and entering the correct timecode for the current frame. Once the timecode is entered, PFTrack will try to search through the Cooke data file again to see if it can be matched to any data.

Once timecodes have been matched, the lower part of the window will display the information for the current frame that has been gathered from the Cooke data file. This data includes:

The camera name
The camera serial number
The lens name
The lens serial number
The reel number
The timestamp for the data chunk
The focus values (including near and far focus)
The Zoom and Iris values
The horizontal field of view, in degrees
Entrance pupil position
Normalized zoom value
Camera status

Note: not all of this data may be available.

When importing XML Camera data, the information displayed will depend on the contents of the XML file, but will typically contain:

The 'Info' string stored in the file;
The film back dimensions;
The pixel aspect ratio;
The current frame number;
The focal length of the camera.

Clicking the Okay button will store the synchronized data with the footage and ensure that this information is used during the camera solve. In particular, the horizontal field of view from the Cooke S4i data file, or the focal length values from the XML Camera data file will be used to set the camera's focal length. This can be done for constant or variable focal length cameras.

The Cancel button will discard the data and close the window.

The F-Curve editor and Dope-Sheet can be used to visualize and edit the stored focal lengths. Open either editor and then click on the Focal Length item in the shot overview.

Synchronize Footage

The synchronize footage window is displayed when replacing a piece of footage with one containing more frames. To replace a piece of footage, right-click on the footage name in the shot overview and select *Replace...* from the popup menu. After browsing for the new footage on disk, the synchronize footage window will open if the new footage needs to be synchronized with the old. In the case of both pieces of footage having the same number of frames, this synchronization is not necessary.

The window displays both the old and new footage, side-by-side. The Sync Frame edit boxes below should be used to identify the same frame in each piece of footage. Changing the frame numbers in these edit boxes will alter the image displayed in the window.

Sync to start: Synchronize the footage assuming that the first frame of each footage is the same.
Sync to end: Synchronize assuming that the last frame of each is the same.
Okay: Close the window and synchronize the pieces of footage.

Footage Manipulation

These footage manipulation window can be used to affect the way the footage is processed before tracking and optical flow calculations. Note that these manipulations are non-destructive and will not affect the image data stored on disk. To open the footage manipulation window, right-click on the footage name in the shot overview and select *Manipulate...* from the popup menu.

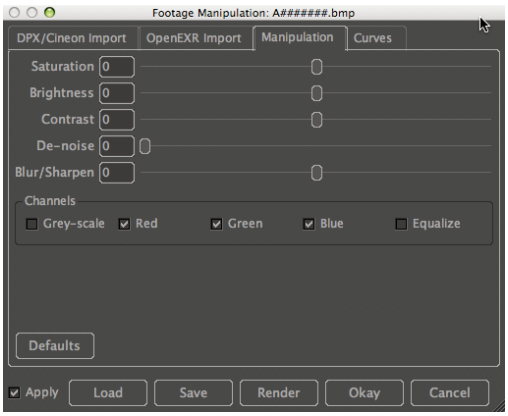

The following buttons are available along the bottom of the footage manipulation window:

Apply: This can be used to quickly enable/disable the footage manipulation.
Load: Load a previously saved manipulation file from disk.
Save: Saves the current manipulation state in a simple ASCII format.
Render: Renders the manipulated footage to disk.
Okay: Apply the current settings to the active image and close the window.
Cancel: Cancel all current adjustments and close the window.

DPX/Cineon Import:

These controls affect the way that Cineon and DPX images are imported into PFTrack (note, these options will have no effect on footage that is not Cineon or DPX).

Ref Black: Adjust the reference black level.
Ref White: Adjust the reference white level.
Soft Clip: Adjust the soft clip parameter.
N-Gamma: Adjust the negative gamma.
D-Gamma: Adjust the display gamma.
Use LUT: Change the lookup table (LUT) used to import DPX/Cineon images. Options are Standard for the standard conversion, None for no conversion, and User Defined.

Load LUT: Loads a previously saved DPX/Cineon LUT.
Save LUT: Saves the current DPX/Cineon LUT in a simple ASCII text format.
Remove LUT: Remove the currently active LUT from the list.
Defaults: Reset all DPX/Cineon import parameters to their default values.

OpenEXR Import:

These controls affect the way that OpenEXR images are imported into PFTrack (note, these options will have no effect on footage that is not OpenEXR).

Black Point: Adjust the black point. This is the floating-point value that will be mapped to a pixel value of 0 inside PFTrack. Floating-point values below this will be clamped.
White Point: Adjust the white point. This is the floating-point value that will be mapped to a pixel value of 1 inside PFTrack. Floating-point values above this will be clamped.
Gamma: Adjust the linearity of the translation between floating-point values and pixel values.
Defaults: Resets OpenEXR controls back to their default values.

Manipulation:
These controls can be used to adjust various image processing parameters.

Saturation: Adjust the saturation of the footage.
Brightness: Adjust the brightness of the footage.
Contrast: Adjust the contrast of the footage.
De-noise: Remove noise from the footage. This will attempt to remove noise without blurring object edges.
Blur/Sharpen: Blur or sharpen the footage.
Greyscale: Toggle between a colour and grey-scale image.
Red: Toggle the red channel on/off
Green: Toggle the green channel on/off.
Blue: Toggle the blue channel on/off.
Equalise: Performs a histogram equalization operation to ensure the footage contains an even spread of light and dark pixels.
Defaults: Resets manipulation controls back to their default values.

Curves:
A curve editor is provided allowing you to manipulate footage curves using linear, smooth or freehand splines. Adjustments can be made to individual channels or the overall footage brightness (Value). To edit a spline, click and drag with the left mouse button in the spline editor. Spline points can be removed by hovering over them with the mouse and pressing the Delete key.

Channel: Specify the channel to edit. Options are 'Red', 'Green', 'Blue' and 'Value'.
Curve: Toggle between 'Linear', 'Smooth' or 'Freehand' curve styles.

Export

The scene export window is where camera, object, feature and geometry data can be exported from PFTrack. To open the scene export window, right-click on an export type in the Shot Overview (the 'Tracking Exports' for 2D tracking exports, 'Camera Exports' for 3D camera/object exports, or 'PFPlugin Exports' items) and select New... from the popup menu to create a new export. Alternatively, an existing export can be edited by right-clicking on it's name and selecting *Edit...*

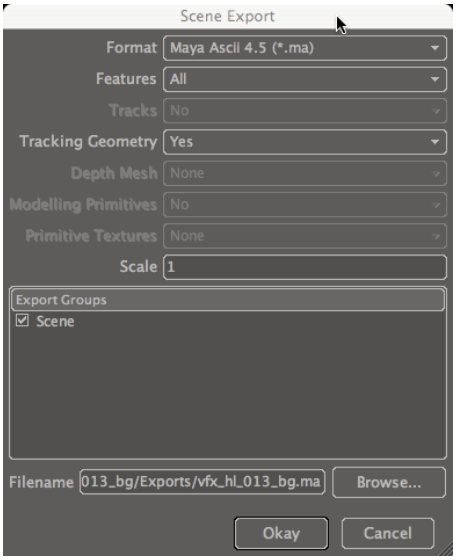

Format: The format that the data will be saved in. To change the format, select an option from the drop-down menu *.
Features: Control which features will be exported. Options are 'All' for all features, 'None' for no features, or 'Tagged' to export only the tagged features. Features can be tagged for export by right-clicking on them in either the shot overview or image window and selecting Tag For Export from the popup menu.
Tracks: Control whether 2D feature tracks will be included in the export. This option is only available for export formats that support 3D and 2D data in the same file (e.g. XML export).
Tracking Geometry: Specify whether tracking geometry will be exported
Depth Mesh: Specify whether a geometry mesh will be build from depth map data *.
Modelling Primitives: Specify whether modelling primitives will be exported *.
Primitive Textures: When exporting textured modelling primitives, this allows the format of the texture maps to be specified *.
Scale: Specify an overall scale to apply to the scene before exporting. A scale of 10.0 will expand the scene by a factor of ten. This may be useful if you need to work in a third party application at a different scale to the one used in PFTrack.

Export Groups: Specify which motion groups will be exported.
Filename: The filename for the export.
Browse: Display a file dialog to allow a new export location to be chosen.
Save: Export the chosen data and close the window.
Cancel: Close the window without exporting any data.

* Note that not third party applications support all of the data that can be exported from PFTrack. As of PFTrack 5.0, tracking geometry can be exported to 3ds Max, FBX, Flame, Lightwave, Maya, Softimage XSI and as Human-Readable and XML text. Depth meshes, modelling primitives and motion capture data can be exported to 3ds Max, FBX, Lightwave, Maya and Softimage XSI and XML text.

Preferences

The preferences window is where user preferences are specified for different aspects of the PFTrack interface. It is opened by selecting File>Preferences... from the menu (Note that on Mac OSX, the Preferences menu item is located in the PFTrack menu, instead of the File menu).

The preference options are organized into groups, with each group presented as a different tab. To change tabs, click on the tab name at the top of the preferences window. The 'Defaults' button will reset the preferences for the current tab to their default values. The 'Okay' button will accept any changes and close the window. The 'Cancel' button will reject changes and close the window.

General Tab:
Remember folders: With this selected PFTrack will remember which folders have been used for loading and saving data.
Load last project: With this selected PFTrack will load the last project that was open when the application is started.
Load Locked Projects: With this selected, PFTrack will warn if another user is using the project you are trying to open, but will still allow you to delete the lock file and open the project if you wish.
Zero-based frame numbers: With this selected, PFTrack will number all frames in your footage from zero rather than use the frame number from the image file.
Name shot from footage: When selected, new shots will be automatically renamed to match any imported footage.
No spaces in filenames: When selected, shot folder names will have spaces replaced by under-scores.

Autosave: With this selected PFTrack will automatically save a copy of the project at user-specified intervals.
Save after track/solve: Automatically saves the project after a tracking and solving operations are completed.
Autosave Interval: This sets the time in minutes between auto-saves.

Auto-feature display type: Select how auto-features are displayed on screen. Options are 'Box', 'Cross' and 'Circle'.

Automatic version check: When selected, PFTrack will check to see if a new version is available when the application is launched.
Check now: Clicking this button will perform the version check immediately.
Hide new progress windows: When selected, new progress dialog windows will be hidden, and the progress will be displayed as a bar above the timeline. Hidden progress windows can be re-shown by clicking on the icon that appears next to the "log window" button at the bottom-left of the timeline.

Interface Tab:
Mouse forward is zoom: This allows you to select if forward mouse movement corresponds to zooming in or out of either the image window, 3D viewer window or timeline.
Spline snap distance: This specifies the distance (in pixels) that you have to move the handles of a mask boundary point before they are snapped together to form a linear point.
Levels of Undo: Sets the number of undo levels available in the application.
Timeline scrub distance: This specifies the distance (in pixels) that the mouse must be moved before the current frame is changed when holding the timeline scrub keyboard shortcut (backslash '\' by default).
Graph lock sensitivity: This is the sensitivity level that determines whether mouse motion in the F-Curve graph is locked to either the horizontal or vertical directions.
Tablet sensitivity: This controls how tablet pen pressure is used to adjust the softness of the optical flow brush.

OS native file dialog: With this selected PFTrack will use the operating system's native file dialog for imports and exports. This enables operating system specific items such as 'Network Neighbours', 'My Recent Documents' etc. but removes the preview window and image thumbnails from the footage import window.
Guess frame boundary: When checked, PFTrack will automatically estimate the image boundary by examining the first frame of footage. The image boundary can be edited by dragging the border lines whilst the View>Frame Boundary menu option is active.

Store image thumbnails: With this selected PFTrack will generate and display thumbnails in the footage import window, and store them on disk. It can take time to process the files before any thumbnails are displayed, especially when footage is located on a network, so switching this option off may increase the speed of browsing through your filesystem.
Check image sequences: PFTrack will automatically check the integrity of image sequences when imported. This involves making sure that all frames can be read from disk and are the same width and height.
Thumbnail folder: The location where image thumbnails are stored. When the folder edit box is empty, thumbnails will be stored in a '.thumbnails' folder in the same location as the footage.
Browse: Clicking this button will open a file dialog to let you choose a new folder for thumbnail storage.
Fullscreen: Toggle the interface between full screen and windowed modes.
Maximized: Maximizes the application window to fit the full screen resolution, without removing the window border.
Viewers stay on top: Control whether 3D Viewer windows always remain on top of image windows. On certain operating systems, toggling this value will only affect newly created viewer windows.

Performance Tab:
Use memory cache: When selected, PFTrack will store frames of your footage in RAM to increase the speed of playback.
Cache on import: Load the footage into the cache as soon as it is imported, rather than on the fly.
Cache size (Mb): Set the amount of RAM in megabytes (Mb) that is used to store cached frames. The more you can allocate the better, but be careful to leave some RAM free for other operations such as feature tracking and the camera solver. As a rule, try to leave at least 500Mb free for the system.

Compression Quality: This controls the amount of compression that is used to store footage previews in RAM. Compression will significantly reduce the amount of RAM that is needed to store footage previews, but will also increase the time it takes to display each frame.
Resolution: This controls the resolution of footage previews. Decreasing resolution will reduce both the amount of RAM needed to store the preview, as well as the time it takes to display.
Reserved RAM (Mb): This specifies the amount of RAM that is reserved for storing footage previews.
Clear Previews: Clicking this button will clear all footage previews from RAM.
Prefetch during playback: When selected, multi-threading will be used where available to accelerate the playback of footage previews.
Store on disk: Selecting this option will store the footage previews in the Shot folder on disk, allowing them to be re-loaded into RAM when the project is opened again.

Multi-Threading: Enabling this option will mean that dual CPUs are used to accelerate auto-feature tracking. This option is not available if you only have one CPU on your system.

Display Tab:
Always texture mapped: When selected, PFTrack will use the texture mapping facilities of your graphics card to display footage in the image window. Depending on your graphics card, enabling this option could speed up or slow down playback rates.
Maximum texture size: This specifies the largest texture map that the system will use for image display and modelling primitive textures.
Frustum size: This controls the size of the camera frustum that is displayed in 3D perspective and orthographic viewer windows.

Image Transparency: This slider controls the transparency of the frustum image displayed in 3D perspective and orthographic viewer windows. The frustum image can be turned on/off using the *View>Frustum Image* menu option.
Antialias points and lines: When selected, PFTrack will anti-alias points and lines in image and viewer windows. Disabling this option is required on some older graphics cards to ensure correct functionality.
Use hardware off-screen buffer: When selected, PFTrack will use graphics hardware to accelerate the rendering of off-screen buffer that are used for certain tasks in PFTrack such as geometry tracking and texture extraction. Hardware off-screen are not available on all platforms.

Coordinate Axis Size: This controls the size of the coordinate axis lines that are drawn in the viewer windows.
Ground Plane Size: This value is used to increase or decrease the size of the ground-plane.

Batch Manager Tab:
Batch manager: This is the machine name where PFTrack will find the batch manager. It can be a fully qualified name, e.g. frog.pixelfarm.co.uk, or the machine's IP address, e.g. 192.0.0.1, or localhost if the batch manager is running on the same machine as PFTrack.
Username: This parameter specifies a login username for the Batch Manager.
Password: The password for the above user.
Cluster: Tells the Batch Manager which machine cluster your shots should be sent to.
Program Alias: This is an advanced option that allows you to specify which version of PFTrack should be used to process the shot if multiple versions are installed.
Required OS: This is the Operating System your shot needs to be processed on. This can be used to restrict processing to certain platforms. For instance, there is no point sending a shot that has Quicktime footage for processing on a Linux machine because PFTrack does not support reading of Quicktime files on Linux.
Data access method: This specifies how a remote batch processing machine can access the data required to process the job (files such as the footage, any masks etc., and the shot file itself). Options are 'All data on shared drive' (the default) which assumes that the batch processing machine has access to the file system where all your shot data is stored – typically by having the data stored on a shared network drive. 'Footage on shared drive' means that the batch processing machine has access to the file system where the footage is located, but all other data comprising the shot will be transmitted to the batch processing network. Finally, 'No data on shared drive' means that all footage and shot data will be transferred to the batch processing network.

File System Translation Tab:
This allows you to view the conversion methods to translate file names between one operating system and another, for example: changing *'C:\my_project_for_windows' to '/mnt/my_project_for_Linux'*. See the section on batch processing for a complete description of how to set these translation options.

Calibration Tab:
Use Z-up: With this selected PFTrack will use Z, rather than Y, as the vertical (up) direction when displaying information or data in the interface.
Inlier Threshold (in pixels): The maximum allowable error (in pixels) before a feature is considered outlying (i.e. a bad feature track).
Estimated percentage outliers: This is an estimate of the number of bad feature tracks in a shot, and is used to help the camera solver.
New user-features default to: This drop-down menu controls the default constraint setting for new user features. The menu can be used to make user features default to either 'Hard Constraint' or 'Soft Constraint'.

Exports Tab:
These controls allows you to set up default export options that will be included automatically in each shot. For example, if you always want to export camera or tracking data to a particular 3D package, you can create a default export in this window, instead of creating one for each shot you track.
Add: Open the export window to specify an export format and location.
Edit: Open the export window to change an existing export format or location.
Remove: Remove a selected default export.

Keys Tab:
This is a list of keyboard short-cuts for all buttons and menu options in PFTrack. The full list of default values is available at the end of this manual. Double click with the left mouse button on a shortcut in the Accelerator column to enter a new value. Modifier keys need to be spelt out character-by-character. For example, Ctrl+A should be entered literally as 'C' 't' 'r' 'l' '+' 'A' rather than trying to simulate the key press. Cursor directions are also entered literally ('D' 'o' 'w' 'n') as are the escape key ('Esc') and the delete/backspace key ('Del' and 'Backspace').

Note: you can redefine the delete key to 'Backspace' if necessary, for correct operation with some notebooks or laptops.

Colours Tab:
This contains an editable list of all interface colours that can be changed by the user. Double click on a colour to display a colour chooser window.

Popup menus

This section contains explanations of the many menu options that are available by right-clicking on an item in an image window or the shot overview. For example, right-clicking on an auto-feature will display the auto-feature popup menu. To get online help for a specific menu, open the menu with a right-click and then hover the mouse over each item and press Shift+F1 to activate the popup help system. Right-click outside the menu to close it afterwards if you do not want to activate any of the menu options.

Footage

Right click on the footage name in the Shot Overview to display a popup menu with the following options:

Edit Format...: Opens the footage format window.
Proxy: Allows switching to a proxy image resolution. Tracking a lower resolution proxy can be significantly faster than tracking the full resolution footage, and the data can be automatically scaled back up to full resolution once tracking has completed. This facility allows you to quickly perform a test track and camera solve at proxy resolution to give you a feel for the results before processing at full resolution. The popup menu also contains an option to open the Proxy Manager window.

Import Lens File...: Loads a lens file from another shot. Lens files are stored in the Lens folder inside the shot.
Import Cooke S4i Data File...: Load a Cooke S4i data file containing per-frame lens information. This will display the Cooke S4i Data Sync window.
Import XML Camera File...: Load an XML Camera data file containing per-frame focal length

information. This will display the XML Camera Data Sync window. More information about the format of the XML Camera data file is given in the Footage And Camera Options section of this manual.
Replace...: Replaces the footage with another sequence that has at least the same number of frames. Typically, this is used to replace specially manipulated tracking plates with the original footage, or to bring in an extended image sequence for further tracking. If the replacement sequence has more frames, the Synchronize Footage window will appear.
Delete: Deletes the footage.
Render Review Movie...: Allow the rendering of a movies showing solved features.
Render Anaglyph Movie...: Render a red/cyan stereo anaglyph showing the camera solve and feature points/3D test objects etc.. Preferences are available to change the left/right eye anaglyph colours.
Manipulate...: Opens the Footage Manipulation window.
Properties...: Opens the properties window which displays useful information about the footage.

Lens Distortion

Right-click on the 'Distortion' entry in the Shot Overview to display a popup menu with the following options:

Auto-Detect: Automatically detect distortion edges in the current frame.
Clear Auto Lines: Remove all automatic distortion lines in the current frame.
Clear All Auto Lines: Remove all automatic distortion lines in the current shot.

Reset: Reset the distortion parameters to their default values (no distortion).
Edit Parameters...: Display the Lens Distortion window. This is the same as choosing **Camera>Lens Distortion...** from the main menu bar.

When distortion lines have been drawn manually, they can be de-activated by right-clicking on the line name in the shot overview or image window and un-selecting the Active menu option.

Mask

Right-click with the mouse on a mask in the Shot Overview, or inside a mask in an image window to display a popup menu with the following options:

Insert Point: Insert another mask boundary point at the current mouse position.
Attach To: Attach the mask to a user feature. The mask will be translated in the horizontal and vertical directions to match the feature track.

Boolean Operations: This menu controls the way in which masks interact with each other. The order that masks are displayed is the same as their order in the Masks entry of the shot overview window (the first mask is furthest away and displayed first, and the last is nearest to the camera and displayed last). Each mask interacts with others according to the Boolean Operation setting, which can by 'AND', 'OR', 'XOR' and 'DIFF'. The AND operation will only display pixels that are covered by this mask and by previous ones. The OR operation will display pixels that are covered by either this mask or a previous one (this is the default). The XOR operation displays pixels that do not overlap another mask, and DIFF will subtract this mask from previous ones.

Ordering>Move Away: Adjust the depth ordering to move the mask one step away from the camera. Boolean operations are affected by the depth ordering of masks.
Ordering>Move Closer: Adjust the depth ordering to move the mask one step towards the camera.

Ordering>Send to back: Move this mask so it is furthest from the camera.
Ordering>Bring to front: Move this mask so it is nearest to the camera.
Flow Guide>No Guide: The mask will not influence the optical flow calculation (the default).
Flow Guide>Motion Boundary: The mask will be used to specify the boundary between a foreground and background object.
Flow Guide>Ignore Pixels: The pixels covered by the mask will be ignored by the optical flow calculation.

Previous Keyframe: Move to the previous keyframe.
Next keyframe: Moves to the next keyframe.
Remove keyframe: Remove the keyframes for the mask from the current frame.

Ignore: Ignore the mask, effectively removing it from the image. This option is only available when further Boolean operations are not used.
Hide/Show: Hide or show this mask in the image window.
Lock/Unlock: Disables/Enables modifications to the mask.
Invert: Invert the mask so areas outside the boundary become masked, instead of areas inside the boundary.
Delete: Deletes this mask and all its keys. You can also delete a mask by clicking on the mask name in the Shot Overview and pressing the Delete key.

Outside Frame Range: This menu affects the behaviour of an image-based mask when the image sequence contains fewer frames than the shot. Options are Repeat to repeat the image sequence throughout the shot, Clamp to clamp to the first/last frame, and Filled/Empty to use a full/empty mask where no image is available.

Mask Boundary Point
Right-clicking with the mouse whilst the cursor is hovering over a mask boundary point will display a popup menu with the following options:

Linear: Use a linear spline for the mask boundary around this point. When unchecked, Bezier spline handles will be displayed.
Continuous: When Bezier spline handles are displayed, this option controls whether the boundary is continuous at this point. Non-continuous points can have their in and out Bezier handles adjusted independently.

Edit Keyframes: Opens the Dope Sheet to show the frames where this boundary point is keyframed.
Previous Keyframe: Moves to the last frame containing keyed boundary point.
Next Keyframe: Moves to the next frame containing keyed boundary point.
Remove Keyframe: Remove the keyframe for this boundary point in the current frame.
Delete Point: Removes the selected point from the mask. You can also remove a boundary point by hovering the mouse over it and pressing the Delete key.

Group

Right-click on a group name in the Shot Overview to display a popup menu with the following options:

Remove: Remove this group, but place all features into the main Scene group. This option is not available for the main Scene group.
Delete: Delete this group and all features it contains.
Hide Other/Show All Groups: Hide all other groups apart from this one, or show all groups.
Hide/Show: Hide/show this group.

Edit: Display the group window.
Rename: Rename this group in the Shot Overview. This can also be achieved by double-clicking on the group name.
Solve/Unsolve/Extend Motion: If the group represents an independent object motion, this option is used to solve/unsolved that motion, or extend the motion solve into additional frames.

User-feature

Right-click with the mouse on a user-feature in the Shot Overview or an Image/Viewer window to display the following options:

Tracking Guide: Toggles the user-feature as a guide for auto-feature tracking. Tracking guides will be used to help identify the correct camera motion when tracking auto-features.
Edit Tracking Parameters...: Opens the User-Feature tracking parameters window to change the tracking parameters for this feature.
Hard Constraint: Specify whether the user feature is used as a hard or soft constraint during the camera solve. A hard constraint assumes that the feature track is accurate everywhere, meaning that the camera solve will try to lock onto the track in all frames. The default hard/soft state of a user-feature can be set in the Calibration tab of the Preferences window.
Overwrite Track: When selected, tracking a feature forwards and backwards over the same range of frames will overwrite the track each time. When unchecked, the track will be averaged with the previous track.

Remove From Frame: Removes the feature from this frame.
Remove From Frame+Forwards: Removes the feature from the current frame and all higher frame numbers.
Remove From Frame+Backwards: Removes the feature from the current frame and all lower frame numbers.

Rename: Allows the renaming of the selected feature in the Shot Overview. This can also be achieved by double-clicking on a user-feature to display the name edit box.
Ignore During Camera Solve: Force PFTrack to ignore this feature when solving for camera or object motion.
Lock/Unlock: Disables/Enables modifications to the user-feature track.

Solve/Unsolve: If you create and track a user-feature after camera motion has been solved, you can use this option to solve for the 3D feature position in space. This allows you to add extra features into an existing scene without re-solving the entire camera path and other feature positions.
Set As Origin: Sets the origin of the coordinate system to this feature position.
Set Scale: When multiple features are selected, this provides a shortcut to the Scale Scene window.
Tag for Export: Mark this feature as one of the ones that will be exported when the 'Tagged Features' export option is chosen in the export window.

New Test Object...: Open the 3D test object browser to create a test object at this feature position.
Snap Selected Object: Snaps a selected test object to this feature position.

Survey...: Open the survey data window with this feature.
Groups...: Open the groups window with this feature.
Constraints...: Open the constraints window with this feature.
Show F-Curves: Displays the X/Y 2D tracking curves for the feature in the F-Curve editor.

Auto-Feature

Right-click with the mouse on an auto-feature in the Shot Overview or an Image/Viewer window to display the following options:

Remove From Frame: Removes the feature from this frame.
Remove From Frame+Forwards: Removes the feature from the current frame and all higher frame numbers.
Remove From Frame+Backwards: Removes the feature from the current frame and all lower frame numbers.

Convert To User-Feature: Convert this auto-feature into a user-feature, maintaining the existing feature track.

Merge Tracks: When multiple auto-features are selected, they can be merged into a single auto-feature using this option. This is useful for connecting separate feature tracks that are partially occluded by a moving object.
Split Track Before Frame: Splits the auto-feature into two. The feature track is split at the point immediately before the current frame, and an additional auto-feature is created containing the remaining portion of the track.
Split Track After Frame: Splits the auto-feature into two. The feature track is split at the point immediately after the current frame, and an additional auto-feature is created containing the remaining portion of the track.

Rename: Allows an auto-feature to be renamed in the Shot Overview. This can also be achieved by double-clicking on an auto-feature to display the name edit box.
Hard Constraint: Sets the feature track to 'Hard Constraint', which indicates that the feature track is assumed to be accurate in all frames. Hard constraint features can be used to lock the camera solve down onto specific features that are found to drift from their feature tracks.
Ignore During Camera Solve: forces PFTrack to ignore the feature when solving for camera or object motion.

Set as Origin: Sets the origin of the coordinate system to coincide with the 3D feature location.
Set Scale...: When multiple features are selected, this provides a shortcut to the Scale Scene window.
Tag for Export: Tag this feature so it is exported when the 'Tagged Features' export option is chosen in the export window.

New Test Object...: Open the 3D test object browser to create a test object at this feature position.
Snap Selected Object: Snaps a selected test object to this feature position.

Survey...: Open the survey data window with this feature.
Groups...: Open the groups window with this feature.
Constraints...: Open the feature constraints window with this feature.

Show F-Curves: Displays the X/Y 2D tracking curves for the feature in the F-Curve editor.

Constraint

Right-click on a constraint name in the Shot Overview to display a popup menu to display the Constraints window with the constraint highlighted.

Modelling Primitives

Right-click on a modelling primitive to display a popup menu with the following options:

Delete: Delete the modelling primitive from the shot.
Extract Textures: Extract textures for the primitive using the current settings in the Modelling menu.
Clear Textures: Clear all textures from the primitive.
Remove All Pins: Remove image pins from all vertices of the modelling primitive.

Optical Flow

Right-click on the Optical Flow entry in the Shot Overview to display a popup menu with the following options (Note: this entry is only visible after optical flow has been calculated).

Edit: Display the optical flow edit window.
Delete: Delete the optical flow data files from disk.
Export: Export the optical flow data as a floating point or RGB image sequence. Optical flow data can also be exported using the OpenEXR file format.

Depth Map

Right-click on a depth map item in the Shot Overview to display a popup menu with the following options. Depth map items are named after the footage they were generated from.

Delete All: Delete all the depth map data files from disk.
Delete Frame: Delete the depth map data file for the current frame.
Export Image Sequence: Export the depth map as an image sequence. Depth maps can be exported either as grey-scale images, or in the OpenEXR image format.
Export Point Cloud: Export the depth map as an ASCII file containing X,Y,Z,Depth point information.

Exports

Right-click on either the 'Tracking Exports', 'Camera Exports' or 'PFPlugin Exports' entry in the Shot Overview, or on an existing export name to display a popup menu with the following options:

Export All/Export: Save the export (or all the exports) to disk.
New: Add a new export. This will display the Export window.
Edit: Edit the export. This will display the Export window.
Remove: Remove the export from the list.

Toolbar

Right-click in an empty part of the toolbar to display a popup menu. This menu allows you to temporarily open or close different toolbar menus, or to automatically re-align them into the toolbar.

Part 2 - Camera And Footage Options

Importing footage

To import footage into a shot, either click on the Import Footage button, or select *File>Import Footage...* from the menu bar to open the Import Footage window. Navigate to the footage you require, click on your chosen sequence and then click Load. Using the preview button you can playback a clip in the import window to ensure it is the correct one. This can save a great deal of time when working on large and complex project with hundreds of shots.

Depending on your operating system, PFTrack can import a number of different movie file formats, such as AVI and QuickTime. PFTrack can also import image sequences (regardless of OS) stored in several image formats such as Targa, DPX, JPEG, TIFF, OpenEXR etc.. A complete list of supported image formats is given in the next section.

If you are loading a sequence of individual frames but only wish to use a selection from the full sequence then you can enter the start and end frame number in the From Frame and To Frame edit boxes before loading. Alternatively, the entire sequence can be imported and you can mark in and out points within the timeline.

PFTrack will check the footage as it is being imported to make sure all the frames are present and correct. This is useful when working with image sequences as a bad file can cause the tracking to fail. If a frame is missing then you should correct this before starting by either duplicating another frame or rebuilding the image sequence. This checking procedure can be disabled from the Footage Import section of the Preferences window.

PFTrack can also use a RAM cache to store frames of the footage in RAM to enable faster playback. This option can be accessed via the Performance tab in the Preferences window (File>Preferences... or PFTrack>Preferences... on Mac OSX) or by choosing Properties from the cache menu at the bottom-right of the interface. The Preference window allows you to define whether or not to use a cache, specify cache size in megabytes (Mb) and whether footage should be cached upon import. In most cases it is wise to set cache size to the highest level you can after considering the needs of your OS and any other applications you are running. Also be aware that operations inside PFTrack such as camera solving and optical flow also require RAM, so make sure you leave at least 500 Mb available for these tasks. Using the memory cache can greatly speed up working within PFTrack and is very effective for high-resolution sequences.

Footage previews can be generated at any resolution or compression ratio, and stored in RAM to accelerate the display of high resolution clips even further. To generate a preview, click the Preview button at the bottom-right of the screen and select Enable from the popup menu. The footage will be compressed and stored in RAM, and playback and scrubbing in the timeline will be greatly accelerated. Note that Preview compression options are available in the Performance tab of the preferences window. Also, footage previews will not affect the quality of feature or geometry tracking – they are used purely for on-screen display, and all tracking operations will still be done using the original plates.

PFTrack stores a reference to the footage in its project files. It does not store the footage itself because these are typically very large in size. This means that if the footage is subsequently moved the project many no longer load correctly. In these situations, PFTrack will prompt you to locate the missing footage manually. PFTrack also provides a number of

footage manipulation options (right click on the footage name in the Shot Overview and select Manipulate... from the popup menu). These operations are applied internally within PFTrack, and will not affect the footage stored on disk. This provides an easy way to adjust your footage for tracking or optical flow without editing the data stored on disk.

Note: If the footage you wish to access is stored on a network drive then it may be difficult, or indeed impossible, to navigate to this location using the standard footage import window. You can overcome this problem by using file dialog boxes that are native to the operating system to give you access to 'Network Neighbours', 'Recent Documents' etc.. The use of this feature is controlled by the 'OS native file dialog' option in Interface tab of the Preferences window (File>Preferences...). However, this comes at a cost – the operating system's native file dialog boxes are of fixed layout and cannot be customized to show footage previews.

Once footage is loaded it can be replaced at any time by another sequence of the same length, or another sequence with additional frames. This is achieved by right-clicking on the footage name in the shot overview and selecting *Replace...* from the popup menu. If the footage is being replaced by a longer sequence, the Synchronize Footage window will be displayed in which you can synchronize the new sequence to the old one.

Supported Image/Movie Formats

Cineon/DPX
OpenEXR
Irix RGB
TIFF
IFF for Maya and Shake
Softimage PIC
Targa
PNG
Windows BMP
JPEG

Quicktime movie support on Windows and OS X platforms.
AVI movie support on Windows and OS X platforms
SGI mov, Quicktime, AVI and MPEG support on SGI platforms.

Note: When loading DPX, Cineon or OpenEXR images into PFTrack, they will be converted and stored internally as 8-bit images. Exporting an un-distorted image sequence of the same type from within the Lens Distortion window will keep the original 10-bit or floating-point image data, rather than using PFTrack's internal 8-bit representations.

Footage Format

After importing footage into PFTrack you should check the Footage Format, to ensure you are viewing the image correctly. To open the Footage Format window, right click on the footage name in the Shot Overview window and select Edit Format... from the popup menu. From this window, you can set the format of the sequence (PAL, NTSC or Film) to change the frame-rate. If you are planning on importing Cooke S4i lens data into the shot, it is important to make sure that the frame-rate is correct in order to ensure that timecode calculations performed by PFTrack will be accurate.

You can also change the type of de-interlacing that is applied to video sequences from within the Footage Format window. The De-interlace options are: None, which is the default and should be used when your footage is not interlaced; Field Averaging will make a frame from

the average between the two fields; Interpolate Lower Field fills in the lower field using information from the upper field; Interpolate Upper Field fills in the upper field using information from the lower field; and finally, Field Separation treats the two fields as if they are separate frames. The field is interpolated so as to maintain the correct resolution but the number of frames and the frame-rate is doubled. If a de-interlacing mode is selected, the Field ordering option becomes active allowing you to set the field order for your image sequence. In low-quality footage such as that captured by a DV camera it will often help to de-interlace footage before feature tracking.

The 3:2 Pulldown option becomes active if NTSC is selected as the footage format. The 3:2 Pulldown field options are: None, WWSSW, WSSWW, SSWWW, SWWWS and WWWSS where 'W' indicates a whole frame, and 'S' indicates a split frame. To simplify the selection of the 3:2 pulldown option, there is also an Analyze Pulldown button that will calculate which option best fits your footage.

Footage Manipulation

You may wish to make some changes to the appearance of your footage using the controls illustrated here from the Footage Manipulation window. Such changes are used to assist the feature/geometry tracking, optical flow and depth map calculations in PFTrack. If your footage is very dark, then it is often beneficial to increase the brightness or contrast before tracking, as this will help PFTrack find a good set of tracking points. Be assured that any changes you do make will only affect your immediate viewing experience and will not alter the original footage in any way.

To open the Footage Manipulation window, right-click on the footage name in the shot overview window and select *Manipulate...* from the popup menu. From here you can adjust saturation, brightness, contrast and sharpness using sliders for fine control. You can also reduce noise in the image, as well as ignore information from one or more colour channels, or change the image to grey-scale. There is also a check box for Histogram Equalization. This will adjust the colour values in an image to provide a good spread between dark and light colours.

PFTrack also allows the manipulation and generation of colour look up tables (LUTs) using standard parameters for the control of Cineon and DPX file import. Custom LUTs can be created and saved for use with other shots, providing better tracking conditions when using Cineon or DPX image sequences.

When working with OpenEXR image files, the black and white-points can be adjusted from within the Footage Manipulation window, as well as the linearity used to map floating-point values to pixel colours for display.

A curve editor is provided to allow even greater control over image manipulations. This lets you to draw linear, smooth or freehand curves that control the red, green, blue and value channels.

All settings can be saved to disk as ASCII files, and reloaded from within the Footage Manipulation window.

Footage Properties

Should you at any point need a quick reminder of any of the basic details regarding a piece of footage you are dealing with within PFTrack, you can obtain this information by right-clicking on the footage name in the shot overview window and selecting Properties... from the popup menu. A window will display a short list of information such as size and duration. If Cooke S4i data has been imported and synchronized with the footage, this data will also be displayed in the properties window.

Importing Cooke S4i Camera Data

Camera data files recorded on-set for Cooke lenses can be loaded by right-clicking on the footage name in the Shot Overview and selecting Import Cooke S4i Data File.. from the popup menu. After loading the data file, a synchronization window will appear, allowing you to check the data contents associated with each frame, and set missing timecodes for each frame of your footage if necessary.

The Smooth Focal Length button can be used to apply a simple 5-tap temporal filter to the focal length values loaded from the file.

Pressing the Okay button will store the focal length data with the camera. The focal length F-Curves can be viewed and edited by clicking the Focal Length item in the Shot Overview.

Importing XML Camera Data

If Cooke lens data files are not available, XML Camera data can also be imported by right-clicking on the footage name in the Shot Overview and selecting *Import XML Camera Data File*.. from the popup menu. After loading the data file, a synchronization window will appear, allowing you to check the data contents associated with each frame.

The Smooth Focal Length button can be used to apply a simple 5-tap temporal filter to the focal length values loaded from the file.

Pressing the Okay button will store the focal length data with the camera. The focal length F-Curves can be viewed and edited by clicking the Focal Length item in the Shot Overview.

An example XML Camera Data file is shown below, containing an animated focal length, indexed by frame number.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<xml>
<info>Test file containing an animated focal length, synced to
frame number</info>
<filmBack units="mm">8.8 6.6</filmBack>
<pixelAspect>1.066667</pixelAspect>
<focalLength frame="1000" units="mm"> 10 </focalLength>
<focalLength frame="1001" units="mm"> 10.1 </focalLength>
<focalLength frame="1002" units="mm"> 10.2 </focalLength>
<focalLength frame="1003" units="mm"> 10.2 </focalLength>
<focalLength frame="1004" units="mm"> 10.3 </focalLength>
<focalLength frame="1005" units="mm"> 10.3 </focalLength>
<focalLength frame="1006" units="mm"> 10.3 </focalLength>
<focalLength frame="1007" units="mm"> 10.4 </focalLength>
<focalLength frame="1008" units="mm"> 10.4 </focalLength>
<focalLength frame="1009" units="mm"> 10.4 </focalLength>
</xml>
```

The <info> tag is optional, and can contain arbitrary text which will be shown in the synchronization window. The <filmBack> and <pixelAspect> tags are optional, although if they are not present, you must ensure that the Camera Parameters window contains the correct preset data before loading the XML file, otherwise the conversion of the focal length values into a field of view may not be correct. Units for the film-back size can be either 'mm' for millimetres, or 'in' for inches.

The <focalLength> tag contains the important camera focal length information. Units can be either pixels ('pix'), millimetres ('mm'), inches ('in') or degrees ('deg', where the focal length value represents the camera's horizontal field-of-view). The frame attribute inside this tag is optional, and when missing, the focal length is assumed to be constant throughout the shot. The frame numbers stored in this attribute must match the frame numbers of your image sequence or movie (Note: most movie files imported into PFTrack have their frame numbering starting at zero). Alternatively, the frame attribute can be replaced by a timecode attribute, such as:

```
<focalLength timecode="00:03:11:05" units="mm"> 10.58
</focalLength>
```

In this case, the timecode attribute can be matched to the timecode of your footage from within the same synchronization window that is used for Cooke data file import.

Note that timecode and frame attributes cannot be mixed inside the same XML data file. It is also not necessary to have a focal length measurement for every frame of your shot. Frames that do not have any measurement are assumed to have an “unknown” focal length during the camera solve.

Lens Distortion

All camera lens combinations introduce some level of geometric distortion into the final image. The amount of distortion will vary depending on the quality of the lens, but if the distortion is significant then straight lines in the scene will appear to be slightly curved in the image. While this may be just the effect a Director required, a significant amount of lens distortion will affect the accuracy of auto-feature tracking and the camera solve.

Typically, consumer grade cameras experience barrel distortion at short focal lengths, and small amounts of pincushion distortion at long focal lengths. This is also true of professional level lens but to a much smaller extent.

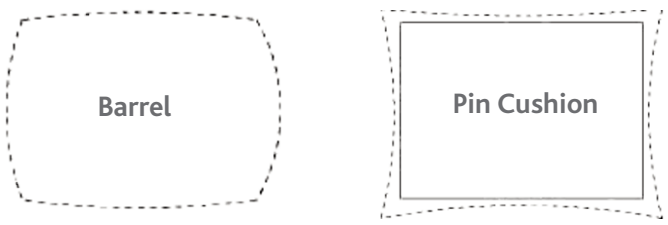

These simple diagrams illustrate the shape of basic types of lens distortions. PFTrack is quite capable of estimating distortion within a sequence but first it has to know what criteria to use to estimate the amount and type of distortion. To provide this guidance it is necessary to identify a distorted line within the footage that you know should be straight. To do this you can either use the semi automatic distortion tools within PFTrack, or trace a line over what should be a straight edge to provide a guide for constructing a lens distortion model. PFTrack can also use known distortion models for Cooke Optics S4 Prime and CXX Zoom lenses. Generally, the amount of lens distortion in an image will vary according to the camera’s focal

length. If the image sequence has been captured at an approximately constant focal length, a single set of parameters will suffice, so you only need to estimate distortion in a single frame. Alternatively, if the image sequence has been captured so that the focal length changes throughout, then distortion should be estimated in two or more frames, and PFTrack will interpolate the data between them.

Semi-Automatic Distortion Correction

PFTrack can identify distinct edges within a frame automatically and these can be used to build up a distortion model. The success of auto-detection will depend on the contents of your shot, and in particular on whether it contains enough long straight lines.

To use line auto-detection, select *Camera>Lens Distortion...* from the menu or right click on the Distortion icon in the shot overview and select the Auto-Detect option. You will notice a number of yellow lines appear over the image. These yellow lines are the ones PFTrack thinks could be straight. To estimate the amount of distortion, find a line that you know should be straight and click on it with the left mouse button. PFTrack will estimate suitable distortion parameters to straighten the line, and the image window will update to show the results of removing lens distortion from the frame.

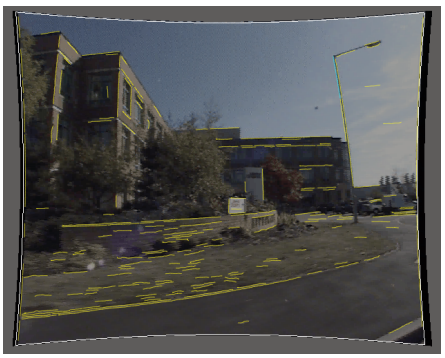

You can use multiple lines at the same time by holding the Shift key whilst clicking with the left mouse button. If the results are somewhat unexpected you can select another line until you get the results you need. The longer the line, the more likely a good distortion model can be estimated. The distortion parameters will be automatically entered into the Lens Distortion window and can be edited afterwards if needs be. Each parameter is activated using the 'Enabled' check box, and its value changed by clicking and dragging with the left mouse button inside the text box. Whilst doing this, you will get visual feedback of the distortion correction in the image window.

Once you are happy with the results you can export un-distorted footage using the Export button. The Reset button will reset the lens distortion parameters to their default state. The yellow distortion lines can be switched off using from the *View>Distortion Lines* menu option, or can be deleted by right-clicking on the Distortion entry in the shot overview and selecting Clear Auto Lines or Clear All Auto Lines from the popup menu. Distortion lines can be safely removed once lens distortion parameters have been calculated, if you wish.

Drawing Distortion Lines

It is also possible to draw the distortion lines manually. To do this you must first put the interface in Distortion Mode by clicking the appropriate toolbar icon (shown left).

Identify an edge that should be straight then draw a line along it. To create end points for your line, click and hold the left mouse button at one end, move to the other end and release the mouse button. This will draw a straight yellow line. If you press shift whilst placing the end points PFTrack tries to snap the point to the nearest significant edge in the image. If you need to adjust the start/end points after they have been placed, they can be dragged around by clicking with the left mouse button. Holding the shift key whilst doing this will display a zoom window, allowing you to place the point more accurately. If you make an error simply select *Edit>Undo* from the menu or press the undo keyboard shortcut (Ctrl+Z by default).

You will notice that when a distortion line is highlighted in the Shot Overview window (as it

will be now) it turns from yellow to light blue. Within the Shot Overview window your line will be identified as Line 1, rising numerically with each new line, although you can change this name if you wish to by clicking on an already selected line and entering a new name.

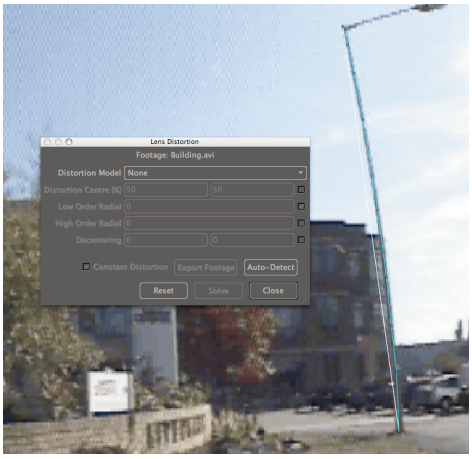

You will see that what should be a straight line in your image will actually curve away from this yellow/blue line. You now have to subdivide your distortion line to make it follow this curve in the image more accurately. To do this, move your mouse to the middle of the distortion line so that a small circle appears, indicating where you can create a new subdivision point. To place that point, press and hold the left mouse button and move the point to a new location that better matches the curve of the line in your image. Again, holding the Shift key before pressing the left mouse button will display a zoom window for greater accuracy. If you want to pan or zoom the image at any point whilst in Distortion mode, press and hold the Navigation mode shortcut (the 'N' key by default) and then use the middle or right mouse buttons. Releasing the 'N' key will revert back to Distortion mode so you can continue editing your lines.

Continue to place a few more subdivision points so your distortion line follows the curve in the image. You don't need to match it at every pixel, but try to make sure that there is no more than a one or two of pixels difference between your distortion line and the edge you have identified in the footage. When you're happy with your line, open the Lens Distortion window shown below by selecting *Camera>Lens Distortion...* from the menu or by right-clicking on the Distortion entry in the shot overview and selecting Edit Parameters... from the popup menu.

The various options in this window are described in more detail in the Lens Distortion Window section of this manual. Generally, you only need to use the low order and high order radial distortion parameters, unless you've used a very special lens to shoot your footage. If your focal length is constant (and so your distortion parameters aren't changing through the shot) then make sure the Constant Distortion box is ticked. Clicking the Solve button will then calculate the lens distortion by attempting to straighten the distortion lines you have drawn in each frame.

PFBarrel

If a shot is subsequently going to have have computer graphics (CG) composited in, it will also be necessary to provide a distortion model to the compositor to allow the CG to be distorted to match the footage. PFTrack can render un-distorted footage to disk, so a compositor or animator can accurately match the CG using a pinhole camera model. A more flexible approach is to use the PFBarrel plugin for third-party compositing applications.

PFBarrel works within compositing applications to apply or remove the lens distortion calculated by PFTrack to your footage. It provides a number of major advantages over other methods and guarantees consistency of correction for all footage in a shot. It can even correct for distortion which changes throughout the shot.

PFBarrel is currently available for Shake, Digital Fusion, Media Illusion and After Effects. For more information on PFBarrel and other application plug-ins, visit The Pixel Farm web site at http://www.thepixelfarm.co.uk

Distortion Grids

The Pixel Farm has produced two distortion grids for you to use if you wish. A distortion grid helps with camera lens distortion by giving clear horizontal and vertical edges that PFTrack distortion lines can then be drawn against. One is a black grid on a white background and the other is a white grid on a black background. Both these grids are included as PDF files in the Docs folder where PFTrack is installed.

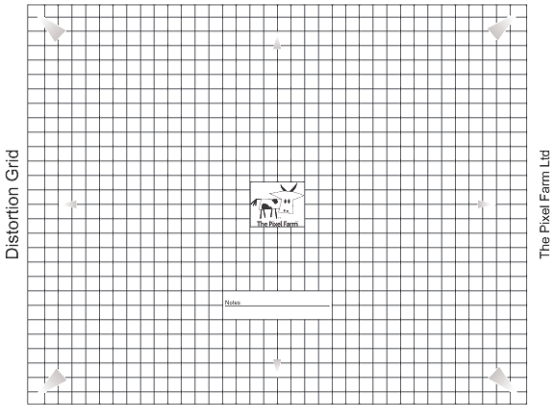

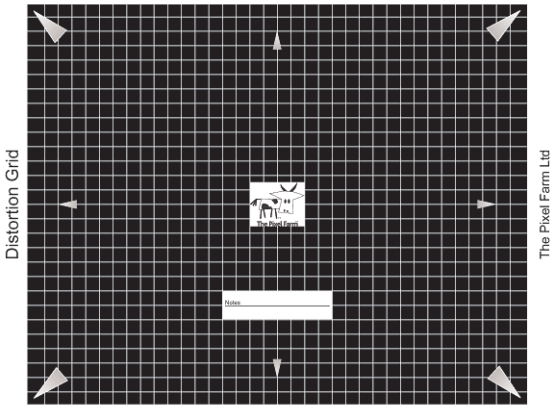

Camera Parameters

To obtain the best results when tracking a camera, you must specify information such as the type of camera used to shoot the footage and the type of camera motion the camera is undergoing. Details of camera settings are often recorded during shooting, so if you have advanced notice that a shot is going to be tracked, it is worth making sure as much information as possible is recorded on-set.

Camera information can by entered in the Camera Parameters window. To open this window select *Camera>Camera Parameters…* from the main menu bar. PFTrack provides three basic types of cameras: Single Camera, Stereo Camera and Motion Capture which are available from the Camera Type drop-down menu. Motion capture and stereo cameras are covered elsewhere in this manual.

The camera's motion can also be specified, and choices are between Free Motion, Translation Only, Rotation Only, Linear Translation + Rotation, Planar Translation + Rotation and Stationary. Different camera motion types can also be specified for segments of the camera path, using the Segment Motion Types window which will be discussed later.

Free Motion is where the camera is free to move, changing its position and orientation. Rotation Only is where the camera is not moving (typically, locked down on a tripod or other mount) but is free to change orientation. The Stationary option is mainly used in the case of object or geometry tracking where you are interested in an object's motion relative to a static camera with a fixed, known focal length.

Note: If the camera is technically moving, but the motion is 'insignificant' compared to the distance from the camera to the scene, then the Rotation Only motion should be specified. For example, if the camera is only moving a few centimetres, and viewing a scene many metres away then there will be no discernible parallax in the sequence. When in Free Motion mode, PFTrack uses parallax to calculate the camera translation – if there is no parallax then the estimate of camera translation and 3D feature positions may be incorrect. In these cases, you should try setting the camera motion to Rotation Only and see if this gives a better quality solve. Also, stationary cameras that have a changing focal length should be tracked using the Rotation Only motion model.

From the Camera Preset pull-down menu you must also select the type of camera used to shoot the footage. A comprehensive list of options is supplied, covering just about every format currently in use today. If, however, your camera type is not shown, or you want to change some of the preset parameters such as film back size or pixel aspect ratio, you can use the 'Custom' preset and enter your own values from a camera sheet. If you wish, you can use the Save Preset button to save the preset to disk so it can be re-used in the future. Note that to ensure the preset file is loaded automatically when PFTrack is next launched, it must be placed in the Presets directory at the location where PFTrack was installed.

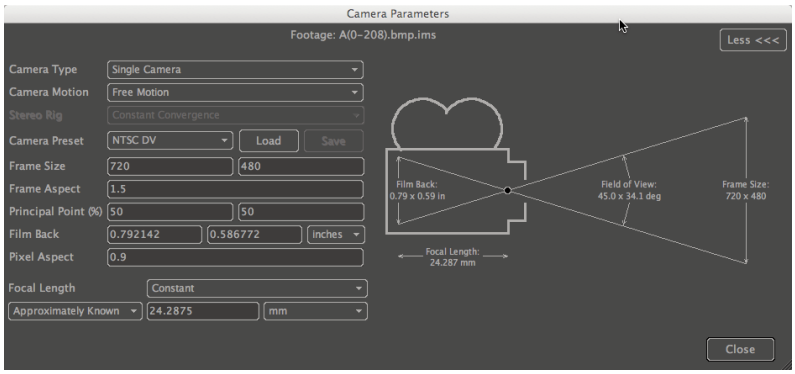

The Camera Parameters window also provides the opportunity to specify information about focal length characteristics. The drop-down menu can be used to choose between Constant or Variable focal length. If the focal length of the camera changes during the shot, it is important to select Variable, otherwise the camera solver may run into trouble. If the focal length remains constant throughout the shot, then select Constant from the menu (the default value). For constant focal lengths, you can also specify whether the focal length is Known, Unknown (the default) or Approximately Known. When it is known or approximately known, you can enter the value in either millimetres, inches, pixels or degrees. Values in degrees correspond to the horizontal field of view.

When known or approximately known focal lengths are entered, the camera diagram on the right of the window will update to illustrate the values that have been entered. This can be used to ensure, for example, that the focal length value you have entered corresponds to a sensible field of view when combined with your camera film-back. If you have no idea what the correct focal length value is, PFTrack will attempt to estimate the value during the camera solve. It is also possible to measure it using the focal length estimation tools described in more detail later in this manual.

Note: Available camera presets are limited by the frame size of your footage. Selecting a camera preset will define the pixel aspect ratio and film back size of your camera. The film back size is only used to convert focal lengths specified in user-friendly units of millimetres into PFTrack's native camera-independent unit of pixels. Therefore, setting the correct film back is not strictly necessary before tracking, unless you want to deal with focal length in units of millimetres or inches, as described above.

The Principal Point is the point about which any focal length changes take place. Typically this will be the centre of the frame. One notable exception to this is if the footage has been cropped, because the centre of the original image will not necessarily be the centre of the cropped image and thus the principle point for the cropped image will be offset. Be aware that when exporting camera data, it is generally the case that a non-central principal point is not supported by third party applications. For this reason, we recommend tracking footage before it is cropped in any way.

Part 3 - Tracking And Solving

Tracking

PFTrack is capable of producing extremely accurate 3D data from film and video footage by estimating the motion of the camera and the 3D position of the tracking points. PFTrack uses two main methods for achieving this: feature tracking and geometry tracking. The feature tracking process is divided into two stages: tracking feature points and solving for camera motion. Breaking the process down this way permits a great deal of user control. If needed, it allows for extensive review and modification of feature tracks before they are used to solve the camera motion.

The second method, geometry tracking, uses imported 3D geometry placed over an image to extract motion data as a one-step process. This is achieved by working out how the geometry should move In order to match the image data in each frame. This is an ideal tool for use in head and face replacement, or for tracking digital set models into an image sequence. Geometry tracking can also provide a very robust means to extract motion information for shots that will not track using feature tracking.

The easiest way of tracking a shot is to use the combined Track and Solve button illustrated on the left. This allows you to track features and solve for the camera motion with one button press, producing 3D data ready for export. In many cases taking this route will produce high quality and accurate 3D data that requires no manual adjustments. If the results are not what you would expect, or the camera solve fails, you are able to run the tracking and camera solve phases separately. This means that feature tracks can be edited and cleaned, and the camera solver can be controlled to produce a better solution.

PFTrack also offers two approaches to the feature tracking procedure itself: Auto-Feature tracking requires no user intervention but still gives you the opportunity to review your feature tracks before you start the solve, as well as providing some level of feature editing. User-Feature tracking allows you to place your own tracking features within the sequence. User-features can also be tracked in addition to auto-features, to provide extra hints about camera movement which can be particularly useful in problematic sections of your footage.

PFTrack is able to track multiple object motions within a sequence and represents these as individual groups. Each group has its own set of motion parameters within the shot overview window. For example you may have a Camera group that contains the camera motion relative to the background scene (this is the default situation), and then a second group that contains the motion of a separate moving object, such as a car, relative to the camera.

Auto-Feature Tracking

Auto-feature tracking is a way of automatically identifying image features and tracking these through the sequence, and requires no user intervention. All you need do is select Tracking>Auto Track from the menu or click on the Auto Track Button illustrated on the left. If auto-features have already been tracked, clicking this button will display a popup menu asking whether you want to discard the existing tracks, or add more. If auto-features have not been tracked in all the frames between the timeline in and out points, you will also receive an option to extend the existing set of tracks. This is useful if you need to extend an existing camera solve into more frames, and don't want to re-track the entire shot from scratch.

As auto tracking proceeds, you will see the features being tracked frame-by-frame throughout your footage in the image window. Green features indicate a good track, and red not so

good. You can also follow progress using the progress bar or the log text window by clicking on the Display Log button at the bottom-left of the timeline. Before tracking, it may help to switch on the *View>Darken* menu option so you can see the feature tracks more clearly against the background.

After tracking has finished, the auto-features will be listed inside the Auto-Features container of the group that has been tracked (the 'Camera' group by default). To change the order the features are listed, right-click on the Auto-Features container in the shot overview and select a sorting option from the popup menu. You can select auto-features by clicking on their names in the shot overview. Holding the Shift or Control keys will allow multiple selections to be made. Auto-feature can also be selected by clicking on them in the image window, whilst the interface is in 'Selection' mode.

Right-clicking on a specific auto-feature in the list will also display a popup menu with additional items to help edit the track and specify how the feature track will be used during the camera solve. You can select a group of features in the image window by clicking the Selection Mode button (or by holding the selection shortcut key, 'L') and then drawing a selection lasso with the left mouse button. All features inside the lasso will be selected. If the Shift key is held, multiple selections can be made at the same time. If you held the 'L' shortcut, release the key once you have selected features to switch back to the previous mode you were in. Once features are selected, they can be deleted by pressing the 'Del' key (or whatever shortcut is assigned to the *Edit>Delete* menu option).

Changing Auto Tracking Parameters

Although you will often not need to alter the default settings used for Auto-tracking, it is nonetheless possible to do so by opening the Tracking Parameters window from the *Tracking>Tracking Parameters...* menu. Full details of each parameter can be found in the Windows section of this manual. The parameters that will often be important are the 'Group', the 'Window Width' and 'Window Height', and the 'Num Features' value.

The Group drop-down menu specifies the motion group into which new auto-features will be placed, and this will be described in more detail in the section on object tracking. If you have not created any additional groups, only 'Scene' will be available, and all new auto-feature tracks will be placed in this group. The Window Width and Height controls the size of the pixel window that will be used to track each feature. Adjusting the window size can sometimes improve the accuracy of feature tracking, but making it too large will mean that features can be rejected too quickly because too many pixels inside the window do not match up from frame-to-frame.

The number of auto tracking points is set to 150, but can be increased or decreased where necessary if the results are not acceptable. In most cases this will be sufficient, but you can increase or decrease this value if you feel you need more or less feature points. PFTrack requires a minimum of 6 points tracked between frames to solve for camera or object motion. However, we recommend you try to ensure that more are tracked if possible for sufficient accuracy, especially when using auto-features.

Increasing the number of auto features does, however, have a direct impact on the speed of the tracking and camera solving procedures and will not always improve results, especially if the additional features are placed at locations that cannot be tracked accurately. In some cases reducing the number of points can improve the solution. In the case of tracking small areas or objects it is a good idea to reduce the number of tracking points because the percentage of bad points with a confined space will increase and produce less accurate results.

When tracking small areas or objects it is a good idea to reduce the number of tracking points because the percentage of bad points with a confined space will increase and produce less accurate results. Alternatively, you can reduce the number automatically by selecting the Area Limit option. This will reduce the number of target features in proportion to the area of the image that is being tracked.

When choosing the number of feature points, it may help to click the Preview button at the bottom-left of the Tracking Parameters window. This will display small dots in the image indicating where features will be placed in the current frame. The preview will be updated automatically after you change any of the Selection Parameters, or any other parameter that affects initial feature placement.

Tracking Length Graph (Track-L)

You can see how many frames are spanned by each feature, as well as find out which frames do not have enough features in them to solve by clicking the Track-L button at the top-left of the timeline to display the tracking length graph. As features are tracked, they will appear in the graph as horizontal bars, spanning the frames in which they have been tracked and colour-coded according to the track error for each frame. Frames that do not contain enough features for the camera solve will be highlighted in red.

You can pan and zoom the graph in the usual way, using the middle and right mouse buttons. To constrain the pan or zoom to either the horizontal or vertical directions, hold the 'X' or 'Y' keys whilst dragging the mouse button. Clicking the small arrow buttons on the right of the timeline will reset the horizontal and vertical graph ranges to show all available data.

The 'A' button can be used to toggle between displaying all features, or only the features which are selected, and clicking the 'O' button will display the features in order of their starting frame. Each horizontal bar in the graph also shows the feature name, but you may need to zoom in to the window to see this, especially if you have a large amount of feature tracks.

Clicking and dragging with the left mouse button will enable features to be selected from directly within the graph. Selected features are drawn with an additional coloured line around each error bar.

Tracking Error Graph (Track-E)

Once features are tracked, you can check the tracking errors by clicking the Track-E button at the top-left of the timeline. This displays a colour-coded graph in the timeline, representing the tracking error for each feature throughout the shot. The error is a measure of how well the pixel window surrounding the tracking point matches the window where the feature was initially placed. A small error (drawn as green in the graph) indicates that the pixel pattern matches very well, and a large error (drawn in red) shows that the pixel pattern could not be matched accurately.

You can pan and zoom the graph in the usual way, using the middle and right mouse buttons. To constrain the pan or zoom to either the horizontal or vertical directions, hold the 'X' or 'Y' keys whilst dragging the mouse button. Clicking the small arrow buttons on the right of the timeline will reset the horizontal and vertical graph ranges to show all available data. The 'A' button can be used to toggle between displaying all features, or only the features which are selected.

Features can be selected from the graph by drawing a selection rectangle with the left mouse button or by clicking on a feature as the mouse hovers over the graph. This allows poorly tracked features to be quickly identified, and either edited or deleted as appropriate. Ensuring

you have a good quality set of feature tracks is the first step to producing an accurate camera solve.

Note that just because a feature has a small tracking error, this does not mean that the feature can be solved accurately in 3D space. For example, a well tracked feature may be placed on a moving object or at the visual junction between two edges of different depth. In these situations, the feature may still track well, so take care and examine your feature tracks if you run into problems during the camera solve.

The Tracking Threshold line drawn in the graph can be used to quickly remove features from frames where they are poorly tracked. The initial value of the threshold will be the same as the 'Max. Residual' value in the Tracking Parameters window: any feature track that exceeds this value will have already been removed. By dragging the threshold line downwards, you can remove additional features if you wish. The threshold line can also be controlled from within the Clean Auto-Feature Tracks window, as described below.

Cleaning Auto-Feature Tracks

For shots containing large amounts of complex motion, it is likely that some of the auto-feature tracks will be incorrect. You can filter out some of the poorly tracked features using the Clean Auto-Feature Tracks window, which is opened from the *Tracking>Clean Auto-Feature Tracks...* menu.

The Tracking Threshold is used to decide whether to remove an auto-feature track from a particular frame. A threshold of 100% corresponds to the 'Max. Residual' value in the Tracking Parameters window. If the tracking error for a feature in one frame is larger than this threshold, the auto-feature will be removed from that frame. By dragging the slider down to a value less than 100%, you will see more and more tracks are removed from the image window. The Tracking Threshold can also be adjusted by dragging the horizontal line in the Track-E window, as described above.

The Rejection Count also takes values between 1% and 100%. This controls the removal of entire auto-feature tracks. An auto-feature will be removed entirely if the percentage of frames where the Tracking Threshold is exceeded is larger than the Rejection Count. For example, if this is set to 50% then any auto-feature that exceeds the Tracking Threshold in more than half of its frames will be removed entirely.

The further you move these sliders to the left, the more bad tracks will be rejected. By careful not to remove too many good tracks whilst doing this, as this will affect the quality of the camera solve.

A new feature in PFTrack 5.0 is the Acceleration Filter. Acceleration is the change in feature track velocity from frame-to-frame, and by enabling the filter, poorly tracked features that zig-zag over the image can be identified and quickly removed. The Max. Acceleration edit box contains the maximum allowable acceleration value that a feature can have before it is filtered entirely. The graph shows a plot of all feature accelerations, sorted by increasing value. Typically, zig-zagging features will appear as a sharp spike on the right-hand side of the graph. By adjusting the Max. Acceleration value, or dragging the filter line inside the graph, features can be filtered as desired. The graph can be zoomed vertically by clicking and dragging with the middle mouse button.

Clicking the A button will automatically estimate a filter level. Note that this level may need to be adjusted further, depending on the type of motion that the feature tracks are undergoing. The F button can be used to reset the viewpoint to show the entire acceleration graph.

User-Feature Tracking

User feature tracking refers to the situation when you manually place feature tracks at specific locations in the footage. Typically, user-feature tracks are used to supplement auto-feature tracks, or to track specific points in the scene that you wish to export for future compositing or animation work.

User-feature tracks are sufficient by themselves to perform a camera solve if enough are placed in the shot. You need at least six user or auto-features tracked between frames to solve for the camera motion, but such a small number may not produce as good a result as full auto tracking. When attempting to solve with user-features, we recommend that eight or more are tracked in the shot, and that they are spread around the image as much as possible. Also, beware that the features should be placed so that they do not all lie on a single flat surface. Doing so may affect the quality of the camera solve, especially when the camera focal length is varying or unknown.

User-features can, however, be very useful in many circumstances. For example, adding one or more user-features to a specific area of interest and then marking them as tracking guides can help the auto-feature tracker 'lock-down' on a particular area of motion. This can be used to help out in situations where you find that the auto-feature tracker is not placing features in a desired area. You can also employ user-features to pick out and name specific points within the scene to track. These points will be represented in the final 3D data model you export for animation or compositing.

There are two ways to create new user-features. The first is to select *Tracking>New User-Feature* from the menu (Shift+F3 keyboard shortcut by default). This will add a new feature into the User-Features container in the shot overview. The feature will already be selected, and clicking with the left mouse button inside the image window will place the feature. Whilst it is selected, the feature position can be changed by clicking and dragging inside the search window with the left mouse button. To place the feature more accurately, hold the Shift key whilst clicking the left mouse button to display a zoom window with a close-up view of the area around the mouse pointer. To place the feature, release the mouse button at the desired position. If you wish to move the close-up view around, just drag the mouse cursor outside the zoom window, and it will update to match your new mouse position.

The second way to create new user-features is to click the User Feature mode button in the toolbar to place the interface in User-Feature mode. When in this mode, clicking in the image window will create a new selected feature at that location (the Shift key can be used again to display a zoom window). Each new feature will be shown in the shot overview, and will be selected by default. Whilst a feature is selected, it can be moved around the image as described above. To create a new feature, click elsewhere in the shot overview (for example, on the “User-Features” container item, which will keep you in User Feature mode) to deselect the user-features, and then click again in the image window.

When a feature is selected, the size of the tracking and search windows can be adjusted by dragging their borders with the left mouse button. By default, user-features will be created with the window size and search size that is specified in the Tracking>Tracking Parameters... window.

Once a feature has been placed in one frame, it needs to be tracked through the shot. After making sure the feature is selected, click the “<<” or “>>” buttons at the bottom of the image window to track the feature backwards or forwards. This will start tracking the feature in the direction you have chosen. You can track more than one user-feature at a time by selecting multiple features before pressing the buttons.

User-features can also be tracked one frame at a time by clicking the “<” or “>” buttons. When tracking complex shots, it may also help to switch on the *View>Centre* On Selected Feature option from the main menu. When this is active, the image window will be automatically positioned so the feature remains at the centre. This can be useful to see when the feature track drifts from its intended location.

User-feature tracking will continue until either the tracking error has exceeded the Max. Residual tracking parameter, the feature moves out of frame, or you press the Escape key. If the feature tracking fails, you can re-position the user-feature by hand to create a keyframe at the correct location, and continue tracking. The Dope-Sheet can be used to see where user-features have been keyframed. A light-blue keyframe marker will appear in the Dope-Sheet at each frame where you have placed a user-feature by hand. If multiple keyframes have been placed, tracking is performed from one key-frame to the next. Keyframes can be deleted by right-clicking on a feature and selecting Remove From Frame from the popup menu, or by selecting the range of frames in the Dope Sheet by clicking and dragging with the left mouse button, followed by pressing the Delete key.

The remaining buttons at the bottom of the image window are used to remove the feature from the current frame (the 'R' button), or from the current frame and frames before/after (the 'R-' and 'R+' buttons). The Re-centre button is used to centre the user-feature inside a tracking marker, typically one that has been placed specifically in the scene to assist with tracking. The tracking marker must exhibit significant contrast compared to its background. Re-centering may not function correctly when the tracking marker is particularly complex, or placed against a complex background. Uniformly coloured markers set against a uniform background will re-centre most accurately.

You can select user-features by clicking on their names in the shot overview. Holding the Shift or Control keys will allow multiple selections to be made. User-feature can also be selected by clicking on them in the image or 3D Viewer window, whilst the interface is in Selection or User Feature mode.

Right-clicking on a specific user-feature in the list will display a popup menu with additional items to help edit the track and specify how the feature track will be used during the camera solve. User-features can be lassoed and deleted in the same way as auto-features.

When creating user-features for tracking additional motion groups, the user-feature will be placed in the currently selected motion group. To select a different motion group, click on the group name in the shot overview.

Editing Feature Track F-Curves

Feature track F-Curves can be viewed and edited after a feature is tracked by opening the F-Curve window using the F-Curve button at the top-left of the timeline, and selecting either a user or auto-feature from the shot overview. Clicking the buttons on the right of the timeline will reset the graph viewpoint to show the X and Y tracking curves in red and green respectively.

All the F-Curve editing tools are available, such as curve smoothing and keyframe interpolation. Further details of F-Curve editing are given elsewhere in this manual.

Masks

In order to get a good quality camera solve, it is important when using auto-features that the features are placed on stationary objects. If you have actors or vehicles moving around in your shot, any auto-feature that is placed on these may have an adverse affect on the quality of the camera solve. PFTrack is fairly tolerant of incorrectly placed auto-features, and if there are only a few in the scene then they will automatically be identified and ignored. However, if there are a large amount of these features, they can end up dominating the camera motion estimate and will produce incorrect results because it becomes increasingly difficult to distinguish the motion of the camera from the motion of the object. If your footage contains moving objects and these occupy a significant proportion of the frame, then you can use masks to ignore those parts of the image and ensure that no auto-features are tracked over them.

If you are trying to track additional object motions then the opposite is true: you will want to mask out the stationary parts of the background scene and ensure that PFTrack placed auto-features correctly over the object.

Creating Masks

The simplest form that a mask in PFTrack can take is the white Frame Boundary lines that appear at the edge of the image sequence. These can be used to very quickly crop the edges of an image by clicking and dragging them to the required position with the left mouse button. This form of mask is designed to quickly crop borders from your footage without the need to build additional mask by hand. The frame boundary can be switched on and off using the *View>Frame Boundary* option in the main menu.

The second type of mask used in PFTrack is a spline-based mask that can be drawn by hand around specific objects. There are two ways to create spline-based masks: the first is to select *Tracking>New Mask* from the main menu (or press the keyboard shortcut, Ctrl+F4). This will create a new mask and display its name in the Masks container of the shot overview. When the mask is selected, clicking with the left mouse button in an image window will draw the mask boundary. Each click with the left mouse button will place a new boundary point. To stop placing boundary points, click the right mouse button.

The second way to create a mask is to put the interface into Mask Mode by clicking the Mask button in the toolbar. Whilst in mask mode, if there is no mask selected in the shot overview, clicking in the image window with the left mouse button will create a new mask and add a boundary point at the mouse location. Additional clicks with the left mouse button will place additional boundary points, and clicking the right button will finish the mask. If you want to create an additional mask, de-select the current mask by clicking on another item in the shot overview (for example, the 'Masks' container item) and then click again in the image window. These three pictures show the mask being created.

Once a mask is created, the masked pixels (i.e. The ones that will be ignored during auto-feature tracking) are coloured pink. Boundary points and edges can be edited whilst in Mask Mode by clicking and dragging with the left mouse button. Right-clicking on a finished mask in the image window or the shot overview will display a popup menu showing different options to control the behaviour of the mask. An additional popup menu is available for each boundary point, and you can display this by right-clicking in the image window whilst the mouse is hovering over a point.

Extra points can be inserted into an existing mask by hovering the mouse cursor over the inside of the mask near the place you require the extra point, right-clicking and selecting Insert Point from the popup menu. Once the point is created, it can be positioned more accu-

rately at the correct place in the image, as normal.

Masks can also be imported from other shots by right-clicking on the 'Masks' container item in the shot overview and selecting Import Mask File... from the popup menu. Mask files are contained in the Masks folder where the shot was saved.

Rounding the corners

By default, masks are created using straight lines between each boundary point. In some cases, you may wish to use Bezier curves instead, especially if you are drawing a mask around a curved object.

To draw a Bezier curve whilst creating the mask, click and hold the left mouse button when placing a new boundary point and move the mouse. Two Bezier control arms will appear, and dragging the mouse whilst the left button is held will adjust the arm positions. Release the mouse button when the arms are in place to move on to the next boundary point.

Boundary points can also be changed from linear to Bezier by right-clicking on the point in the image window and selecting or un-selecting the Linear menu option. You can also use the Continuous menu option to toggle between moving the Bezier controls arms together or independently. Whilst in Continuous mode, adjustments to one arm will also affect the other to ensure that the mask boundary that passes through the point is smooth .

Animating Masks

After you have drawn a mask in one frame and play through your image sequence, it may become apparent that the mask does not cover your moving object in all frames. To correct this, you can move the boundary points and edges around in additional frames to ensure the object is always covered. As each point is moved, it will turn from white to yellow, indicating that the point has been keyframed at that position. It is not necessary to position the boundary points in every frame – PFTrack will interpolate between their positions in-between each keyframe, so you can animate a mask around a moving object very quickly.

If you want to move the entire mask as a whole, rather than moving every point individually, then make sure the mask is selected in the Shot Overview (click on the mask name if not) and press and hold the Shift key to display a transformation widget. Clicking anywhere inside the dotted boundary and dragging the mouse will move all the mask points at the same time. You can also rotate or re-size the entire mask by clicking and dragging the square handles around the widget edge, or in the centre. The centre of rotation can be adjusted by dragging the cross with the left mouse button. You can also use the lasso tool when in Selection Mode to lasso a set of mask boundary points and move these as a group. Mask boundary points will turn blue when selected.

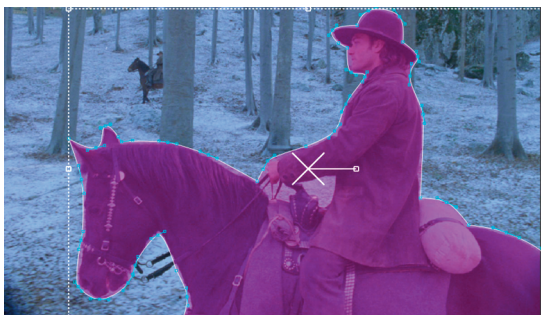

Remember that masks don't have to be too accurate to perform the task of removing feature points from a moving object. When auto-feature tracking, the results from a quickly drawn mask will often be no better than from a mask that is very accurately rotoscoped around the object boundary, so don't spend too much time refining your masks unless it is absolutely necessary.
After you have finished drawing a mask, you may wish to lock it to prevent any further accidental adjustments. You can do this by right-clicking on the mask and selecting the Lock option from the popup menu.

The Dope-Sheet will draw yellow keyframe markers in the frames where the mask has been edited. You can open the Dope-Sheet by clicking on the Dope-S button at the top-left of the timeline. Keyframes can be moved around, if necessary by clicking and dragging with the left mouse button. Right-clicking on a keyframe will display a popup menu allowing you to copy, paste or delete a keyframe. To paste a keyframe after copying, you must first select a frame to paste into by clicking and dragging with the left mouse button (the selected range of frames will be highlighted in yellow). You can also move a group of selected keyframes by clicking and dragging within the dope sheet using the left mouse button. To close the dope sheet and move back to the standard timeline view, click on the Dope-S button once again.

Note: if you wish to scrub to a different frame in the timeline whilst the dope sheet is open, you must click the left mouse button in the cache indicator area just underneath (this is the area containing the white diamond part of the current frame indicator).

Tracking Masks

An alternative to keyframing the movement of a mask in several frames is to attach a mask to a user feature that has already been tracked onto an object. This can be used in conjunction with key framing to compensate for the changing shape or scale of an object.

To attach a mask to a feature, the first step is to track a user feature onto the object you wish to mask out (it is best if the feature is tracked over as many frames as possible). Once this is done, draw your mask in the usual manner and then right-click on the mask and select 'Attach to' from the popup menu. If you have a number of user features in your sequence then select the one you want to attach the mask to. You can then play through your sequence and the mask should stay in position relative to the object and the user feature.

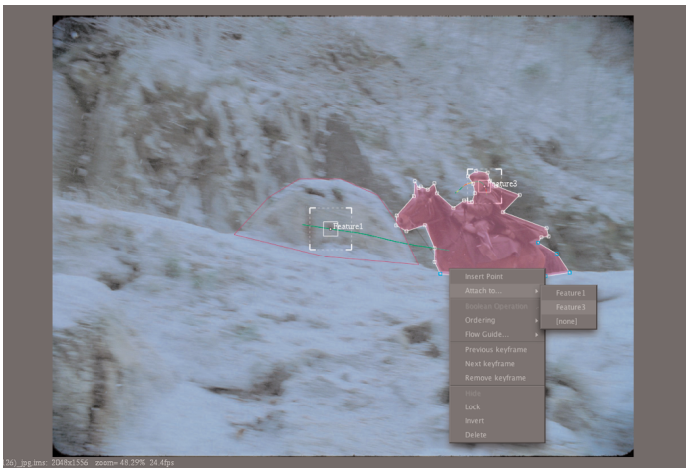

Remember that if you attach a mask to a user-feature, you may not want that user-feature to be included in a later camera solve. In this case, you can make sure the user-feature is ignored by right-clicking on it and selecting Ignore During Camera Solve from the popup menu.

Inverting Masks

Masks can be inverted so that the area inside the boundary defines the space that tracking points will be selected in, rather than excluded from. To invert a mask, right-click on the mask and select Invert from the popup menu. Inverting masks can be useful when tracking a single moving object as well as camera motion: The mask can be used to exclude the object from whilst tracking auto-features for the camera, and then inverted to exclude the rest of the scene when tracking features onto the object.

Boolean Operations on Masks

In PFTrack it is possible to alter the way that multiple masks interact with each other by choosing a Boolean operation from the mask popup menu. The order that the Boolean operations are applied depends on the order of the masks in the Shot Overview. Masks are created and listed in furthest-to-nearest order of depth from the eye, so the first mask in the list is the one that is furthest from the camera, and the last mask in the list is the one that is nearest.

The Boolean operation assigned to the second mask will affect how it interacts with the first mask. Similarly, the Boolean operation assigned to a third mask will affect how it interacts with the first two masks etc.. Boolean operations are not available for the furthest mask from the camera (the first one that was created).

After an additional mask has been drawn, right click on the mask name in the Shot Overview (or right click on the mask in the image window), and from the Boolean operation menu select either **AND, OR, XOR** or **DIF**. With AND selected, the mask will be coloured only where the mask intersects with the previous masks. With OR selected, the mask will be coloured where either this or the previous masks are drawn. With XOR selected, the masks will be coloured only where this mask does not intersect with the previous masks. With DIFF selected, the mask is used to remove an area from the previous masks.

You can change the order of the masks in the shot overview by using the Move Away, Move Closer, Send to Back and Bring to Front options in the Ordering popup menu. This will also affect the way that the Boolean operations are applied.

Image Masks

The third type of mask used in PFTrack is an image mask, which is simply a movie or image sequence that can be loaded from disk. You can import an image mask using the File>Import Mask Image... menu option. Import the mask image sequence as you would any other footage. If you import a black and white image, white areas will be used to mask out regions in the frame. If you import a colour image, it will be converted to black and white internally. Note that your mask image does not have to be the same size as your background footage: it will be resized automatically when imported.

If the image sequence that you use as a mask is not the same length as your background footage, you can adjust the start frame of the mask by dragging the keyframe indicator in the Dope Sheet. The behaviour of the mask before the first frame and after the last frame can also be adjusted by right-clicking on the mask name in the shot overview and choosing an option from the Outside Frame Range... popup menu. Options are Clamp to clamp the first/last frame, Repeat to repeat the image sequence, or Filled/Empty to use a full/empty mask when no image is available.

Colour Key

The final type of mask that can be used in PFTrack is the Colour Key. This allows specific ranges of colours in the image to be used to define masked pixels, when tracking green or blue-screen shots for example. You can open the Colour Key window by selecting *Tracking>Colour Key...* from the menu or by pressing the Colour Key button in the toolbar.

The colour key is able to perform in four different colour spaces, namely HLS, YUV, YIQ and RGB. It is only possible to have one type of colour space per colour key but it is possible to have more than one colour key active at the same time.

To create a colour key, first select the colour space you wish to use (the default is HLS). Once you have selected the colour space, click the Add button, and paint in the image window with the left mouse button to select the colours to be included in the mask. A single click on the image will select the colour of an individual pixel, and dragging the mouse whilst the left button is held will select a range of colours. You can paint multiple times to extend the range of colours you require for the key. You can also drag the vertical bars in the colour ranges with the left mouse button to increase or decrease the range of colours.

When tracking a green or blue-screen shot with tracking markers, two colour keys can be used: one for the green/blue-screen background and another for your tracking markers (assuming the tracking markers are different colour). Once the key has been created, you will have everything except the actor masked out. Clicking the 'Invert' button in the colour key window will then invert the entire mask, ensuring that the actor is masked out and the tracking markers on the screen are visible and available for auto-feature tracking.

Solving Camera Motion

Solving for the camera motion will attempt to estimate the path of the camera through the scene, as well as estimating the 3D position of auto and/or user-features. Taking all the 2D tracking points along with any further information you may have supplied relating to the camera parameters, PFTrack will calculate all the missing camera parameters (such as focal length) and motion, and construct a point-based 3D model of the scene depicted in your footage.

The time this takes depends upon the length of your sequence, the complexity of the camera motion and the speed of your CPU. The camera solver requires 6 or more features to be tracked between each frame, but will of course function better with more features provided they have been tracked accurately. These can be either auto-features, placed automatically, or user-features that are placed and tracked by hand.
Before solving for the camera motion, it is very important to ensure that the camera parameters are correct, as described in the Footage and Camera Options section of this manual. The current camera parameters can be viewed by selecting the Camera>Camera Parameters... menu option. In particular, make sure that the Camera Motion and Camera Preset are correct, including the Pixel Aspect ratio. If your camera is zooming, also make sure that the Focal Length control is set to Variable.

To start solving for the camera motion, select *Camera>Solve Motion...* from the main menu, or click the Solve button in the toolbar. The Solver Window will appear, showing controls that affect the way that the camera motion is generated. For most shots, you will not need to change these options, and can press the Start button to begin the camera solve. A popup window will appear showing the progress of the solve (Note: these popup windows can be deactivated using the Hide new progress windows option in the General tab of the Preferences window). Once the solve completes, the feature tracks will turn into small dots to showing their 3D position, and the structure of the scene can be viewed by opening a 3D Viewer window where you can navigate around the scene using the mouse buttons (the Window>New Viewer menu option, or press the New Viewer button on the right of the toolbar).

When the camera solve is completed you will see that your footage is now covered with coloured dots instead of feature tracks. The green dots should make up the majority as these represent those 3D features which match their track position well. The red dots show a poor 3D point, which deviates from its track position. You can see exactly how poor by switching on the View>Projection Errors menu option. This will draw a small line from the 3D feature point to the corresponding 2D track point. You can also get a numerical value of the error by hovering the mouse pointer over a feature until a popup menu appears. White dots indicate where a feature wasn't tracked into that frame but is visible after the camera solve.

It is also possible to carry out the tracking process and camera solve process in a single operation by choosing the Tracking>Track and Solve menu option (or by pressing the Track and Solve button, or using the keyboard shortcut Shift+F6). The camera solve in this case will be performed using the default Solver Control settings.

Controlling The Camera Solve

Assuming that your feature tracks are accurate, the most important parameter to control the quality of the camera solve is the Keyframe method, available in the Solver Control window. PFTrack uses three keyframes to build an initial camera path for part of the shot, and getting this initial solution correct is very important. Using the Automatic keyframe method, PFTrack will try to choose the best place to position these three keyframes, but for complex shots, placing the keyframes manually can often improve the overall quality of the solution.

To view the position of the keyframes PFTrack has chosen after the camera has solved, click the Dope-S button under the main view window, and make sure that no feature or object is selected (for example, by clicking on the footage name in the shot overview) . This will open the Dope Sheet in the time line area and display the keyframes as orange bars. To remove a keyframe, simply drag-select a set of frames in the timeline with the left mouse button and press the 'Delete' key.

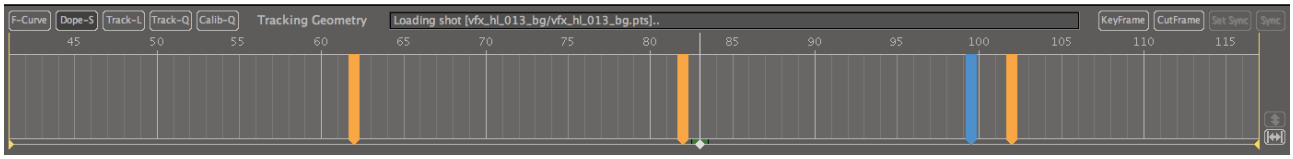

If the solve has failed or is not of sufficient quality, you can place one, two or three keyframes manually by pressing the KeyFrame button at the top-right of the timeline. Try to place them in positions where you can see significant camera motion and parallax in the shot. If you place one keyframe, PFTrack will automatically choose two more on either side. If you place two keyframes, PFTrack will automatically pick a third. It is important that plenty of features are tracked between these keyframes, so the keyframe positions must be chosen to balance the amount of camera translation (as much as possible) without reducing the number of features that are common to all keyframes. Placing the keyframes too far apart may mean there aren't enough features to build the initial solution, and placing them too close together may mean there is not enough camera translation to get a good solve. A small amount of experimentation can be helpful here, and in most situations, placing a single keyframe at a position where there is significant translation and parallax will suffice. Once keyframes are placed, make sure the Keyframe Method is set to Manual in the Solver Controls before solving again. Keyframes can be removed by clicking the KeyFrame button again.

If the camera solve fails at any point (or is aborted), PFTrack will still give you a partial camera track for as much of the shot as possible. PFTrack can also generate an initial solution (often very quickly) which can be refined and edited, and used as a basis for the complete solution by extending the solve into the remaining frames using the Extend Solution option available in the Solver Controls.

Additional Solver Control options are available that affect the way that the camera motion is estimated. The High accuracy and Initial solution only options are mainly used when trying to track problematic shots. Initial solution only will halt the camera solve once the three camera keyframes have been identified and the camera path between them has been estimated. The initial solution is always used as the starting point for the entire camera path, so ensuring that the initial solution is accurate before the remaining frames are added is important.

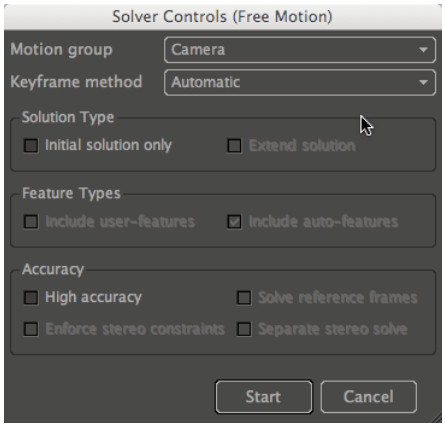

The High accuracy option controls how additional frames are added to the initial solution. By default,

these frames are added in chunks, according to how well PFTrack thinks the solution is progressing. Selecting this option will ensure that additional frames are added one at a time, which will increase overall accuracy, but will also increase the time it takes to finish the camera solve. These options are discussed in more detail in the section on Trouble-Shooting A Camera Solve.

The other options in the Solver Control window are used to include (or discard) user or auto-features from the solve. If you have placed user-features to guide auto-feature tracking, but don't want them included in the solution then you can ignore them by un-ticking the Include user-features option.

Finally, if the camera path has not been solved for all frames between the in and out-points in the timeline, the Extend solution option will be available. Selecting this will extend the existing solution by adding missing frames until the entire camera path is complete. This is useful in situations where you need to extend a camera track and don't want to re-generate it from scratch. The Extend solution option can also be useful when trying to track difficult shots, because it allows the initial solution between three keyframes to be extended throughout the shot.

3D Viewer

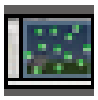

After the camera has been solved, the 'New Viewer' button becomes active at the right of the toolbar. Click this to create a new 3D viewing window in the main image viewing area (or select *Window>New Viewer* from the menu). This new window displays a 3D representation of your animated camera path and the point cloud using the same colour coding as is displayed in the footage window. If the footage has a Rotation Only camera motion, then the 3D feature points will be located on the surface of a sphere which is centred at the camera position. This is because a Rotation Only camera motion does not contain any parallax in the image data, which means it is only possible to calculate the direction of the features and not their distance from the camera.

If you find that both windows do not fit fully within the main image viewing area, simply click the Tile Windows button to automatically re-size the windows to fit the available space, or select *Window>Tile Windows* from the menu. You can also use this to cycle between different window layouts by repeated clicks on the button.

You can navigate the 3D image window in any direction to view the point cloud from whatever angle you desire using the left, middle and right mouse buttons to rotate, zoom and pan the viewpoint. You will notice the camera path is shown as a line the 3D window. This line gives you a good idea of how well your shot had tracked and allows you to see any jumps or badly tracked frames. You can open any number of 3D windows, and each can be changed from a 3D perspective to orthographic viewpoint using the popup menu at the bottom-left of the window. To reset the viewpoint position for the active window, click the 'Fit To Window' button in the toolbar or select *Window>Fit To Window* from the menu.

The viewer windows can also display camera clipping planes if required, by selecting the *View>Clip Planes* option from the main menu. The 'Near' and 'Far' edit boxes in the 3D Viewer windows can be used to adjust the clip-plane distances.

Note: Although it is usual today for 3D applications to work on the basis that the Y direction is up, you can change to using Z as up from within the Calibration tab of the Preferences Window.

Camera Error Graph (Cam-E)

Clicking the 'Cam-E' button above the timeline will display graph showing the overall RMS (Root-Mean-Squared) error for the camera solve. This is a measure of how well the 3D feature points match their 2D track positions when viewed by the camera. Two graphs are displayed: the first shows the RMS Error for all feature points in each frame, and the second shows the Inlier RMS Error, which is the error for only the 'inliers' in each frame. An 'inlier' is a point which projects below the threshold used to distinguish good points from bad: inliers are the green points mentioned above, and 'outliers' are the remaining red points (the inlier threshold can be changed in the Calibration tab of the Preference window).

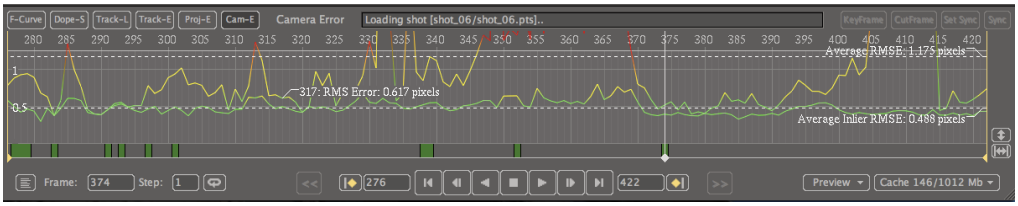

Hovering the mouse over a graph will display a readout showing the frame number and the error value. You can zoom or pan around the graph using the middle and right mouse buttons. To restrict motion to either the horizontal or vertical directions, hold the 'X' or 'Y' keys whilst dragging the mouse. The buttons at the right of the timeline can be used to reset the horizontal and vertical range of the graph to their default values.

The Camera Error graph can be used to quickly identify problems with your feature tracks and camera solve, if the situation arises. For example, the difference between the RMS Error graph and the Inlier RMS Error graph can be used as an indicator of which frames have a significant amount of outliers.

Projection Error Graph (Proj-E)

The Projection Error graph is displayed by clicking the 'Proj-E' button above the timeline. This will display the projection error for each feature in each frame. The projection error is a measure of how well the 3D feature point matches its 2D track position when viewed by the camera. Assuming the 2D track is correct, a large projection error means that the feature point was not placed in the correct 3D position, or the position of the camera is wrong.

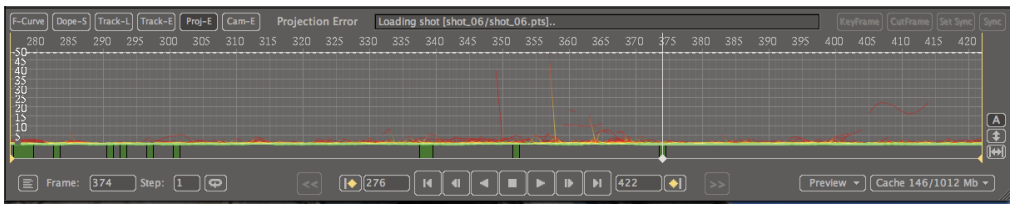

You can zoom or pan around the graph using the middle and right mouse buttons. To restrict motion to either the horizontal or vertical directions, hold the 'X' or 'Y' keys whilst dragging the mouse. The buttons at the right of the timeline can be used to reset the horizontal and vertical range of the graph to their default values.

Hovering the mouse over a graph will display the feature name as well as the error value. You can select features directly from the graph by clicking on them or drawing a selection rectangle with the left mouse button. This provides a quick way to identifying which features are not matching their 2D track positions correctly. You can switch the image window to display either the solved 3D feature positions or their 2D track positions using the *View>Solved Features* menu option. Alternatively, switching on the *View>Projection Errors* menu option will draw a red line between the 3D feature point and the 2D track. You can also get a numerical readout of the error by hovering the mouse over a feature point until a popup box is displayed.

Once poorly solved features have been selected, they can be deleted if necessary by pressing the 'Delete' key, or selecting *Edit>Delete* from the main menu. If a 2D auto-feature track looks correct, but the feature still has a large projection error then the error may be caused by the camera solver not concentrating enough on positioning that feature. This may be because there are lots of other features elsewhere that are overwhelming the camera solver. To correct this, you can switch the auto-feature to 'Hard Constraint' and then use the Improve Solution button to try and force the camera solver to lock-down onto the feature more accurately. This is covered in more detail in the section on Trouble-Shooting a Camera Solve.

Removing Badly Solved Features

After solving for camera motion, you can remove poorly solved feature points from the solution in a similar manner to cleaning auto-feature tracks. Select *Camera>Clean Auto-Features...* from the main menu to display the cleaning window (shown here).

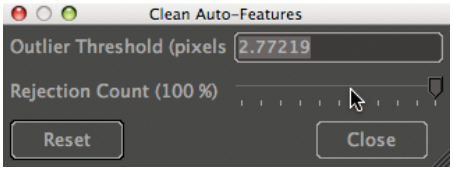

The Outlier Threshold is the threshold used to decide whether to remove an auto-feature from individual frames. By default, the threshold is set to the largest projection error of all features. The threshold corresponds to the horizontal "Outlier Threshold" line that is drawn in the Proj-E timeline graph. Reducing this threshold (or dragging the line in the Proj-E graph) will remove auto-features from frames in which their projection error is too high.

The Rejection Count is measured as a percentage from 1% to 100%, and controls the removal of entire auto-features from the solution. An auto-feature will be removed entirely if the percentage of frames where the Outlier Threshold is exceeded is larger than this value. For example, if this is set to 50%, then any auto-features that exceed the Outlier Threshold in more than half of their frames will be removed entirely.

Individual feature points can also be deleted by hovering the mouse over them and pressing the 'Delete' key whilst the interface is in Navigation or Selection mode.

Viewing and Editing Camera F-Curves

To find the value of any of the camera parameters or examine how they vary throughout the sequence, open the 'Camera F-Curves' container in the shot overview panel and click on a parameter name to display a graph in the F-Curve editor window. You can choose to display the camera translation, rotation or focal length separately, or display them together by clicking on one of the container names. The graphs are colour coded to help distinguish between them when more than one is displayed at the same time. You may wish to click the two buttons on the right of the timeline to set the horizontal and vertical range of the window so the entire graph is displayed.

Hovering the mouse over a point in the graph will display a readout showing the frame number and the value of the graph at that frame. Clicking the middle mouse button within the graph area and dragging the mouse allows you to zoom horizontally and vertically. If you click and drag with the right mouse button in the graph area, you can scroll horizontally and vertically To restrict the motion to either horizontal or vertical direction, hold the 'X' or 'Y' modifier keys. Clicking the F-curves button again will close the graph window and switch back to the standard timeline.

Editing F-Curves

The F-Curves for each camera parameter can also be edited as well as viewed. Editing an F-Curve directly does not affect the current camera position, but is instead used as a target when improving the camera solution. In this way, if you see a kink or jump in one of the camera motion parameters, you can edit the curve to remove the error, and then let PFTrack update the overall solution to try and match the curve you have drawn.

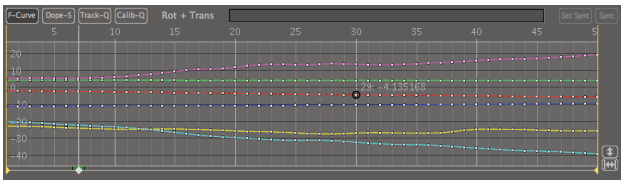

The two main F-Curve editing modes are 'smoothing' and 'interpolation'. In addition to this you can also simply click on a point and drag it into position. As you edit F-Curves you will notice a grey dotted line appearing. This line shows the path of the original F-Curve before editing. Pressing the undo shortcut (Ctrl+Z) at any point will undo the last edit to the F-Curves.

To smooth an F-Curve, first select the range of frames that require smoothing by clicking and dragging with the left mouse button. The frames that are selected will highlight in yellow. If you only have one F-Curve visible, pressing the '-' or '+' keys will smooth or un-smooth the F-Curve over the frames you have chosen. On some systems it may be necessary to hold the 'Shift' key down to access the '+' key. As F-Curves vertices edited, they will turn white. If you have multiple F-Curves displayed at the same time, you must hover the mouse over the curve you wish to smooth before pressing the '-' or '+' keys.

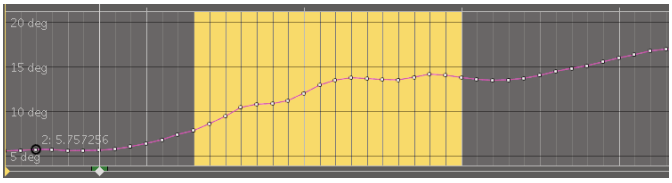

F-Curves can also be interpolated between keyframes using Bezier splines. To create spline keyframes, right-click on a point on the F-Curve and choose Add Key from the popup menu. This will place a keyframe at that point, indicated by a small square box. After two or more keyframes are placed, spline handles will appear that can be used to adjust the spline position. Again, note that the F-Curve has turned white in the frames where it has been edited. Keyframes can be removed by selecting Remove Key from the popup menu. If you wish to adjust the Bezier handles independently, de-select the Continuous option from the menu.

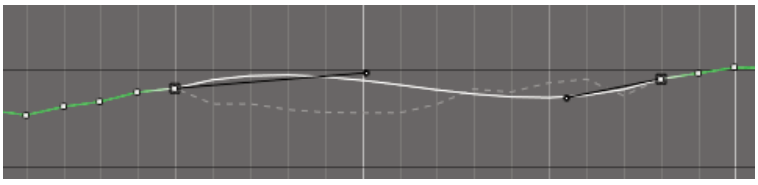

To revert the F-Curve back to its original state, and ignore the F-Curve target whilst improving the camera solution, select a range of frames by dragging with the left mouse button, right-click on a point and select Remove Target from the popup menu. Doing this will change the F-Curve from white (which indicates the curve was edited) back to its original colour.

Finally, you can also convert an F-Curve graph into a set of Bezier splines automatically. First select a range of frames by clicking and dragging with the left mouse button. Then right-click on a point and select Make Spline from the popup menu. PFTrack will then produce a Bezier spline covering the range of selected frames that closely matches the original F-Curve. This spline can then be edited as necessary.

Using F-Curves to Improve the Solution

Once the F-Curves targets have been edited, the camera solution must be updated to account for their new positions.

To copy an edited F-Curve directly into the proper camera motion, right-click on an edited point and select Copy Target To Camera from the popup menu. Remember that after doing this, you may want to switch off *View>F-Curve Target Motion* from the menu in order to revert to the proper camera viewpoint. Copying F-Curves into the main camera will remove the F-Curve targets from the graph.

Alternatively, the entire solution can be adjusted to try and meet your F-Curve targets. This has the benefit of adjusting the other camera parameters and feature positions automatically. For example, if you have edited the camera translation to remove a spike in the camera path, it may be that the camera rotation must also be adjusted to ensure that feature points are matched correctly in the image. This is done by choosing Camera>Improve Solution.. from the main menu, or by pressing the Improve Solution button in the toolbar.

To ensure the F-Curve targets are properly accounted for, make sure the F-Curve box is ticked before pressing the Start button. It is possible to alter the influence of the elements taken into account when improving the solution. This is done by moving the individual sliders to alter the relative weighting for each type of constraint. Further details of these parameters are given in the section describing the Improve Solution window.

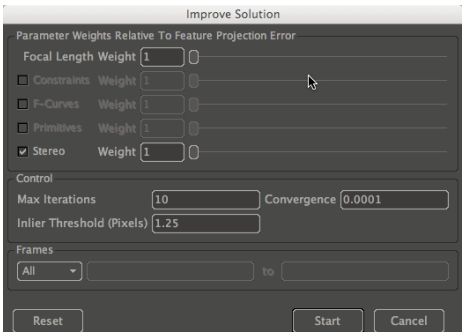

Note: don't make too many large edits to multiple F-Curves at the same time, otherwise the solver may not be able to satisfy all the constraints correctly. If the solution has not been improved after running Improve Solution, you can press Ctrl+Z or choose Edit>Undo from the main menu to revert back to the previous camera solution. In such cases, you may be able to obtain a better result by copying the F-Curve targets directly into the camera, as described above.

Hard and Soft Constraints

Both user and auto-features can be set to either Hard or Soft constraints to change the way they are used by the solver. A hard constraint means that the 2D feature track is assumed to be accurate in all frames. Soft constraint means that the 2D feature track may have errors in certain frames. A feature set to 'Soft Constraint' may have part or all of its feature track ignored by the solver if it thinks the feature track is bad. By default, auto-features are set to 'Soft Constraint', and user-features are set to 'Hard Constraint', although the default state of user-features can be changed from the Calibration tab in the preference window.

Changing an auto-feature to 'Hard Constraint' will have the effect of forcing the camera solver to try and "lock-down" onto that feature throughout the shot. This can be useful, for example, if the feature projection is drifting from the track position, even though the track position is accurate. Similarly, if a certain auto-feature is not appearing at all in the solution, switching it to 'Hard Constraint' can will make the solver concentrate more on the feature and try to introduce it into the overall solution.

You can change the constraint state by right-clicking on a feature in either the shot overview or an image window, and selecting the appropriate menu option. If your camera track is drifting in places, and you can see a particular auto-feature moving out of alignment with its 2D

track, try switching that auto-feature to 'Hard Constraint' and either press the 'Improve Solution' button (or the menu option) or solve the entire camera path from scratch. This will try to ensure the projection of the feature point into each image matches the 2D track as accurately as possible throughout the entire shot.

Creating Feature Constraints

Feature Constraints are a method of restricting the relative position of feature points whilst improving the camera solution. For example, if you have tracked some features on a planar surface, but after the camera solve has completed the features are not lying perfectly on the plane, you can create a Planar constraint which will try to re-position them correctly. Such errors may be caused by noise in the 2D feature tracks, or not enough parallax in the camera motion to get a reliable estimate of 3D position.

To open the Feature Constraints window, shown here on the right, select the *Tracking>Constraints...* option from the main menu. A new constraint can be created by clicking the New button under the top-left window. This will create an unnamed planar feature constraint. To rename the constraint, click again on its name to display an edit box. Right clicking on the feature constraint will display a popup menu allowing you to change the constraint type. Supported constraint types are Linear, Planar, and Distance. None will disable the current constraint, whereas Linear and Planar represent constraints to keep the features on a straight line or a flat plane respectively. A Distance constraint will try to maintain a fixed distance between all features in the constraint.

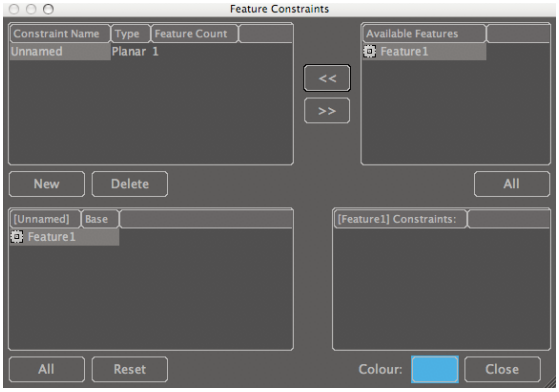

The feature constraints window operates in a similar way as the Groups window. To add a feature to a constraint, highlight the constraint you wish to use in the top-left hand window, then highlight the features you wish to add in the top right-hand window and click the add button '<<'. The feature points will now appear in the bottom left-hand window, while the bottom right-hand window will show the constraints that are applicable for any selected feature point. A number of feature points can be selected at the same time by holding the Shift or Control keys whilst clicking in the top right-hand window.

Note: If a set of features has been highlighted in the main interface window, perhaps selected by the lasso tool for example, then only these features will be listed in the top right-hand window of the feature constraints window. This, combined with the All button below the feature list allows you to quickly create feature constraints for a specific set of features. Individual features can also be added to a constraint by right-clicking on the feature and selecting the Constraints... option from the popup menu.

To remove a particular feature from a feature constraint, select the feature constraint in the top-left hand window, then select the point to be removed in the bottom left-hand window, and click the remove button '>>'.

As well as showing the features contained in a constraint, the window at the bottom-left also allow you to control how that constraint will be enforced. In order to constrain features to a plane for example, PFTrack must first establish exactly which plane to use. Each feature in a constraint can be tagged as being a Basis feature, and only the basis features are used to calculate the parameters of the plane or line, even though the constraint will be applied to all its features. By default all features are basis features, which means that the plane or line to

which they will be constrained will be the best-fit plane or line for all the features. To change the basis state of each feature, right click on the feature name in the bottom left-hand window and select Yes or No from the popup menu.

Different constraints can be assigned different colours to make them easier to distinguish. To display the constraint colours for each feature in an image window, select View>Constraint Assignments from the main menu.

Segmented Camera Paths

If you have problems solving shots with complex motions such as combinations of camera pans and free-motion, PFTrack's camera solver will let you partition a shot into different segments, and specify a different type of camera motion for each segment. This should be done before solving, but only if you are having problems with the standard solver. As the camera is solved, it will try to obey the type of motion you have specified for each segment. Note that segmenting the solver in this way will affect the standard feature-based camera solver, not tracking geometry.

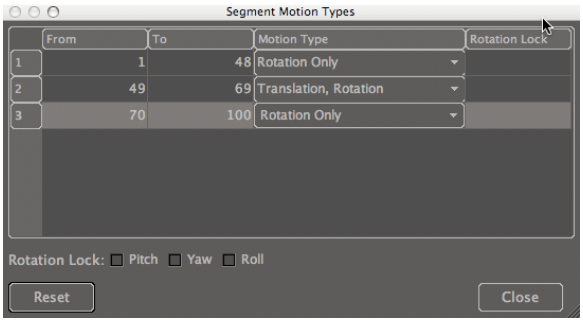

To use the segmented solver, the main Camera Motion must be set to 'Free Motion' in the *Camera>Camera Parameters* window. Open the Dope-Sheet and use the 'CutFrame' button to place cut-frames at the points where the motion of the camera changes. These will be shown as blue indicators in the Dope-Sheet. As an example, assume the camera performs the following motions:

Frames 1-49: Rotation Only
Frames 50-70: Free Motion (translation and rotation)
Frames 71-99: Rotation Only

In this case, you would place 2 cut-frames: one in-between frames 49 and 50, and another between frames 70 and 71 (Note, pressing the CutFrame button will place the cut immediately before the current frame). The cut between frames 49 and 50 signifies that the motion type changes between frames, and similarly for the cut between frames 70 and 71.

To examine the segments, select *Camera>Segment Motion Types...* from the menu and for this example you would see a popup window showing 3 segments: frames 1-49, frames 50-69 and frames 70-99, and alongside is a menu where you can change the motion type for each segment. In this example, you would set 'Rotation Only' for the first and last segments, and 'Translation + Rotation' for the middle segment.

The seven different types of motion that are available are:

Translation + Rotation: this is the same as the main "Free Motion" type;
Translation Only: Camera translation with no rotation;
Linear Translation, Rotation: Straight line camera motion, a dolly shot could produce this motion;
Planar Translation, Rotation: Camera is moving in a flat plane;
Rotation Only: Camera can rotate, but not translate;
Rotation + Small Translation: Useful for shots where the camera is not moving far, relative to the distance from the set;
Interpolate: This will initially ignore the segment during the camera solve, and then interpolate the camera motion after the solve has completed.

To change the motion type for a segment, click in the Motion Type column with the left mouse button and choose an option from the drop-down menu.

In addition to specifying the motion type, you can also lock the camera rotation around any axis within a segment. To do this, highlight the segment you wish to lock, and click the Pitch, Yaw or Roll boxes underneath the segment table. The Rotation Lock column in the table will update to reflect which of the rotation axes are locked (Pitch corresponds to rotation around the X axis of the camera, Yaw corresponds to rotation around the Y (vertical) axis of the camera, and Roll is rotation around the Z axis). For example, if a segment has the camera rotating around its Y axis only, locking Pitch and Roll will stop rotations around the X and Z axes.

Estimating Focal Length

After a camera solve, you may notice that the field of view of the camera looks too big or too small. In these cases, estimating the camera's focal length before performing the camera solve can often help. PFTrack provides two different tools for doing this. This can be particularly useful for tricky shots where there is not enough parallax to get a reliable estimate of the focal length. The focal length estimation tools can also be used to automatically orient the scene during the camera solve, or provide targets to meet when running the "improve solution" process.

Creating a Focal Length Tool

To estimate a focal length using the focal length tool, select the *Camera>Estimate Focal Length* menu option. A wire-frame box and ground-plane will appear in the image window, along with some control buttons and a readout of the current camera field of view and focal length. To estimate focal length, this box must be aligned with a rectangular object in the image (such as the corner of a building or a window-frame) in order to help establish some vanishing points. You can adjust the position of the corners of the box by clicking and dragging with the left mouse button. The vertices will turn blue as they are positioned, indicating that they are "pinned" to the image. Once four or more vertices are pinned, you will see the focal length and field of view change accordingly. You can also use the ground-plane and horizon line to help ensure that the box is positioned accurately. To remove a pin from a vertex, right-click it and select the Remove Pin menu option.

To get the best results, try to position the box in a frame where the rectangular object is not viewed directly from one side: choosing a frame where two-point perspective or three-point of perspective is available will help. The frame where the vertices have been pinned is shown in the Dope-Sheet as a yellow marker. Once the vertices are positioned and the focal length is estimated, opening the *Camera>Camera Parameters* window will show that the camera's focal length has been set to 'Known' if it is constant throughout the shot. For cameras with variable focal length, the tool can be used to get an estimate for the focal length in a single frame.

The 'Lock Size' button can be used to fix the size of the box whilst dragging the vertices. This means that only the focal length and orientation of the camera will be affected. The 'Reset' button will reset the tool back to its default value, and the 'Hide' button can be used to hide the tool from view.

Drawing Vanishing Point Lines

You can also aid PFTrack by estimating the focal length of the camera if the sequence has an easily identifiable orthogonal pair of vanishing points in one or more frames. This would normally be the case if you have buildings or man-made structures within you shot. Orthogonal vanishing points are those that belong to directions in the scene that are at right-angles to each other. For example, the X and Y pair of direction in the scene are orthogonal, as are the X and Z, and the Y and Z pairs.

To mark vanishing points, put the interface into Coordinate Axis Line mode by clicking the Coordinate Axis Line button in the toolbar, shown left. Select which axis you wish to draw from the shot overview by clicking on the 'X Axis', 'Y Axis' or 'Z Axis' items inside the Scene Orientation container. To draw a line in an image window, press and hold the left mouse button to place one end point of the line, and then move the mouse to the other end point before releasing the mouse button. You can change which axis a line represents by right-clicking on an end point and selecting X, Y or Z from the popup menu. By default, X-axis lines are drawn in red, Y-axis lines in green and Z-axis lines in blue. When in Coordinate Axis Line mode, the dope sheet (opened by clicking the Dope-S button) will show the frames where axis lines have been marked.

In order to estimate the focal length for a single frame, you need to draw at least two lines for each of at least two axes within that frame (e.g. two lines for the X axis and two lines for the Z axis). Once this is done, select *Camera>Estimate Focal Length* from the menu to estimate the camera focal length.

A popup window will display the estimate of the focal length that has been calculated.
Answering 'Yes' to the question 'Store this focal length?" will store the focal length for later use by the camera solver. Note that a focal length measured in millimetres or inches is only meaningful if the camera's film-back is the correct size. You can change the film-back size from within the Camera Parameters window. Frames that have a stored focal length estimate are identified in the Dope-Sheet with purple markers. To remove markers from the Dope-Sheet, select a range of frames by clicking and dragging with the left mouse button and press the 'Delete' key or use the Edit>Delete menu option.

Accuracy in placing the coordinate lines is important. To assist with this, hold the 'Shift' key whilst placing points to display a close-up view of the image. The point will be placed upon release of the mouse button, and the zoomed window will close. You can move the end points of the line by clicking and dragging with the left mouse button, holding Shift if you wish to open up the zoom window.

The accuracy of the focal length estimate depends strongly on how well the lines are marked by the user, and on which lines have been selected. For each axis, PFTrack will find the intersection of each set of lines to identify a vanishing point in the image. Because there will always be a small amount of error when positioning these lines, the focal length that is calculated will only be an estimate of the true focal length. The location of the vanishing points determine the camera's focal length. This means that small errors when placing lines that are almost parallel to each other will translate into large variations in the location of the vanishing point, and therefore large variations in the resulting focal length. Small errors in placing non-parallel lines, however, are going to have little relative effect on the position of the vanishing point. This means that you should avoid drawing lines that are very short, or nearly parallel: this is why choosing frames that exhibit two or three-point perspective is important.

Survey Data

To improve PFTrack's camera solving calculations even further, it is possible to enter survey data measured directly from the set. You can open the survey data window, shown below, by choosing the Tracking>Survey Data.. menu option, which is available once features have been tracked. When you open the survey data window, the Available Features list shows all the features in the shot. You will often need to track user-features at the location of your survey points, because the placement of auto-features may not correspond to your survey data.

To enter survey data for a feature, click on a feature in the Available Features list to select it, and then click on the add feature button ('>>') and you will see a new entry in the survey table. Enter the relevant survey location for that feature in the SurveyX, SurveyY and SurveyZ boxes in the right hand table. You need to specify at least 6 survey points before they can be used to solve for the camera motion.

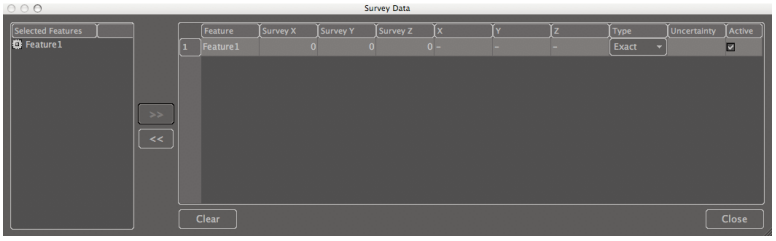

Survey points can be specified as either Exact or Approximate. Exact points are assumed to be measured exactly. Approximate survey points have an associated 'uncertainty', so for example, you can specify that the point is approximately at position (1,2,3) with an uncertainty of 0.01. This would mean the the true survey measurement is somewhere between 0.99 and 1.01 for the x coordinate, and similarly for the y and z coordinates.

If you create a survey point after camera motion has been solved, the survey coordinates will default to the current feature location. This can be used to help solve complex shots, because it allows you to create a set of survey data from one part of the shot and then use that data to solve for the rest of the shot, if necessary, by extending the in/out points and then extending the camera solution.

Reference Frames

Reference frames are a new feature in PFTrack 5.0 that can be used to help estimate 3D feature positions for rotation-only shots. Without reference frames, it is only possible to estimate the direction of each feature, and the 3D point cloud consists only of points on the surface of a sphere that surrounds the centre of camera rotation. By tracking user-features in the main shot and also placing them in reference frames, true 3D feature positions can be estimated.

Reference frames must be loaded before the camera is solved, by selecting the *File>Import Reference Frame* menu option. Note that depending on how your reference frame files are named, it may be necessary to click the Single Frames box in the import window because numbered image sequences will not be displayed by default.

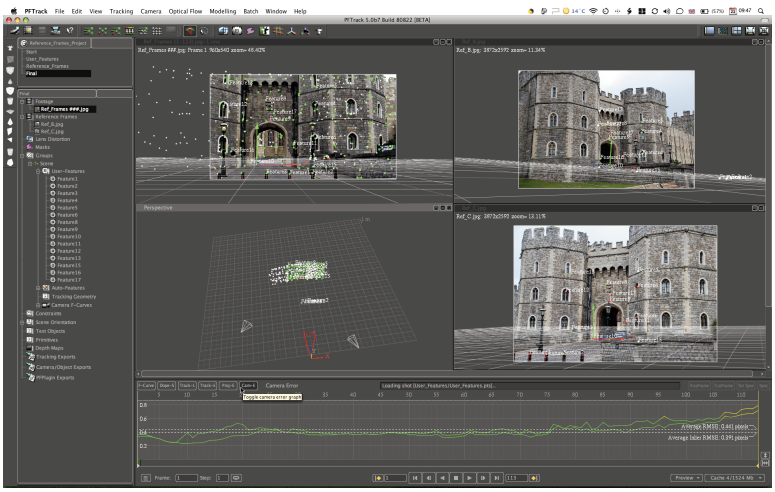

To solve for the position of each reference frame, at least eight user-features must be tracked in the main footage and also placed in the reference frame. At least one user-feature must also

be shared between all reference frames and the main footage. After user-features have been placed, the position of reference frames can be estimated during the camera solve by selecting the Solve Reference Frames option in the camera solver window.

The position of reference frames and the 3D user-feature locations can be examined by opening a 3D viewer window after the camera solve has finished. If the reference frames are not in the correct position, user-features can be edited by switching off the View>Solved Features menu option, or additional user-features can be added where necessary. Reference frames can then be re-solved by right-clicking on their name in the shot overview and selecting a menu option.

Reference frames can also be used for free-motion cameras, although they will not affect the 3D position of user-features until the camera solution is improved using the Camera>Improve Solution menu option (or the Improve Solution button).

When using the image-based modelling tools available in PFTrack, it is also possible to extract texture maps from reference frames. This means that high-resolution reference frames can be loaded to provide more accurate texture maps than are available when extracting textures from the main camera footage.

Scene Orientation

Once the camera motion has been solved you will notice a grid appears in both the footage and 3D viewing windows. This represents the ground plane (Y = 0, or Z = 0 when using Z-up), and its orientation is indicated by X, Y, Z axis directions. Unless survey data has been used, It is unlikely the ground-plane will match the position or orientation of the real ground plane in your shot initially. It is important that you spend some time orientating the ground plane, as this affects the orientation of the point cloud that will be exported to third-party applications. While it is possible to orientate the scene within other 3D systems, it is advisable to orient the data in PFTrack initially. The ground plane is also used when carrying out image based modelling within PFTrack, and poor positioning can make building models more difficult..

Setting The Origin Point

First you need to decide upon a suitable location for your point of origin. This is a point around which you will move the ground-plane and it can be anywhere in your scene. You can place the origin at a solved feature point and it makes sense to select a location which will be of the most practical use for the purposes of your footage, e.g. a spot on an obvious ground plane if your footage has one, which will provide a good reference within the data for export to your 3D package.

While in Navigation Mode, find a well solved feature (a green dot) and move the mouse pointer over this dot/feature. You'll see that the dot expands slightly and information about the feature appears. Now right click on the feature and from the pop-up menu select Set As Origin. You'll see that the ground-plane origin now snaps to this selected point. This is now your point of origin and PFTrack has switched to Scene Orientation Mode as indicated by the Scene Orientation button shown here, and now darkened.

Manual Orientation

When in Scene Orientation mode, you can orient the ground-plane until it is arranged to match the real orientation of your shot. It is often easiest to do this in the footage window, and the best tool for this is your eye. Use the left mouse button to rotate the ground-plane around the origin so it looks correct. If you need to change the origin point to a better posi-

tion you can do this at any time by selecting a new origin point, as described above. Use the mouse to orientate as follows:

Left button = rotate
Middle button/wheel = zoom/scale
Right button = pan

To give you more precise control over adjustments to the ground-plane you can constrain the ground plane's rotation about an individual axis using the left mouse button in conjunction with the 'X', 'Y' and 'Z' modifier keys. You can also use the middle and right mouse buttons to zoom and translate the ground-plane. The *View>Horizon Line* menu option will switch on a virtual horizon line that can help when lining up. The size of the ground-plane can also be adjusted from the Display tab in the preference windows if required.

Alternatively, manipulators can be used to change the orientation of the ground-plane. Whilst in Scene Orientation mode, four buttons appear at the bottom-left of the image window. These can be used to create translate, scale and rotate manipulators at the origin point.

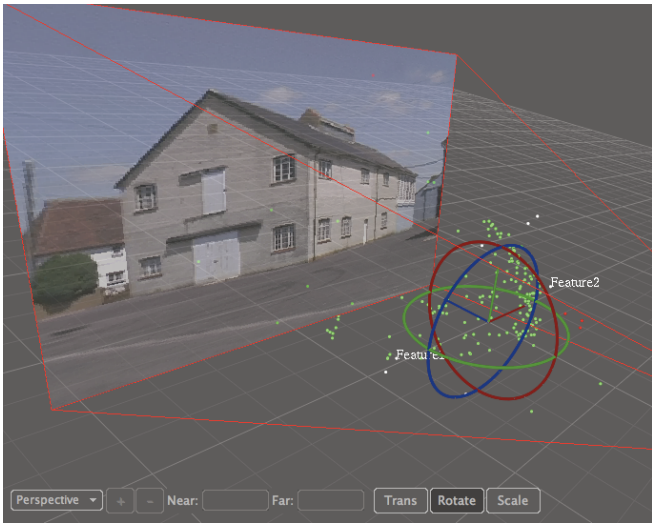

As the position and orientation of the plane is changed in the image window, you will see in the 3D viewer the camera and point cloud move according to your adjustments, while the ground-plane remains stationary. Using new orthographic viewpoints can also help orientate your scene, and the translate, scale and rotate manipulators are also available in these windows.

To switch to an orthographic view, use the drop down menu at the bottom-left of the window to pick a viewpoint orientation. Whilst changing the scene orientation, you may wish to quickly navigate the window viewpoint instead. To do this, hold the Navigation keyboard shortcut (the 'N' key by default) and then release it when you are finished to switch back to Scene Orientation mode.

Once you are satisfied with the position, orientation and scaling of the ground plane and subsequent display of the 3D camera data, you may wish to check your sequence using the playback controls. If you have trouble orienting the ground plane so that it appears locked in all frames of your sequence, it is often because the camera focal length has not been estimated accurately enough. In these situations, fixing the camera's focal length will make it easier to position the ground plane correctly.

Drawing Axis Lines

You can also orient the 3D scene by drawing X, Y and Z coordinate axis lines in a frame, in a similar manner to estimating a focal length (see Drawing Vanishing Point Lines in the section on Estimating Focal Length).

Axis lines can be placed if your sequence has an easily identifiable orthogonal pair of vanishing points in one or more frames. This would normally be the case if you have buildings or man-made structures within you shot. Orthogonal vanishing points are those that belong to directions in the scene that are at right-angles to each other. For example, the X and Y pair of direction in the scene are orthogonal, as are the X and Z, and the Y and Z pairs. By default, X-axis lines are drawn in red, Y-axis lines in green and Z-axis lines in blue. When in Coordinate Axis Line mode, the dope sheet (opened by clicking the Dope-S button) will show the frames where axis lines have been marked.

To draw axis lines, switch into Axis Line mode, either by clicking the Axis Line button in the toolbar. Clicking and dragging with the left mouse button in an image window will create a new axis line. The axis that will be drawn depends on whether the X Axis, Y Axis or Z Axis entry in the shot overview is selected. You can change the axis label for a line after drawing by right-clicking on an end-point and choosing an option from the popup menu.
To fully define the scene orientation, two axis lines must be marked in at least two coordinate directions. For example, you must draw at least two X lines and two Z lines, or two Z lines and two Y lines). After the axis lines have been drawn select *Camera >Orient Scene* from the main menu to orient the scene accordingly.

Orienting To A Plane

The fourth button is used to snap the ground-plane to a group of selected feature points: use the Lasso mode to lasso a set of features that are lying on a flat plane and press the 'Fit' button (or use the *Camera>Orient Scene* menu option) and the ground-plane will be positioned accordingly.

If you have feature positions that lie in a flat plane that you wish to use as the ground plane, then you can automatically orient the scene to best fit the feature positions. To do this, you need to select at least three features by clicking on them or lassoing them in a window. To select features, switch to Selection mode by clicking the Selection mode button in the toolbar. Alternatively, you can hold down the Selection mode shortcut (the 'L' key by default), select your features and then release the key to switch back to orientation mode.

Once three or more features are selected, click the 'Fit' button at the bottom-left of the image window. If the 'Fit' button is not visible, it will be because you are not in Scene Orientation mode. You can also select *Camera>Orient Scene* from the main menu to fit the ground-plane to the feature positions. One of the features will also be selected as the origin point. If you wish to change which feature is the origin after orientation, simply right-click on the feature and choose *Set As Origin* from the popup menu.

Setting The Scale Of The Scene

The overall scale of the scene's coordinate system can be changed in two ways. Firstly, you can switch to Scene Orientation mode by clicking the Scene Orientation button in the toolbar and then using the middle mouse-button to scale the entire 3D point cloud and camera path. This can be used to set a rough overall scale by eye.

Opening a 3D Viewer window will show the ground plane, point cloud and camera path. The size of each ground plane “unit” is shown inside the 3D Viewer window, and defaults to one metre.

If you wish to set an exact scale, open the Scene Scale window from the *Camera>Scale Scene...* menu. This will open a window showing the 3D locations of each feature, allowing you to set a fixed distance between any pair of feature points in the scene. If you select two or more features before opening the Scale Scene window, only those features will be listed in the table.

After highlighting two features in the table by clicking with the left mouse button, the distance between the features is displayed at the bottom of the window. You can edit this distance, or the units it is measured in, and as each edit is made you will see the scale of the scene change in image or 3D Viewer windows.

The Scale Scene window can also be used to set a specific distance between the primary and secondary cameras when tracking stereo shots. This is achieved by editing the value in the Stereo inter-ocular distance edit box.

Transforming The Camera

After solving for camera motion, the camera can be transformed using numerical edit boxes by selecting *Camera>Transform Camera...* from the main menu. The popup window shows the camera position and orientation for the current frame. Changing frames in the timeline will update values in the window accordingly.

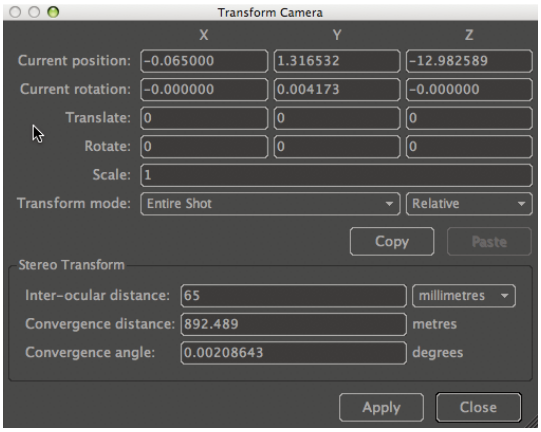

You can enter relative or absolute translation, rotation and scale adjustments in the Translate, Rotate and Scale edit boxes. After typing a value, press Apply to apply the transformation. You can change whether you want to adjust the current frame only, the current segment only, or the entire camera path using the Transform mode box. The Copy and Paste buttons can be used to take a copy of the current camera position and then paste that into a new frame.

For stereo cameras, the Inter-ocular distance, Convergence distance and Convergence angle edit boxes can be used to transform the secondary camera position into a fixed position relative to the primary camera. More details regarding stereo tracking can be found elsewhere in this manual.

The camera can also be transformed when in Camera Orientation mode. This operates in a similar way to scene orientation. Press the Camera Orientation button in the main toolbar, and use the left, middle or right mouse buttons in an image or viewer window to transform the camera around it's current position. By changing the Transform mode in the camera transform window, the adjustments can be made to the current frame, the current segment or the entire shot.

Object Tracking

Object tracking is used in PFTrack to extract the motion of rigid objects in the shot, such as cars driving along a street, in addition to the overall motion of the camera. In order for motion to be solved using object tracking, the object must be large enough to track 6 or more features, and must exhibit enough parallax relative to the camera to be solved using PFTrack's camera solving algorithms.

PFTrack can track as many additional objects in the shot as you require, provided they satisfy the criteria above. For each additional object, you must create a new groups to store the user or auto-features associated with the object's motion. Generally, a mask is employed to separate the moving object from the background when tracking the camera, and vice-versa when tracking the object.

Creating Groups

By default, PFTrack creates a single group called 'Scene' which contains the user and auto-feature tracks that are used to estimate the overall camera motion, relative to a static scene. To create extra groups, open the Groups window by choosing *Tracking>Groups...* from the main menu.

Before creating a new group to track a separate object, you must tell PFTrack how many different motions are present in the scene. If you are tracking one additional object then you will need two motion: one for the camera and one for the object, so enter '2' in the ''Number Of Motions' edit box. After specifying the number of motions, you can create new groups by clicking on the New button under the top-left hand window. This will create a unnamed group, but you can rename them if you wish by clicking on a highlighted group name and entering a new name in the edit box. By default, the group will be assigned to the next available motion number. To change the motion number assigned to a additional group, right-click in the motion column and choose an option from the popup menu.

Tracking Features In Different Groups

Before tracking the object motion, you need to solve for the main camera motion. You can do this in the normal way by tracking features into the 'Camera' group and then solving for the camera motion. To help place auto-features correctly, it is often necessary to build a mask that enclosed your object in all frames, to ensure that auto-features are not placed on the moving object whilst the camera is being tracked.

To track auto-features into the main camera group, highlight the 'Camera' group name in the shot overview and press the Auto-Track button in the toolbar. This will create auto-features in the Camera group, which will be used to solve for camera motion. The group into which auto-features will be placed is shown by the Group drop-down menu in the Tracking Parameters window (opened using the Tracking>Tracking Parameters menu option).

Alternatively, you can also track some user-features if desired. The group into which new user-features is determined by the currently selected group. If no group is selected, the Group choice in the Tracking Parameters window is used. If features are placed into the wrong group, they can be moved easily using tools available in the Groups window (opened using the Tracking>Groups... menu option). It is also possible to use geometry tracking to solve for the main camera motion (see the section on Geometry Tracking for further details).

Solving for the camera motion is done in the usual way, and all the standard controls are available. Be sure to set the camera parameters correctly so you get a good quality camera solve.

Once the main camera motion is obtained, the next stage is to track features into the second group if not already done. If you drew a mask to exclude your object from the camera track during auto-tracking, you can invert the mask by right-clicking on its name in the shot overview and selecting Invert from the popup menu. This will ensure that when you come to track auto-features again, they are placed on your object and not in the background. To track auto-features into the second group, highlight the group name in the shot overview by clicking on it, and press the Auto-Track button. Again, you can track user-features if you wish, or use geometry tracking.

If you object is small, you may wish to reduce the target number of auto-features that are tracked into the group. The Area Limit tracking parameter may also help. Switching this on will automatically limit the number of auto-features that are tracked according to the fraction of the image area that is un-masked.

If you want to clear auto-features from a particular group at any time (for example, if you want to re-track them with different parameters), highlight the group name in the shot overview and select Tracking>Clear Auto-Features from the main menu. A warning box will appear, showing the group that will be cleared and asking you to confirm your choice.

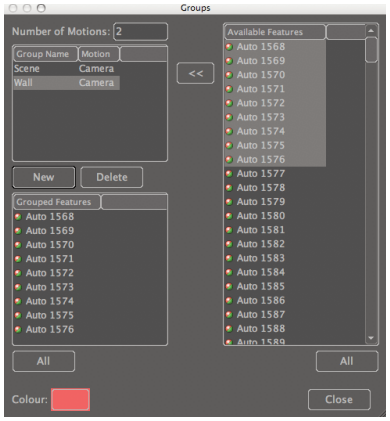

After user or auto-features have been tracked, you may wish to move them into another group. For example, if some auto-features that were tracked into your object group were actually placed on the background, you can move these features to assist with the main camera solve. Right-clicking on a feature and selecting Groups... from the popup menu will display the Group window with the feature shown at the top-right. Select the group you want to add the feature to, and press the Add button ('<<'). The feature will then appear in the Grouped Features list at the bottom-left.

A feature must belong to a group, so if you want to remove a feature from one group you must do this by adding it to another, or by deleting it altogether. When the feature is added to the new group it will automatically be removed from the old group. If one or more features have been selected in an image viewing window then only those features will be displayed in the Group window when it is opened. Opening the group window with nothing selected will display all features in the list.

The feature group assigned to each feature can be easily identified in an image window by the particular colour assigned to the group. This can become quite important if you have several groups in use, each one with a separate purpose. You can switch on the group colour assignments using the View>Group Assignments menu option. This will colour each feature according to the group it is assigned to.

When each group is created, it is assigned a random colour. To change a group colour, make sure the group is selected in the Group window and then click on the colour box at the bottom left-hand corner of the window. This will display a colour chooser that allows you to choose a new colour.

Once you have tracked auto-features for your object, it is important to ensure that you fix any bad tracking points, especially those that are placed in the background. Having too many feature points misplaced on the background will affect the quality of the object motion estimate. If you don't want to move the bad features into the Camera group, just delete them by hovering the mouse over the feature and pressing the 'Delete' key, or lasso them whilst in Selection mode.

Solving For Object Motion

Once features have been tracked into the second group, you can proceed to solve for the motion of the object. To do this, highlight the group name in the shot overview and press the Solve Motion button in the toolbar, or right-click on the group name and select Solve Motion... from the popup menu. Assuming that the camera motion has already been solve using the usual approach, this will display the Solver Controls, and the Motion group drop-down menu should show the name of your second group. If only Camera is displayed in the menu, this will most likely mean that you have not solved for the main camera motion yet.

When solving for additional object motions, the same controls are available for you to use in the Solver Window. See the section on Controlling The Camera Solve for an explanation of these parameters. Pressing the Start button will start the object solve. This process can be repeated for any number of objects within your scene. Make sure to create a separate group for each object, and repeat the steps we have described here as appropriate.

If object solving fails at any point, you can delete the partial motion track by highlighting the group name in the shot overview and selecting Camera>Delete Solution from the main menu. Alternatively, right-click on the group name in the shot overview and select Unsolve Motion from the popup menu.

Once object tracking is complete, object and camera motion can be exported to additional applications. Export of object motion is supported for 3ds Max, FBX, Flame, Lightwave, Maya, Softimage XSI and as XML data..

Adjusting Object Sizes

Once an object is tracked, you will see a wire-frame box appear around the object in the image and viewer windows. This illustrates the orientation of the coordinate system that is used to store the object's 3D feature points. You can adjust the orientation of this box and its origin point by highlighting the group name in the shot overview and pressing the Scene Orientation toolbar button.

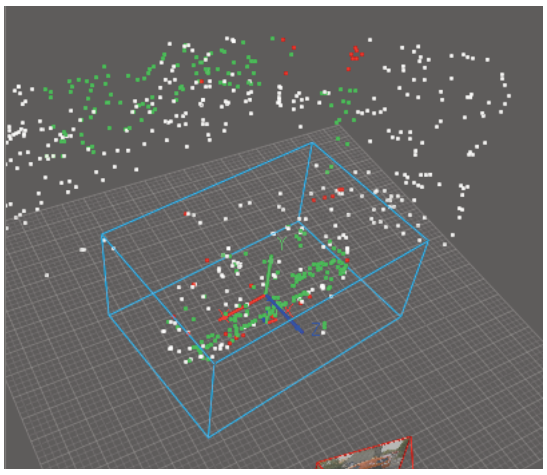

The usual translate, rotate and scale manipulators are available and you can also use the left, middle and right mouse buttons to rotate, scale and translate the object's coordinate system. To move the origin to a specific feature, right-click on the feature and select Set As Origin from the popup menu.

It is also important to set the scale of the object group, which in turn affects its distance from the camera. PFTrack is not able to automatically estimate this scale because it cannot be determined visually from the feature data alone.

When scaling the object group, it will help to have a 3D viewer window open to check the effect of any adjustments, because they will not be visible in an image window when viewed through the camera. If you wish to scale the object so that the distance between two points is a specific value, select those points and open the Camera>Scale Scene window to edit the distance..

Geometry Tracking

In some cases, trying to solve for camera motion using standard 2D features will not give the quality of results that you require, especially if the object is very small or it is hard to place specific features accurately. In these situations, using geometry to obtain motion data will allow more accurate results to be obtained very quickly. The idea here is to import a geometric mesh that you can position accurately in one or more keyframes, and have PFTrack estimate the motion of the object throughout the rest of the shot.

Geometry tracking can be used to estimate camera and object motion. You don't have to have an exact model of your scene or object: if you want to track camera motion relative to the side or corner of a building, you can simply import a box-shaped object, place it in the correct position in one frame and then track the camera motion from there. If you wish to track a flat planar surface, the easiest approach is to simply load a square polygon and track the camera motion relative to the surface using these geometry tracking tools.

Tracking a digital set model into a background plate can provide accurate results very quickly if a surveyed geometric model is available. In more complex examples, geometry tracking is also an ideal tool for head replacement. In this case, simple cube or sphere shapes may not be sufficient to get a good quality track. Having an accurate head model will improve the solution because the geometry is a better representation of the underlying image.

Importing And Viewing Geometry

To load geometry for tracking, select *File>Import Tracking Geometry..* from the menu. Navigate to the Wavefront OBJ model you want to use and click the 'Open' button. PFTrack provides a small set of OBJ models that you can test with if you wish, and they are located in the 'Primitives' folder where PFTrack was installed. When exporting geometry from a 3D application to use for geometry tracking, remember that PFTrack does not currently support textured models or application native formats other than Wavefront OBJ. The complexity of the geometric mesh will also affect tracking speed: the more triangles you have the slower the mesh will track, so try to get a good balance between mesh complexity and accuracy when modelling objects for geometry tracking.

When loaded, the imported model will be placed at the centre of the viewing window, and a set of tracking controls will appear at the bottom of the window. Right-clicking on the geometry name in the shot overview will display a popup menu that can be used to change the coordinate system of the objects, as well as show/hide back-facing polygons. To delete a piece of tracking geometry, highlight its name in the shot overview by clicking with the left button and press the 'Delete' key (or choose *Edit>Delete* from the main menu).

It is important to ensure that the camera's focal length is correct before tracking geometry. You can change the focal length directly in the *Camera>Camera Parameters* window, or by using the up and down arrows in the FOV edit box to alter the Field Of View (measured in degrees). If you are finding it hard to position your model try altering the FOV using the up and down arrows until you find a better match. The focal length can also be estimated whilst positioning geometry in a frame, if necessary, This is achieved by using the Drag manipulation mode, described below.

You can switch the way the model is displayed by using the appearance drop down menu (the default mode is Hidden Line). You can select a number of display modes for the imported geometry to suit your working method. Hidden line is often the most useful when placing an object over footage, especially if you are using the 'Drag' option to position a vertex. The other modes, Faceted, Smooth and Transparent can be switched using the object view drop down menu and are shown to the left.

Note: This release of PFTrack does not support the import animated geometric elements. Also, if you want to track a rigid object that is composed of several individual pieces, you need to import all pieces as a single piece of geometry.

Positioning Geometry In One Frame

Before starting to track the geometry through the sequence, try to find a frame where the geometry can easily be aligned to the image. This should be a frame where you can see as much the object as possible, and the object is as large as possible in the image.

The way that you interact with the model is determined by the manipulation drop-down menu. By default, this is set to Fly Camera and when in this mode you can use the left, middle and right mouse buttons to rotate, scale and translate the model until it matches up with the image data in one frame of the footage you want to track. As you adjust the object position, you will notice that a light-blue key-frame is created in the Dope Sheet.

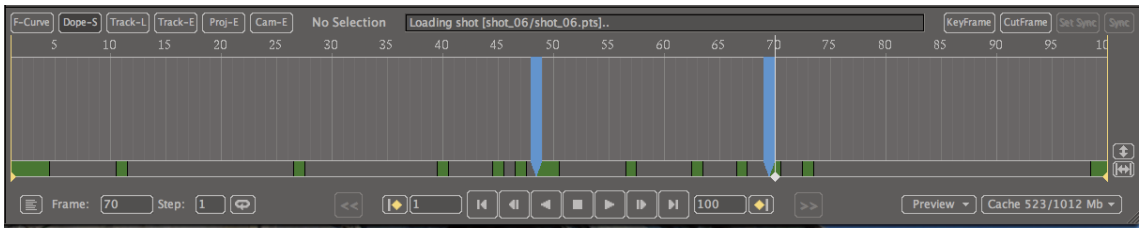

Additional manipulation modes are available that will let you use standard Translate, Scale and Rotate manipulators to position your object, if you wish. It is worth remembering that if you are tracking the camera relative to a static object, these manipulations are really adjusting the camera position, not the object position. When in Translate, Rotate or Scale modes, you can use the Centre button 'C' to re-position the manipulator widget if it moves off screen.

If you want to set the position of a model very accurately in one frame before tracking, as would be the case with a head or face replacement, you can enable the Pin manipulation mode from the drop-down menu and pin vertices of your model to specific pixels in the image. This process will try to lock your model position so the vertex matches the pixel you have attached it to. As you hover the mouse over each vertex, you will notice a red circle drawn around it. Click and drag the vertex into position in the image. If you want to drag a vertex without pinning it, hold the 'Shift' key before clicking on the vertex.

Repeating this for a number of vertices will transform the camera position to try and match all vertices you have pinned. You can delete a pin at any time by right-clicking on the vertex and selecting Remove Pin from the popup menu. When in Pin mode, the 'T' 'R' and 'F' buttons are available to use, and clicking these will affect the type of transformation that is used to match the geometry to the pins. By default, 'T' and 'R' and active, which means the camera will be allowed to translate and rotate. Clicking the 'F' button will also allow the camera's focal length to be changed during the transformation. This can be useful when you do not know the focal length before-hand, and can greatly assist in getting a good quality geometry track.

Tracking Through The Shot

After the object is positioned correctly in one frame of the shot you can automatically track it into other frames. To do this, press the Track Forward (or Backwards) buttons ('>>' or '<<') along the bottom of the image window. You can stop tracking at any time by pressing the 'Escape' key. Once completed, the footage will return to the frame you started tracking from. Unless you positioned your object in the first frame, you will need to track both forwards and backwards to complete the estimate of camera motion.

If the geometry starts to drift from the image underneath, which may happen if you have very fast camera or object motion, the track can be stopped by pressing the 'Escape' key. It is then possible to scrub through the time-line to the point where the drift starts and reposition the model again using whatever manipulation mode you desire. Doing so will create another keyframe in the time-line. This can be repeated as needed. Pressing a track button at any time will then re-track between key-frames, blending the motion of the object appropriately. For example, you may track geometry forwards from the first frame to the last, and then re-position the geometry in the last frame because it has drifted slightly. You would then track backwards from the last frame to the first, and both transformations will be blended together accordingly using linear interpolation. You can disable this interpolation by switching off the interpolation ('I') button at the bottom of the image window.

If you have a very difficult sequence to track you can track geometry forwards or backwards one frame at a time using the '<' and '>' buttons at the bottom of the image window, allowing you to fine tune your model position before tracking into the next frame. In this way even the most difficult shot can be tracked relatively quickly.

To remove keyframes from the Dope Sheet, click and drag with the left mouse button to select a range of frames and press the 'Delete' key (or choose the Edit>Delete menu option). All adjustments to tracking geometry positions can be un-done at any point by pressing the undo shortcut, 'Ctrl+Z', or by selecting Edit>Undo from the main menu.

When tracking geometry, Z-jitter weighting (ZW) can be used to reduce any jittering that occurs along the view direction of the camera. This may occur if the object is small and the tracker is unable to accurately estimate its distance from the camera. Increasing the ZW value (up to a maximum of 100) will reduce amount of jitter only in the view direction without affecting the other estimates of object position.

Adjusting Vertex Weights

The Weights display mode can be used to show the weight assigned to each vertex. Weights can be used to ensure that some parts of the geometry are ignored during tracking, by reducing the influence image motion has over certain areas of the geometry. For example, if you are tracking a talking face, you may need to reduce the weights around the mouth if its movement is affecting the quality of the geometry track. When in Weights display mode, click the '-W' or '+W' buttons and paint over vertices with the left mouse button to decrease or increase weights. Paint strokes can be undone by pressing the 'Ctrl+Z' keyboard shortcut or selecting Edit>Undo from the main menu.

Holding the 'Shift' key whilst dragging with the left button held will change the size of the brush. Weights can be reset to their original values by right-clicking on the geometry name and selecting Reset Weights from the popup menu, and the colours of high/low weights can

be changed from the preferences window.
Weights can be painted in either an image window, or a 3D viewer window to adjust those hard-to-reach places.

Tracking Objects using Geometry

By default, PFTrack will produce a static object and moving camera when geometry is tracked into the 'Camera' group. When doing head replacement or object tracking you may want a moving object and static or moving camera. The quickest way to do this is to first solve for the main camera motion (or set the camera motion to Stationary in the Camera Parameters window if you have a static camera) and then create a new motion group for the additional object, as described in the section on Creating Groups, elsewhere in this manual. Note that you can use either features or tracking geometry to solve for the camera motion.

To import tracking geometry into the second group, select the group the shot overview by clicking on it's name and then import your geometry in the usual way (File>Import Tracking Geometry..). The tracking geometry will appear inside the new motion group. You can then follow the same process of placing the geometry in one and track it throughout your footage. Once complete, you will have the camera motion in the 'Camera' motion group and your object motion in the second group. You can choose to export both or either of these groups when you export data to your 3D application.

Stereo Cameras

A major new feature in PFTrack 5.0 is the ability to track stereo camera rigs. A stereo camera rig has two lenses (denoted primary and secondary) that simulate human binocular vision, and when viewed through special glasses or on an auto-stereoscopic display, provide a three-dimensional view of the scene.

Typically, a stereo camera rig is constructed so the camera lenses are locked at a specific distance apart, and are either pointing in exactly the same direction or angled slightly towards each other. It is often sufficient simply to track one of these cameras, and the other can be positioned automatically using a fixed transformation within your 3d animation package, but in cases where the offset and convergence angle of the cameras are unknown, this may not always provide the most accurate results.

Obviously, it is possible to track the camera for each eye separately, but PFTrack allows both eyes to be tracked at the same time, thereby placing both virtual cameras together in the same scene, and allowing information from both lenses to contribute towards the overall solution. Once tracked, both virtual cameras can be exported to a variety of 3rd-party packages. Note: Feature groups and tracking geometry cannot be used to track stereo cameras.

Preparing Stereo Footage

To enable stereo tracking in PFTrack, load footage for the primary camera and open the Camera Parameters window using the Camera>Camera Parameters menu option. Changing the Camera Type menu to Stereo Camera will enable the Stereo Rig drop-down menu, where you can specify the type of camera rig that has been used to shoot the footage. The default option is Constant Convergence, which means that the primary and secondary cameras are angled towards each other at a fixed angle throughout the shot. If the cameras are both pointing in the same direction (i.e. The convergence angle between them is zero) then this can be changed to Parallel. Finally, if the convergence angle between the cameras is non-zero, but changing throughout the shot, this menu can be set to Variable Convergence.
After setting any other necessary camera parameters and closing the window, the import footage button will be available to load the footage shot by the secondary camera. Note that

the secondary footage must be the same resolution and frame-rate as the primary footage. Make sure you specify the correct camera preset for the secondary footage as well as the primary.

When using stereo camera rigs, the 'Rotation Only' setting for camera motion is taken to mean rotation around the point exactly half-way between each camera's centre-of-projection. For monoscopic rotation-only cameras, only feature directions can be estimated, resulting in 3D feature points scattered on the surface of a sphere centred at the camera. For stereoscopic rotation-only cameras, true 3D feature positions can be estimated, so full scene structure can be recovered in the same way as it can for free-motion cameras.

For stereo tracking, it is generally recommended that you open two windows, each containing one piece of footage. You can do this quickly by selecting Window>Layout>2-View (2xImage) from the main menu bar. This will open two image windows: the left containing the primary footage, and the secondary footage on the right.

If the two pieces of footage are not synchronized, the tools described in the section on Motion Capture can be used to sync the footage together. After making sure the Sync button at the top-right of the timeline is pressed, scrubbing through frames in the timeline will allow both pieces of footage to be viewed together.

Tracking In Stereo

Tracking stereo footage is done in a very similar way to monoscopic footage. For a feature to be included in the camera solve, it must appear in both the primary and secondary footage. Simply clicking the Auto-Track button will identify auto-features in the primary footage, and transfer them automatically into the secondary footage during tracking. When using masks to restrict where auto-features are placed, the mask must be drawn over the primary footage. As features are tracked, if they cannot be placed in the secondary footage they will be removed. This ensures that every feature is placed in both the primary and secondary footage.
User-features can also be tracked into stereo shots. When in User-Feature mode, a Stereo button appears at the bottom-right of each image window. By default, this button is pressed which means that any user-feature adjustment made in one window is automatically transferred into the second. Because of this, tracking a user-feature in a stereo shot can often be done as quickly as in a monoscopic shot.

In certain situations, however, the feature may not be transferred correctly, especially if the inter-ocular distance (the distance between the primary and secondary cameras) is large. In this case, switching the Stereo button off will allow user-features to be adjusted independently in each window. At least six user-features need to be tracked in both the primary and secondary windows for the camera to solve, but as with monoscopic shots, adding more will often increase the accuracy of the solve.

Solving Stereo Camera Motion

Once features have been tracked, clicking the Solve Motion button will open the solver window. There are two additional solver controls that come into play when tracking stereo shots. The first is Enforce Stereo Constraints, which controls whether the inter-ocular distance and convergence constraints are enforced during the camera solve. The inter-ocular constraint restricts the primary and secondary cameras to be separated by a fixed distance, and the convergence constraint enforces the cameras to be angled towards each other according to the Stereo Rig option in the Camera Parameters window (Parallel, Constant Convergence or Variable Convergence).

In certain situations, such as when the cameras that shot the primary and secondary footage were not placed in this configuration, or have slightly different focal lengths, attempting to enforce the stereo constraints during the camera solve may cause the solve to fail. In these cases, switching the Enforce Stereo Constraints option off will often allow a camera solve to be obtained, even though the position of the secondary camera relative to the first may not be consistent throughout the shot.

The second solver option that will affect the stereo camera solve is Separate Stereo Solve. By default this option is switched off, which means that the secondary camera position will be estimated at the same time as the primary position, using information from features in both camera windows. Provided your footage has been shot with a properly configured stereo rig, this will generally produce the most accurate solution, especially when the Enforce Stereo Constraints option is switched on. However, if the stereo rig has not been configured accurately, switching Separate Stereo Solve on will force the camera solver to produce a solution for the primary camera first, and then estimate the position of the secondary camera as an additional step. In situations where the unified solver fails, this can often be used to produce a valid solution.

Once the camera solver has finished, the solution can be examined in the usual ways, typically with a 3D Viewer window showing the position of the primary and secondary cameras. In addition to the camera frustums, two yellow dotted lines are drawn from the camera centres to illustrate the convergence distance. The convergence point can be switched off using the View>Stereo Convergence menu option.

Stereo constraints can also be enforced when improving the solution. Click the Improve Solution button or selecting *Camera>Improve Solution* from the menu bar to open a popup window. By ticking the Stereo box, PFTrack will attempt to enforce a fixed inter-ocular distance, and camera convergence. The relative weighting of these constraints can be adjusted using the slider bar.

Adjusting The Secondary Camera

The position of the secondary camera can be set exactly relative to the first by using the *Camera>Transform Camera* menu option. The inter-ocular distance, convergence distance or convergence angle edit boxes show the values for the current frame. Editing these values will transform the secondary camera accordingly. Note that this transformation assumes a properly configured stereo rig, with the cameras angled towards each other using only a rotation about their vertical axis. Positioning the secondary camera in this way for rigs without this configuration may reduce the matching accuracy between the feature tracks and projections in the secondary footage.

Often, the inter-ocular distance between the primary and secondary cameras is known: typically it is set to around 65mm when constructing the stereo rig. In these cases, the scale of the camera and feature points can be set within PFTrack by scaling the scene to meet this criterion. This is done using the Camera>Scale Scene menu option, which displays a popup window containing the current inter-ocular distance. Changing this value will scale the entire scene accordingly.

Rendering a Stereo Anaglyph

It is possible to render out a red/cyan stereo anaglyph for viewing through tinted 3D glasses by right-clicking on the footage name in the shot overview and selecting Render Anaglyph Movie from the popup menu. After selecting a suitable file location, a movie will be rendered containing the footage and object overlays, as seen through each camera view. The left-hand camera will be rendered in the red channel, and the right-hand camera in the green/blue channels. Preference options are available to adjust the colours used for the left and right eyes, and these can be adjusted according to the colours of the filters in your 3D glasses.

Test Objects

It may be the case that it is hard to judge the quality of the camera solve just by looking at the placement of feature points in the image window. In certain situations there are actually multiple combinations of feature position and camera motion that can satisfy the constraints provided by 2D feature tracks, especially when the camera focal length is unknown or varying.

To remedy this, 3D test objects can be imported and placed in free space or locked onto user/auto-features to test the quality of the camera solve.

Adding Test Objects

To add a test object to the scene click on the New Test Object button in the toolbar, or right click on a user/auto-features when in Navigation, Selection or Scene Orientation Mode and select New Test Object from the popup menu.

A window will be displayed that show a list of possible objects. Click on the object you wish to add and you will see the object appear in the preview area. Press the Create button to add the selected object to the scene.

The Browse button allows other Wavefront OBJ models to be imported into PFTrack to be used as test objects.

Transforming Test Objects

To translate, scale or rotate a test object, select the required object from the Shot Overview (if not already selected) and click the Trans, Rotate or Scale buttons that appear at the bottom-left of the image window to create standard manipulators for adjusting the object's position.

You can snap its position to another feature point by selecting the object in the Shot Overview and then right-clicking on a feature and choosing Snap Selected Object from the popup menu.

Trouble-Shooting A Camera Solve

There are several different approaches you can try to fix a bad camera solve. The most important tools for examining the overall solution are the 3D viewer windows and the Cam-E and Proj-E graph windows, and both should be used to check the quality of your camera solution. In particular, make sure that the 3D structure of your scene and the camera path look okay in the 3D perspective and orthographic views.

Examining The Solution

When the camera solve is completed, the feature tracks will have turned into small dots in the image window. The green dots should make up the majority as these represent those 3D features which match their track position well. The red dots show a poor 3D point, which deviates from its track position. You can see exactly how poor by switching on the *View>Projection Errors* menu option. This will draw a small line from the 3D feature point to the corresponding 2D track point. You can also get a numerical value of the error by hovering the mouse pointer over a feature until a popup menu appears. White dots indicate where a feature wasn't tracked into that frame but is visible after the camera solve.

Remember that a large projection error can be caused by several factors:

1. The 2D feature track is incorrect
2. The 3D feature position is incorrect
3. The camera position or focal length is incorrect

Clicking the 'Cam-E' button above the timeline will display a graph showing the overall RMS (Root-Mean-Squared) error for the solution. This is a measure of how well the 3D feature points match their 2D track positions when viewed by the camera. Two graphs are displayed: the first shows the RMS error for all feature points in each frame, and the second shows the Inlier RMS error, which is the error for only the 'inliers' in each frame. An 'inlier' is a point which projects below the threshold used to distinguish good points from bad: inliers are the green points mentioned above, and 'outliers' are the remaining red points (the inlier threshold can be changed in the Calibration tab of the Preferences window).

You can also click on the 'Proj-E' button to display the error for each feature point individually. Hovering the mouse over a graph will display the feature name as well as the error value. You can select features directly from the graph by drawing a selection rectangle with the left mouse button, providing a quick way of identifying which features are not matching with their 2D track positions.

Fixing Bad Feature Tracks

The Cam-E and Proj-E graph buttons can be used to quickly identify problems with your feature tracks and camera solve, if the situation arises. For example, the difference between the RMS Error graph and the Inlier RMS Error graph can be used as an indicator of which frames have a significant amount of outliers. In these frames, features with a large projection error can be identified in the Proj-E graph and appropriate action taken.

If the feature is not tracked correctly, it should be re-tracked or removed entirely. PFTrack is able to automatically reject a number of badly tracked auto-features (assuming they have not been set to 'Hard Constraints'), but too many bad tracks can seriously affect the quality of the camera solution. Generally, if you are having problems tracking a shot, the first thing to do is make sure that you do not have many badly tracked features.

Remember that badly tracked features are not only those that have parts of their feature track coloured red: For example, a well tracked feature may have been placed on a moving object or at the junction between two edges of different depths. In these situations, the feature may still tracked well but the estimate of its 3D position will not be correct, so take care and examine your feature tracks if you run into problems during the camera solve.

On the other hand, if the feature track looks good but the projection error is large, it may be because the camera solver was not able to 'lock down' sufficiently on the track. This can happen when there are large numbers of feature tracks elsewhere in the image, which overwhelm

the camera solve and make it focus more on other parts of the scene. You can force the solver to lock down onto a specific user or auto-feature by changing the feature to 'Hard Constraint' from within the feature's popup menu. Features that are 'Hard Constraints' must have accurate feature tracks, and the solver will never ignore them in favour of other features in the shot.

Improving The Overall Solution

Once you have edited the feature tracks to ensure that that are of sufficient quality, you can get PFTrack to update the current solution to try and better match the feature data. This can also be done after editing camera parameters, such as the focal length. For example, you might find that the camera's focal length is initially too small. By editing increasing its value in the Camera Parameters window and then improving the solution, you can get PFTrack to try and re-position the cameras and 3D feature points so that the focal length better matches the value you have entered.

To improve the overall solution, select the *Camera>Improve Solution* option from the main menu, or click the Improve Solution button in the toolbar. This will display the Improve Solution window, which is described in more detail in the Windows section of this manual.

When improving the overall solution, you can also edit the camera F-Curves if you wish, to iron out any kinks or wrinkles in the camera path. This is described in more detail in the section on Viewing And Editing Function Curves.

Solving The Camera Motion In Stages

In some rare situations, it may be difficult to get a good quality solution to work from in the first place. In these cases, the Improve Solution window is unlikely to improve the quality of the solution sufficiently and a different approach must be taken, which involves solving for the overall camera motion in different stages.

When tracking a Free Motion camera, the first stage is to get a good initial solution on which to build. The initial solution in this case is a camera path for one part of the shot, covering only the frames between the 3 keyframes that are used to control the camera solve. If a camera track fails, or does not give an accurate result, the first thing to do is check the position of these keyframes and the quality of the initial solution.

After the solve has stopped, open the Dope-Sheet and you will see 3 orange keyframe markers. These indicate the positions of the 3 keyframes that were used to build the initial solution.
To check the quality of the initial solution only, press the Solve Camera button again, and tick the 'Initial Solution Only' box. Now, solve the shot again by pressing the 'Start' button. You will see that rather than solving the entire camera path, PFTrack stops early and only solves the frames between the first and last keyframe. Open up a 3D viewer window and take a look at the solution - does it look okay?

The initial solution needs to be created over frames of the shot where the camera undergoes some translation, and the feature points have exhibit significant parallax. If this is not the case, the initial solution will be poor and the first thing to try is to move the initial keyframes. You can do this quickly by placing a single keyframe and then running the camera solve again. Before doing this, you need to remove the current keyframe positions by selecting the frames in the timeline and pressing the "Delete" key. Now, find a region of the shot where the camera undergoes significant translation and there is significant parallax in the feature tracks and press the 'KeyFrame' button at the top-right of the timeline (If this button is not available, try opening the Dope-Sheet). You will see a single orange indicator appear in

the timeline. Now, press the 'Solve Camera' toolbar button again, but change the KeyFrame Method to Manual before starting the camera solve. This ensures that PFTrack uses the keyframes you have just placed, instead of estimating their positions.

After the initial solution has been built again, look in the dope-sheet and you will see that 2 additional keyframes have been created on either side of the one you positioned earlier. Sometimes, these additional keyframes will be too close, so if the initial solution still isn't working, delete them and then create new ones a bit further away. Note that there needs to be at least 6 features shared between all 3 initial keyframes, so don't move the keyframes too far away from each other.

If you can't find any suitable keyframe positions, or if the perspective in the initial solution seems wrong, try to estimate a focal length for one of the frames, or enter a known or approximately known focal length in the Camera Parameters window and then solve for the initial solution again to see if this improves things.
Once you've got a good initial solution, the next stage it to extend that so the entire shot is successfully tracked. To do this, make sure that the 'Extend' option is selected in the Solver Controls window. This will extend the existing solution into the other frames, without re-generating it from scratch.

If the solve still fails at a particular frame, this may be due to poor quality feature tracks or extreme camera motion. The first thing to try is to switch the solver into 'Single Frame Step' mode by selecting that option in the Solver Controls window. This will run a more accurate (but slower) solve that tries to update the solution one frame at a time. If this still fails, take a look at the frame where the solve stops and see if you can improve the feature tracks around that area.

Other Trouble-Shooting Tips

One useful trick for extending an initial solution out to the full shot is to use survey data. After the initial solution has been generated, select 6 or more solved features and generate some new survey points by opening the Survey points window (*Tracking>Survey Data..*). You will see that the survey locations default to the current solved positions. Once the survey points are created, delete the current solution and solve the shot again using the survey data.

This is also useful in situations where you've positioned the ground plane nicely for the initial solution, and don't want to re-position it again after the full solve.

Exporting Data

In order to use the camera track and feature positions in third-party applications, the data must be exported from PFTrack. PFTrack has several broad categories of exports: those that deal with 2 dimensional data (saving 2D tracking data which could be used, for example, for corner pinning in a compositing application); those that deal with 3 dimensional data (such as the results of the camera solve, i.e. 3D feature positions, camera path and geometry) which can be used with any number of animation packages; and finally exports to Pixel Farm plug-ins, and the export of depth map or optical flow data.

The first three types of export are grouped together at the bottom of the shot overview window. To create an export, right-click on the appropriate item and select New... from the popup menu. A window will open where you can choose the desired export format, along with the items of the scene you wish to export, and the path and name of the exported file. You can have as many exports as required for each shot. To modify an existing export, right-click on its entry in the shot overview and select Edit... from the popup menu. To remove an export, select Remove from the popup menu.

By default, each export will be saved in the Exports folder for the shot, but you can change the export location by clicking the Browse button inside the Export window and browsing to a new location. As each export is created the data is automatically saved to disk, although it is also possible to manually export at any time by right-clicking on a particular export entry or the export folder in the shot overview and selecting Export or Export All from the popup menus.

Export Options

When creating an export, the Export window will be displayed that contains a number of options that you can use to specify what exactly to export. The Format menu is used to set the format for the export. Note that the format you choose will determine which additional export options are available.

The Cameras menu controls which camera will be exported. This will be disabled unless you are exporting a stereo camera track, or motion capture data. When available, this menu can be used to select between the Primary or Secondary cameras. Some export formats also support multiple cameras inside a single file. In this case, selecting All from the menu will save all cameras together.

The Features menu specifies which features will be exported. As well as exporting All or None the tracking points, it is possible to export a subset of the user or auto-features that have been specifically tagged using the Tagged option. This is very useful when you only need a few reference points in the export data, but needed to track a large number in order to get a good quality camera solve. To tag a particular feature for export, right-click on the feature in either the shot overview or an image/viewer window, and select Tag For Export from the popup menu. Tagged user and auto-features are marked as such in the shot overview.

The Tracks menu is only available for exports where 2D and 3D data can be saved in the same file (for example, the XML export). Setting this option to Yes will save the 2D track position for every feature in every frame. Because this can generate a lot of ascii text, you may want to switch the Tracks option off unless to require 2D tracking data as well as 3D camera motion.

The Tracking Geometry option is used to export tracking geometry as well as 3D camera data, assuming that you have imported geometry in the shot and the export format you have chosen supports it.

If you have carried out Z-Depth estimation for your scene, you can use the Depth Mesh menu to select whether to export depth meshes from All Frames or just the Current Frame. Depth mesh export can also be switched off by choosing the None menu option.

Modelling primitives and texture maps can be exported to applications that support them by choosing Yes from the Modelling Primitives menu, and selecting an image format for the texture map files.

The Scale edit box can be used to scale the entire scene independently for each export. By default, this scale is set to 1.0, although for export to Flame/Inferno, Cinema 4D, Motion and AfterEffects, the scale defaults to 100. Note that adjusting this value will only scale the export: the scene inside PFTrack will remain at its current scale.

Finally, if you have tracked multiple object motions using different groups, you can choose which groups to export by ticking or un-ticking the box next to each group name in the Export Groups list.

Note: The Shake Corner Pin export is special in that you can only export four tracking points. Trying to export any other number will fail. The easiest way to do this is by tagging the four features you want to form the corner pin using the Tag For Export menu option, and then change the Features export option to Tagged to ensure that only those four points are saved.

Setting Default Exports

A set of default exports can be defined in the Exports tab in the Preferences window. Any exports defined here are automatically added to a new shot. This feature is very useful when working to a fixed workflow with PFTrack because you do not have to create each export by hand for every shot you track.

Part 04 - Beyond Tracking

Optical Flow Analysis

What is Optical Flow?

The 2D tracking tools in PFTrack allow you to track individual points throughout a shot. Optical Flow describes the 2D motion of every pixel in your footage. PFTrack contains a comprehensive suite of advanced optical flow tools that work alongside the already powerful tracking toolset.

Optical flow tracks all pixels from one frame to the next, building a comprehensive model of all pixel movement. A brush-based set of editing tools also allows you to modify and enhance any anomalies in the flow data, producing an accurate flow field prior to export.

Once PFTrack has computed the optical flow data and you have modified any part of it you wish, you are left with an invaluable resource. The data can be used in many parts off the post production or CGI pipeline, such as driving particle systems, creating optical effects, driving expressions and dynamics etc., the possibilities are only limited by your imagination. The Pixel Farm have developed a number of plug-ins for third-party applications that provide immediate access to the optical flow data, such as re-timing and per-pixel motion blur (PFRetime and PFBlur). Alternatively, the data can be passed to other third-party 3D animation and compositing applications by exporting as a set of standard RGB images. As with all Pixel Farm products, the unified analysis of a single shot provides obvious benefits of absolute consistency through a post job, and the real world time and cost saving benefits of 'analyse once, use many times' can be dramatic.

Calculating Optical Flow

The optical flow tools calculate a flow vector for every pixel, and can use feature tracks and masks as 'hints' to assist the calculations. In most cases, these extra hints are not required, and you can display the optical flow control window by selecting *Optical Flow>Parameters...* from the main menu or by pressing the Optical Flow button in the toolbar.

This window contains options to control the quality of the optical flow as well as the range of frames you wish to calculate.

The Block Size field at the top of the window refers to the size of the area (or window) around each pixel to be used as the sampling area when calculating motion. Smaller block sizes can produce a limited degree of improved accuracy but at the cost of extended computation times.

The quality of the optical flow algorithm can be changes between Standard and High. The Standard quality algorithm will assign a single flow vector to each block of pixels that is found. The High quality algorithm will try to assign a single flow vector to each pixel, if necessary, in order to improve the accuracy of the flow field. This will take more time to calculate, but will produce more accurate results.

The Source Frames options allow you to control which frames you want to calculate optical flow for. The default option is All, which will calculate optical flow for all the frames between the in and out point in the timeline. Changing the menu to Current Frame will compute flow for the current frame only, and From/To Frame will perform the calculation between specific frames. Note that when computing optical flow for a single frame, any flow fields that are

affected by the image data in the frame will be calculated. This means that the flow fields to/from the previous and next frames will also be generated.

These options can also be used when re-computing the flow field as well, so after the initial flow has been calculated you can add/remove/edit the feature tracks or masks in any problem areas and then re-generate part of the flow field to take account of the changes you have made.

Now you are ready to set things in motion by clicking the 'Calculate Flow' button or by selecting the same option from the Optical Flow menu. The length of time it takes to calculate the flow field data is dependent upon the length and resolution of your footage, as well as the type of motion in the shot. Simple scenes that don't contain any moving objects will finish faster than those with lots of complex motion.

As optical flow is calculated, data files for each frame are stored in the 'Optical Flow' folder in the location the project was saved. To delete the optical flow data files, right-click on the Optical Flow entry in the shot overview (which will only appear once optical flow has been calculated) and select Delete from the popup menu.

Feature Hints

Any assistance you can provide to further improve the quality of flow data is obviously a good idea. If you have complex objects that move by large amounts between each frame, you may find that the optical flow algorithm is not able to track the pixels of the object correctly from frame-to-frame. In this case you can track some user or auto-features and use these as 'hints' to tell the optical flow algorithm how pixels are moving. Any features that are present in the shot when optical flow is calculated will be used as hints.

The process of tracking and flow analysis can be carried out in one hit using the 'Track and Flow' button shown here, but in most cases, additional hint features will not be necessary.

Masks For Moving Objects

In certain situations, you may find that the optical flow vectors from one moving object incorrectly influence the vectors in another, especially along the boundary between the two objects. In these cases, you can draw a mask around your object using the standard tools, and use this mask to affect the optical flow calculations. Creating and editing masks is covered in the Masks section of this manual. Note that in contrast to constructing garbage masks for feature tracking purposes, it helps to make sure a mask for optical flow fits as accurately as possible around a moving object.

Once you are satisfied that the mask covers the problem item you need to set some instructions about how the mask should behave. Right click on the mask to reveal a pop-up menu, and select an option from the Flow Guide sub-menu: No Guide instructs PFTrack that the mask will not affect the optical flow calculations (this is the default setting). The second option, Motion Boundary is used when you want to use the mask to define the boundary between one moving object and another. To use this option successfully, it is important that the mask is drawn accurately around the moving object in each frame. Finally you can select Ignore Pixels to avoid calculating motion vectors for pixels inside the mask.

Feature tracks can be used at the same time as masks, but it is important to be aware that mask depth ordering affects influence of feature tracks: When a feature is placed over a mask, its influence will be restricted to pixels covered by the mask that is nearest to the camera plane.

Viewing and Editing Optical Flow

When optical flow has been computed you will see that the image is now overlayed with coloured arrows, indicating the direction and distance that each pixel travels between frames. If you zoom in to the image (middle button + forward/backwards) you will see that the density of these flow vectors increase according to the zoom level.

The direction of the vectors indicates the pixel movement within the image, while their colour corresponds to the usual colour code in PFTrack, so a green flow vector means that the similarity score for this pixel is high (or in other words, the pixel could be matched accurately with another in the image). It may help to darken the background image using the View>Darken menu option to help see the vectors more easily.

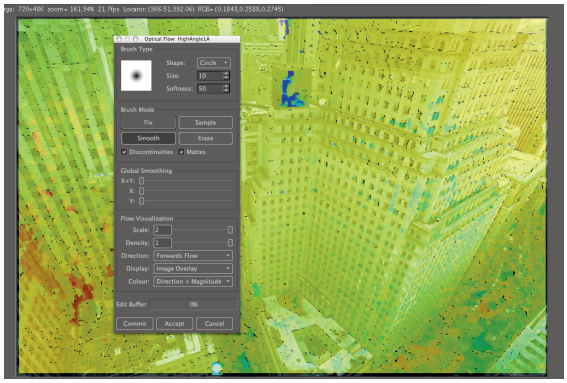

If there are any anomalies you would like to alter or if you simply want to smooth the flow field you will need the Optical Flow editing tools. To access these tools select Optical Flow>Edit... from the main menu, or right-click on the Optical Flow entry in the shot overview and select Edit... from the popup menu. The window illustrated here will appear, containing additional options for viewing and editing the flow data.

To make changes to the flow field, PFTrack employs a brush-based system so that the user can correct the flow field by hand. The brush can be adapted in several ways to suit the requirements of the job, including changing its size, shape and softness, as well as the effect the bush has on the data.

The brush is activated by firstly selecting a brush mode. From the illustration, you can see that there are four brush modes: Fix, Sample, Smooth and Erase:

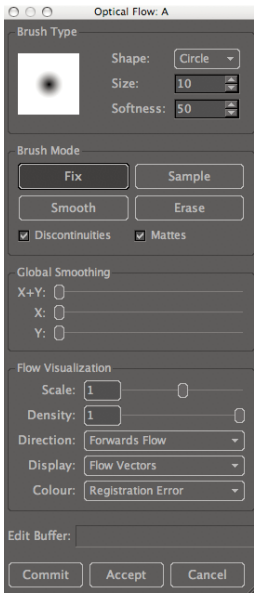

Fix: Painting the flow field with this brush will adjust the flow vectors so they point in the direction of the brush stroke. Brush vectors will not be adjusted if doing so would increase the error.

Sample: Hold the shift key and use the left mouse button to sample the flow field and calculate an average flow vector. Painting without the shift key held will paste this average vector into the flow field.

Smooth: The brush will smooth the flow field in painted areas.

Erase: Paints a discontinuity into the flow field, effectively removing motion vectors for pixels that are occluded in the next/last frame. A discontinuity is a place in the image where a pixel has no match in the next frame. For example: if a region of the image in one frame is obscured in the next by a moving object, the pixels in that region have no valid matches in the next frame. The erase brush will let you place these discontinuities in specific places.

Having made a selection, position your cursor over the area you want to treat and you will notice a circle representing the size of your brush will be drawn in the image window. A left click will now activate the brush, whilst the middle and right mouse buttons can be used to zoom and pan around the image. You will see that the flow vectors change colour within the

area of the brush to make it easier to identify which vectors are will be edited. Wacom graphics tablets are supported, and pen pressure can be used to control the brush softness. You can adjust the pressure sensitivity of your tablet pen from the Interface tab in the Preferences window.

Two further options to consider which alter the behaviour of the brush are Discontinuities and Masks. These appear as check-boxes beneath the main brush mode buttons. Each check-box gives you the option to decide whether the paint tools should respect discontinuities and mask boundaries. If you uncheck discontinuities then you can paint over them in the flow field. In the same way, if you uncheck masks the smooth brush will allow you to smooth the flow field over motion boundaries defined by masks.

Whilst editing flow data, all changes are stored in memory before being committed to disk. The edit buffer indicates how much memory has been used. As this buffer becomes full, you may wish to commit edits back to disk by pressing the 'Commit' button. This will commit edits for the current frame back to disk. Pressing the 'Accept' button will commit edits for all frames to disk and close the window. Finally, the 'Cancel' button will remove all edits and close the window.

Global Smoothing

The Global Smoothing controls apply a smoothing operation to the flow field across the whole sequence. There is a slider for each the X and Y axes as well as for the X+Y axes combined, where Y is the vertical image direction.

Flow Visualization

The items in the bottom section of the window concern themselves with Flow Visualization. These options do not affect the flow data directly, only how it is displayed on screen. The Scale and Density sliders adjust the scale and density of the flow vectors that are displayed in the image window. Adjusting these can help you see detail in the flow field more easily, particularly if you are closely zoomed in to your image.

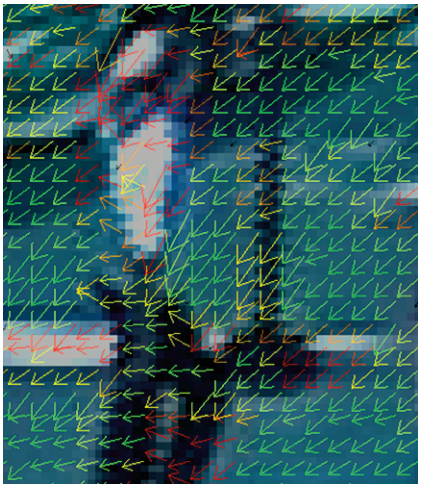

The Direction drop-down menu provides the ability to view either the forwards or backwards flow fields. Changes you make when editing one of the flow fields will automatically reflect in the reverse field data.

The Display menu options allow you to change how the flow field is displayed on screen. The default setting, Flow Vectors, shows the movement of every pixel as an arrow. The second option is Image Overlay, which colour-code of each pixel according to the setting of the Colour drop-down menu.

The final display option is Morphed Frame which is useful for visualizing how accurately two frames can be interpolated using the current flow field. If you are using the optical flow data for retiming with the PFRetime plugin, you can use this display option to quickly check if the flow field will produce an accurately retimed frame. When displaying a morphed frame, the image is generated by blending the current frame into the next one, according to how the pixels are moving, thus producing a new frame at the mid-point between the current and next frames. To quickly toggle this mode on and off, you can use a keyboard shortcut (Shift+M by default, although you can choose your own shortcut from the preferences window). Each time you make an edit to the flow field, the retimed frame will be adjusted automatically.

The Colour menu affects the colouring of both flow vectors and image overlays. The most important colour option is Registration Error, which will colour each pixel according to how well the image data matches from one frame to the next. If the pixel is the same colour, after it is flowed into the next frame then the pixel will be coloured green. Red pixels represent areas where PFTrack could not find a suitable matching pixel in the next frame. You can also choose to colour pixels according to the Direction and/or Magnitude of the flow vectors. The direction of flow is represented by changing hue, and the magnitude by changing brightness.

Exporting Optical Flow Data

Optical flow data as encoded RGB images, or the floating-point OpenEXR image format, and loaded into your animation or compositing application for use in your own pipeline. Optical flow corresponding to whatever flow direction you are displaying on screen will be exported.

To use encoded RGB image sequences in third-party applications, it is often necessary to write scripts or macros that are able to take RGB image pixels and convert them back to flow vectors and then perform whatever task you require. Alternatively, applications such as Apple's Shake are able to read OpenEXR image files, and read the x and y flow channels directly.

To export the flow data as an image sequence, right-click on the Optical Flow item in the shot overview and select Export... from the popup menu. This will allow you to save your flow data as a set of images for use in other applications. Exporting in OpenEXR format will store the true optical flow vectors, along with the RGB image data for each frame.

When exporting the flow data in another format, such as TIFF or Targa, the flow vector for each pixel is encoded as follows:

red channel: flow x value, normalized to fit between [-127..128]
green channel: size of the flow vector, clamped to 255 pixels
blue channel: flow y value, normalized to fit between [-127..128]

PFTrack also provides an option to export the data in a raw floating-point format. The format of these files is as follows:
FLOW
<width> <height>
<flow vectors>

Where <width> and <height> are ASCII integers and <flow vectors> are little-endian ordered binary floating point values (two per pixel, representing the x and y flow vector).

Motion Capture

Motion capture is the process of extracting movement information from a live action event such as a person running or dancing, or a talking face. It is commonly used in game and film production and applied to animated characters to give realistic animation.

PFTrack offers users the ability to undertake motion capture by shooting two or more camera views of the same event, and tracking and matching up user features within each sequence. An example of this would be to set up two or three cameras pointed at an actor's face. The main camera or principal camera would be the view you were going to use in post production, perhaps the front view. The second and third cameras could then be set up to provide right and left hand views. Tracking markers would then be placed on the face of an actor and the sequence filmed. The more markers placed on the face the better the resulting track. Once you have your three video sequences they can be imported into PFTrack and used to solve for the motion of each tracking marker.

Note that in contrast to normal object tracking, motion capture does not assume that the tracking markers are moving as a single rigid object. However, in order to properly capture the motion of an object, the cameras that you use to shoot your subject must be static.

Planning The Shot

When planning to use the motion capture facilities in PFTrack, it would be useful to keep the following points in mind:

Tracking Markers need to be placed on the subject before shooting takes place. The more markers on the subject, the better the resulting motion capture will be (assuming that features are tracked accurately for each marker). Generally, we recommend using small circular markers for motion capture, where each marker is coloured to contrast with the subject (for example, using small black markers on a face for facial motion capture).

When shooting your footage, you can mix-and-match different types of camera for the principal and secondary footage. The images for each view can be a different size and have a different frame-rate, but it is important to make sure the frame-rate value is set correctly for each piece of footage. PFTrack will solve motion using only two cameras but it is often worth using three for increased accuracy. Support for multiple resolutions and frame-rates means you can mix DV cameras for reference with a high resolution or film camera for the principal footage, providing a fast and cheap way to produce motion capture data.

The focal lengths for the different cameras do not have to be the same, but the cameras cannot move – they must be fixed in position whilst the subject is filmed.

Try to record an obvious sync frame, such as a clap or a blink, to help synchronizing clips inside PFTrack. If the clips are not correctly synchronized, then the motion capture results will not be accurate.

Importing Footage

Once you have two or more suitable video sequences they can be imported into PFTrack for tracking. After setting up a new shot in the usual way, you will need to import your principal footage. This is the footage that will be included in the motion capture export to third-party applications.

Import your first sequence in the normal way using the *File>Import Footage...* menu option or by pressing the Import Footage button in the toolbar.

PFTrack needs to be told that motion capture is to be performed. To do this, open the camera parameters window by choosing the *Camera>Camera Parameters...* option from the main menu and change the motion type to 'Motion Capture' using the drop-down menu. This will make the import footage icon active again, allowing more footage to be loaded.

Additional footage can now be imported into the shot. It is advisable to use the PFTrack windows tools to layout the clips in a logical fashion, although this is not essential because you can also work from within a single window. Use the Tile Window button if you need to tidy up the window layout.

Synchronizing Footage

The next stage is to synchronize the image sequences. This is the process of making sure each sequence is playing back exactly the same frame at the same time. A simple and logical toolset has been employed to synchronize the sequences if this operation has not been performed in an editing package before importing them into PFTrack.

To synchronize the footage, find a frame in each sequence where a specific event is happening (for example, a blink if you are filming facial motion). Move to that frame in the first sequence and click the 'Set Sync' button at the top-right of the time line. You will notice a blue marker is placed in the time line at that frame. Now find the frame containing the same event in the other sequences and use the 'Set Sync' button to mark the frame again.

Once you have set the sync frame for each sequence, you can press the 'Sync' button to ensure that when playing through the footage, each window is updated to show the same frame at the same time (You will need to open multiple image windows to see this).

Tracking User Features

The next stage is to track user features over each tracking marker, making sure that the same user feature is tracked over the same marker in each sequence. PFTrack will perform motion capture even if you only have a single user feature on your object, although obviously the more you use, the better quality solution you will have. You should also place enough user features to ensure you capture all the animation you need. For accurate results, we generally recommend using at least 5 or 6 features placed on key points to allow you to match them with any geometry you build in your 3D application.

You can create and track user features in the usual way (see the section on User Feature Tracking for more information). Once you are happy with the placement of a feature in the principal footage, click on the secondary sequence and track the same user feature on the corresponding tracking marker.

You may want to re-name each user-feature as it is created because this can help to easily identify the features and make sure that they are placed at the correct position in each image sequence. It is important to ensure that each feature is tracked in at least two image sequences. If a feature is only tracked in one sequence, it's 3D position cannot be solved.

Setting Rigid Feature Constraints

In order to help the solver generate more accurate motion capture data, you can specify rigid constraints between features. If several of the tracking markers on your subject do not move relative to each other (the end of a nose and the eyes on a human face, for example) then creating these constraints will often help the solve build a better result.

Feature constraints are created in the usual way, by opening the constraints window from the

Tracking>Constraints... menu option, as described in the Creating Feature Constraints section of this manual. For motion capture, 'Rigid' constraints should be used, where each constraint contains all the features that are moving rigidly with respect to each other.

Specifying these constraints is not always necessary, but doing so will often improve the quality of the motion capture solution.

Solving For Motion Capture

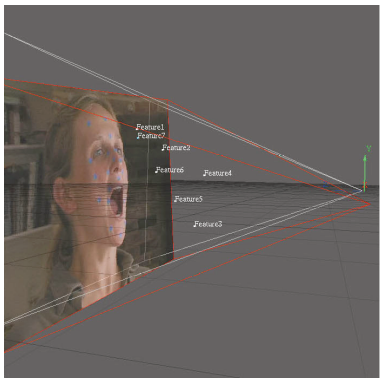

Once you have tracked every user feature in each sequence, you solve for the feature motion and camera positions by selecting *Camera>Solve Camera* from the menu or pressing the Solve button on the toolbar. Once the camera solve is finished, you can view the 3D point cloud in a 3D window before exporting.

After the camera positions and feature trajectories have been estimated, you can solve or un-solve additional points as you wish. Right-click on a user-feature to display a popup menu, and select 'Solve' or 'Unsolve' as necessary. This type of flexibility is common to PFTrack's non linear design and allows you to add extra tracking points to the solution without re-solving for the camera positions.

Z-Depth Extraction

PFTrack includes the ability to automatically derive per-pixel Z depth information from a sequence. The process of Z depth extraction uses 3D feature positions and camera information to produce an estimate of the distance of each pixel from the camera.

The quality of the results depends on the quality of the camera solution, and the type of motion the camera is undergoing. Because image data is being analysed to estimate depth, inconsistencies in a sequence can also give poor results: reflections or specular highlights, for example, can cause problems. Z Depth extraction uses solved feature points and/or modelling primitives as hints, so you can easily influence the solution by adding new user-features or by removing poorly positioned features before estimating depth.

Note that unless a stereo camera rig is being used, depth extraction is only possible for a camera that is translation: a single Rotation Only camera does not provide enough parallax information to estimate depth at each pixel. When using a stereo camera rig, parallax between the left and right eyes is often enough to provide an accurate estimate of depth.

When estimating Z depth from a single camera, it is necessary to mask out any moving objects, because object motion in 3D space will seriously affect the estimate of depth at moving pixels. When using a stereo camera rig, however, this restriction does not come into effect, and the depth of moving objects can be estimated.

To build a depth map, the first stage is to track and solve for the camera motion in your shot. After that, adjust the near and far clip planes in a 3D Viewer window to define the range of depth estimates you desire, and click the Depth Map button or choose Camera>Depth Map... from the menu to open the depth parameter window.

Depth data is displayed as either a grey scale image in an image window, or as a 3D-mesh or point cloud shown in a 3D viewer window. Both the grey scale image and the mesh can be

exported for use in compositing or 3D modelling systems. Depth maps can also be exported using the OpenEXR image format, providing a depth channel in addition to the standard RGB channels.

Setting Depth Parameters

The parameters in the depth window are listed fully in the Windows section of this manual, but here are the most important parameters that you need to consider:

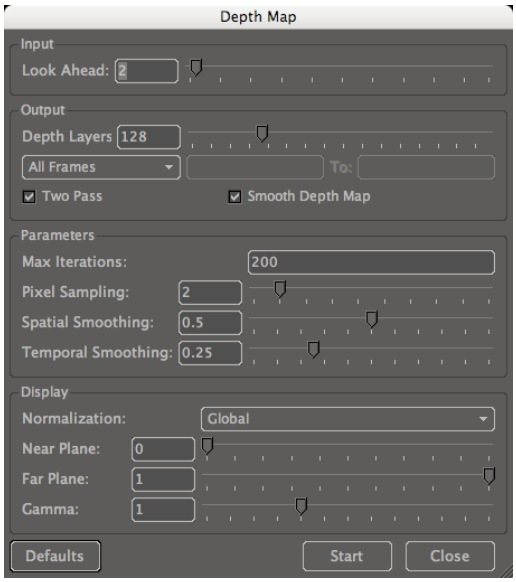

Look-Ahead sets the number of frames PFTrack looks ahead/behind the current frame in order to estimate depth. If you have a sequence with very slow movement you may want to increase this number to try and ensure that the camera undergoes a significant amount of translation between the frames. Lowering the number can also provide better results in fast moving sequences. For example, a look ahead of '2' means that for every frame of the sequence, depth will be estimated by comparing frame 'N' with frame 'N+2' and 'N-2'. It is also worth remembering that a pixel can only have its depth estimated if it is visible in the current frame and the ahead/behind frame. Because of this, the look-ahead value must be carefully chosen so as to produce enough camera translation whilst not moving too many pixels off-screen. If look-ahead is increased too much, it may not be possible to accurately estimate depth for pixels around the edge of the frame.

For stereo camera rigs, the Use Stereo Cameras box can be ticked. This will force PFTrack to compare image data in the primary and secondary cameras for each individual frame. Note that this also means that z-depth data can be generated for shots containing moving actors or other objects, provided there is sufficient parallax between the primary and secondary camera viewpoints. When depth is estimated from stereo cameras, the Look-Ahead parameter is disabled.

The Source menu is active when a stereo camera has been tracked, and the Use Stereo Cameras box is not selected. This menu allows you to choose which of the primary or secondary footage to use in the depth estimate. Selecting All from this menu will produce two independent depth maps, one for each eye.

The Output options allow you to set the number of depth layers (or depth resolution) of the image produced. This effectively controls the number of different grey levels, up to a maximum of 512. As a rule of thumb, increasing this number will produce more accurate depth maps, but will also increase the computation time. These depth layers will be positioned between the near and far clip planes, which can be adjusted using a 3D perspective viewer window. The drop down menu allows you to specify the range of frames to calculate depth for. Select the option you require from the drop down menu, and enter the frame range in the boxes if necessary.

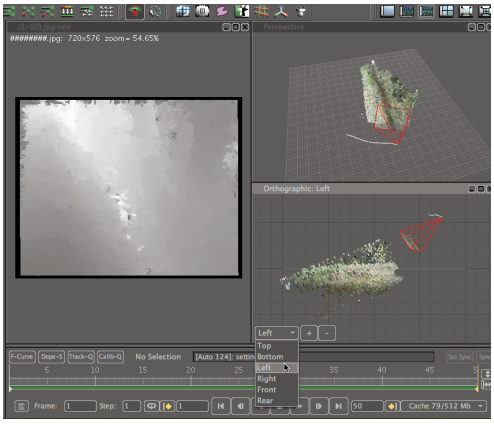

The other settings in the Parameters and Control sections of this window affect the time/quality trade-off for depth estimation. These are described

in more detail in the Windows section of this manual. Click the Defaults button to reset all depth parameters to their default values.

After altering any settings you wish, click the Start button to begin the depth estimation process. Depth estimation is a complex process, and may take a significant amount of time for high resolution image sequences. PFTrack will take advantage of multiple CPUs and CPU cores when estimating depth. Multi-threading can be enabled or disabled from the Performance tab in the Preferences window.

Once generated, you can use a 3D viewer window to visualise the depth data and see where pixels have been placed. If you do not get the results you expect you can alter settings and re-calculate the depth map. One common problem is the pixel streaking occurs at the edges of an image. If this happens adjust the Frame Boundary or create masks to avoid the outer edges of your frame.

The Display options affect the way the depth map is converted from floating-point to grey-scale pixel data for display on screen. The Normalization menu provides a choice of a Local per-frame normalization, in which the nearest/farthest pixels in the frame are mapped to white/black, or Global which maps the the nearest/farthest pixels in all frames to white/black.

The Near Plane and Far Plane settings change the boundaries of the depth map conversion to a grey-scale image. Increasing the Near Plane slider will adjust the white-point of the grey-scale image: setting the slider to 10% will push the near-point of the grey scale depth map back 10%. Decreasing the Far Plane slider has the reverse effect, in that it pulls the far-point or of the grey-scale image back towards the camera. This can be used to increase the accuracy of the limited precision grey-scale image in a specific area.

Once calculated, depth maps for each frame are shown in the Dope-Sheet with grey markers.

Exporting Depth Data

The grey-scale representation of depth maps can be exported using standard RGB image formats. The floating-point depth data can also be exported in the OpenEXR image format, either in a normalized or un-normalized form. To export the depth data as an image sequence, right-click on depth map item for your footage (in the Depth Maps container in the shot overview) and select Export Image Sequence... from the popup menu.

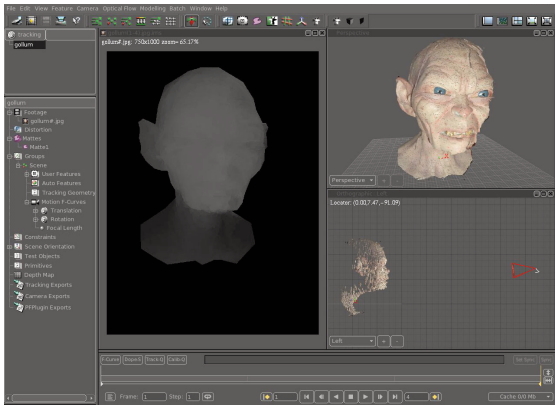

To export a depth map as an ascii-text file containing 3D point cloud data, select Export Point Cloud option will generate an Ascii text file containing the X,Y,Z position, R,G,B colour and the depth value (i.e. Distance from the camera) for each pixel.
Depth maps can also be saved as 3D triangular meshes when exporting to some 3D applications by setting the appropriate values in the Export window.

Image-Based Modelling

PFTrack includes a useful image-based modelling system that can help you to construct simple geometric models directly from the image data and camera path. Image-based modelling is the process of using camera data and visual references to build up 3D geometry. Unlike other systems, PFTrack can build geometry from a single still image, multiple still images or a moving sequence. A simple hierarchical modelling paradigm is employed to allow complex geometry to be built using simple primitive elements. These are manipulated using a vertex-based constraint system.

Image-based modelling primitives can also be used to provide hints for the Z-Depth calculation. This will happen automatically if you have built modelling primitives and then estimate Z-Depth.

Image modelling is an ideal tool for pre-visualisation and animation planning as it allows 3D animators or 2D artists to quickly build up an environment from background plates. In addition to matching the geometric characteristics of a background, you can also extract texture maps from one or more frames and map them onto the geometry.

Prior to carrying out image-based modelling, the camera motion must be solved and the ground-plane correctly oriented. The better the camera solve is, the easier it will be to perform the modelling because the camera focal length and positions will be estimated more accurately. It is also important to ensure that the ground plane is correctly positioned. The ground plane can be used to constrain modelling primitives, so placing it in a sensible position can greatly aid the modelling process.

Creating Primitives

Once you have an accurate camera solve, the main Modelling menu becomes active. From this menu it is possible to add primitives to your scene and extract textures. In order to make the modelling process faster you can switch on the Primitive Toolbar by selecting the option from the main Modelling menu. When activated, a floating menu containing your modelling primitives will be displayed on screen. You can 'dock' this menu to the other toolbars if you wish by dragging it to the location you require.

To create a primitive, either click it from the Primitive Toolbar or select it from the *Modelling>New Primitive* menu. This will place the primitive in the centre of your scene. After creating a suitable primitive you will see it displayed in the main image window, along with a set of control menus at the bottom left-hand corner.

You can import other geometric models using the Modelling>New Primitive>Add Primitive... menu option. PFTrack currently supports the import of Wavefront OBJ files, which can be exported from most common 3D modelling applications.

Once placed within the scene, a primitive can be selected by either clicking on its icon in the Shot Overview, or by clicking it inside the main image window. As primitives are added they will become a 'child' of any currently selected primitive, or if none is selected they will be independent from the other primitives. This structure can be clearly seen in the Shot Overview as a 'child' primitive will be shown linked to the 'parent'. The relationship between 'parent' and 'child' primitives can be used in combination with vertex constraints to restrict the positioning of the object. For example it would make sense to create a 'door' primitive as

a child of a 'wall' primitive. Primitives can be renamed as required by selecting them in the shot overview and then clicking again with the left mouse button.

PFTrack also provides an 'Auto-Mesh' modelling tool. This will triangulate a set of 3D feature points and build an editable mesh from them. To use this, select a set of feature points using the Lasso mode and then press the 'Build Mesh' button in the Modelling toolbar (or select Modelling>Build Mesh from the main menu). The features will be triangulated from the point-of-view of the current camera position to create a 3D mesh. This can be a useful tool when trying to model organic shape or landscapes.

Editing Primitives

Modelling primitives are manipulated by clicking and dragging their vertices to "pin" the object into position in the image. By pinning a vertex in two or more different frames, its 3D position can be determined, placing it correctly in the 3D space of your scene. As vertices are adjusted, the position and shape of each modelling primitive is updated to try and match the pin positions.

After creating a suitable primitive you will see a set of control menus at the bottom left-hand corner of the image window. The first menu controls the way the geometry is displayed: 'Hidden Line' is the default rendering mode with the other options being self-explanatory.

The second drop-down menu controls the current editing mode. The default is Object Edit, which will transform the entire object in a rigid fashion whilst a vertex is positioned. Vertex Edit mode allows you to edit the position of individual vertices in the object. The final option, Deform Edit is a combination of both Object and Vertex edit modes: the overall object shape and orientation of the primitive will be adjusted, whilst the vertex positions will also be changed independently where necessary.

The transformations that affect the primitive whilst it is being edited are controlled by the 'T', 'R' and 'S' buttons, which indicate that the object should translate, rotate and scale when the button is pressed. When in Vertex Edit mode, these transformations do not apply, but you can constrain the movement of a vertex to an axis by holding down the 'X', 'Y' or 'Z' modifier keys.

When you move a vertex around the screen you will see that the primitive is also transformed, and a small red dot appears that indicates the vertex is "pinned" to the image position. If you have oriented your ground plane correctly you can also right-click on a vertex and "pin" it to the ground plane or to another modelling primitive. The next vertex you move will also create a pin, and the primitive shape will be adjusted to attempt to match both pins. To properly specify the 3D position of a vertex you need to pin it in two or more different frames. After a vertex is pinned in one frame, attempting to pin in a different frame will display a blue vertex guide line to help you place the vertex in the correct position. If you want to move a vertex without leaving an anchor point, hold down the 'Shift' key when positioning it with the left mouse button.

You can restrict the way objects are transformed when in Object Edit or Deform Edit mode by selecting any combination of manipulation buttons, 'T', 'R' and 'S' for Translate, Rotate and Scale respectively. By default all manipulation buttons are active, and clicking on the buttons will switch the transformation off. For example clicking the 'T' button will prevent translation in any axis but because the 'S', scale and the 'R' rotate button are still active the object will still scale and rotate when necessary.

The Edit Mesh button allows you to add additional vertices into a primitive by clicking the

left mouse button at the surface point where you want a new vertex to be placed. Edges and triangles that can be subdivided will be highlighted in red as you move the mouse over them. This new vertex can then be edited in the same way as any other.

Vertex Constraints

Right clicking on a vertex of any modelling primitive will display a popup menu allowing you to constrain the vertex in a number of different ways. Selecting Pin To Vertex will allow a vertex to be pinned to another on another primitive. When this is done, a red dotted target line appears. Clicking on another vertex will then 'pin' the two together. Pin To Edge allows a vertex point to be pinned to an edge on another object. In this case, the vertex point is free to move along the edge in any direction. Similarly, Pin To Face will pin a vertex to a planar face of another object. Once a face is chosen, a circular target will be drawn to show the face that has been targeted.

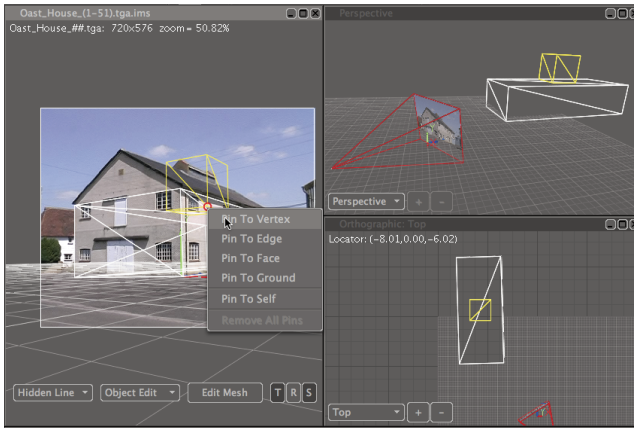

The remaining options in the menu are Pin to Ground and Pin To Self. Pin To Ground operates in a similar way to Pin To Face, but will constrain the vertex to lie on the ground-plane. Pin To Self operates similarly to Pin To Vertex, but allows a vertex to be pinned to another in the same primitive. Finally, Remove All Pins can be used to remove pins for a particular vertex from all frames in the sequence.

Extracting Texture Maps

Once primitives have been created, you can extract texture maps from the image sequence and have those maps automatically mapped on the surface geometry. To extract textures for all primitives at the same time, use the *Modelling>Extract Textures* menu option. Textures can also be extracted for one particular primitive by right-clicking on the primitive name in the shot overview and selecting Extract Textures from the popup menu.

Texture maps will be displayed in the image and 3D viewer windows whenever the display menu at the bottom-left of the image window is set to Textured. To remove all texture maps, use the *Modelling>Clear Textures* menu option (or the Clear Textures option from a primitive's popup menu to remove only the texture maps for that primitive).

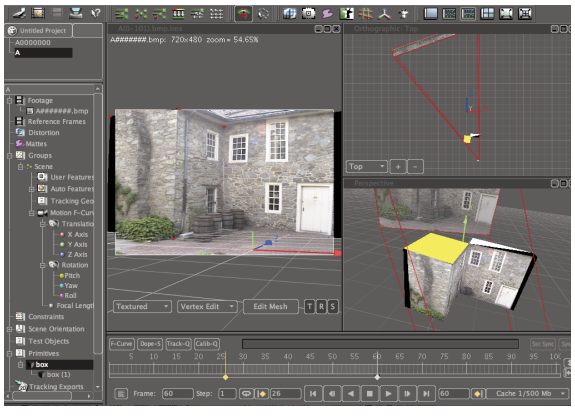

Additional menu options can be used to control the way that texture maps are created. The *Modelling>Extract Textures* From menu allows you to choose which images are used to construct the texture maps. The options are Current Frame which will use the current frame displayed in the image window to extract textures, and Best Frame to make PFTrack extract a texture map for each triangle from the frame where the triangle covers the largest number of pixels.

The *Modelling>Texture Fill Method* menu controls how PFTrack fills holes in texture maps. The choices are None to ignore hole filling, Uniform to fill holes with the average texture map colour, and Blend to blend colours together from nearly pixels.

Part 05 - Keyboard Shortcuts

Default Keyboard Shortcuts for PFTrack

This section lists the default keyboard shortcuts for PFTrack 5.0. You can change these if you wish by accessing the Keys tab in the Preferences window. Click on the shortcut you wish to edit and enter a new value. Modifier keys (such as Shift or Ctrl) need to be spelt out, e.g. "Ctrl+A" should be entered literally as C, t, r, l, +, A.

Menu	Action	Shortcut
File	New Project	none
File	Open Project	ctrl+O
File	Save Project	ctrl+S
File	Save Project As..	none
File	New Shot	ctrl+N
File	Open Shot	none
File	Duplicate Shot	none
File	Reset Shot	none
File	Close Shot	none
File	Delete Shot	none
File	Import Footage	ctrl+I
File	Import 2D Tracking	none
File	Import Reference Frame	none
File	Import Image Mask	none
File	Import Tracking Geometry	none
File	Preferences	none
File	Quit	ctrl+Q
Edit	Undo	ctrl+Z
Edit	Cut	ctrl+X
Edit	Copy	ctrl+C
Edit	Paste	ctrl+V
Edit	Delete	DEL
View	Darken	alt+K
View	Actual Size	alt+1
View	Unit Pixel Aspect Ratio	alt+S
View	Frustum Image	alt+X
View	Frame Boundary	alt+9
View	Axis Lines	alt+L
View	Distortion Lines	alt+I
View	Masks	alt+M
View	Colour Key	alt+H
View	User-Features	alt+U
View	Auto-Features	alt+A
View	Centre On Feature	alt+N
View	Feature Tracks	alt+C
View	Feature Labels	alt+B
View	Search Windows	alt+R
View	Groups Assignments	alt+O
View	Constraint Assignments	alt+P
View	Tracking Geometry	alt+2

View	Projection Errors	alt+E
View	Optical Flow	alt+Q
View	Solved Features	alt+Y
View	Missing Features	ctrl+M
View	Ground Plane	alt+G
View	Horizon Line	alt+5
View	Group Transforms	alt+8
View	F-Curve Targets	alt+7
View	Vertex Guides	alt+0
View	Test Objects	alt+3
View	Clip Planes	alt+D
View	Depth Cue	alt+Z
View	Depth Map	alt+6
View	Depth Mesh	alt+J
Tracking	Tracking Parameters	Shift+F2
Tracking	New User-Features	Shift+F3
Tracking	New Mask	Shift+F4
Tracking	Colour Key	none
Tracking	Auto Track	Shift+F5
Tracking	Track & Solve	Shift+F6
Tracking	Clean Auto-Feature Tracks	none
Tracking	Survey Data	none
Tracking	Groups	none
Tracking	Constraints	none
Tracking	Clear User-Features	none
Tracking	Clear Auto-Features	none
Camera	Camera Parameters	ctrl+P
Camera	Segment Motion Types	ctrl+Q
Camera	Lens Distortion	ctrl+L
Camera	Estimate Focal Length	ctrl+E
Camera	Solve Motion	Shift+F7
Camera	Statistics	none
Camera	Clean Auto-Features	none
Camera	Improve Solution	Shift+F8
Camera	Depth Map	none
Camera	Orient Scene	Shift+F9
Camera	Scale Scene	none
Camera	Transform Camera	ctrl+R
Camera	Un-Solve Frames	none
Camera	Delete Calibration	none
Optical Flow	Parameters	none
Optical Flow	Track and Flow	none
Optical Flow	Calculate Flow	Shift+F10
Optical Flow	Edit	Shift+F11
Modelling	Primative Toolbar	none
Modelling	Build Mesh	none
Modelling	Clear Primatives	none
Modelling	Extract Textures From..Current Frame	none
Modelling	Extract Textures From..Best Frame	none

Modelling	Texture Fill > None	none
Modelling	Texture Fill > Uniform	none
Modelling	Texture Fill > Blend	none
Modelling	Extract Textures	none
Modelling	Clear Textures	none
Batch	Open Script	none
Batch	Edit Script	none
Batch	Execute Script	none
Batch	Clear Script	none
Batch	Launch Batch Manager	none
Batch	Track off-line	none
Batch	Track and Solve off-line	none
Batch	Track and Optical Flow off-line	none
Batch	Solve off-line	none
Batch	Optical Flow off-line	none
Window	Layout->Store	None
Window	Layout->Rename	None
Window	Layout->Delete	None
Window	New Window	alt+W
Window	New Viewer	alt+V
Window	Tile Windows	alt+T
Window	Fit To Window	alt+F
Help	Log Output	none
Help	Show Log	none
Help	About	none
Help	What's This?	shift+F1

Playback Options	**Shortcut**
Play/Stop Movie	space
First Frame	down arrow
Last Frame	up arrow
Back One Frame	left arrow
Forward One Frame	right arrow

Mouse Button Emulation	**Shortcut**
Left Mouse Button Emulation*	A
Middle Mouse Button Emulation*	S
Right Mouse Button Emulation*	D
Mouse Wheel Up Emulation*	F
Mouse Wheel Down Emulation*	G

User Feature Selection	**Shortcut**
Previous User-Feature/Mask	Q
Next User-Feature/Mask	W
Previous User-Feature/Mask Keyframe	[
Next User-Feature/Mask Keyframe	]

Tracking	**Shortcut**
Remove From Frame	;
Remove From Frame + Forwards	Shift+>

Remove From Frame + Backwards	Shift+<
Track User-Feature Forwards	P
Track User-Feature Backwards	O
Decrease Vertex Weight	,
Increase Vertex Weight	.

Axis Controls	**Shortcut**
X Axis Lock	X
Y Axis Lock	Y
Z Axis Lock	Z

Toolbar Button Modes	**Shortcut**
Navigation Mode	N
Selection Mode	L
Distortion Mode	none
User-Feature Mode	none
Mask Mode	none
Colour Key	none
Coordinate Axis Mode	none
Scene Orientation Mode	none
New Test Object	none
New Modelling Primitive	none

Optical Flow Edit	**Shortcut**
Commit Optical Flow Edits	shift+C
Toggle Morphed Frame Display	shift+M

General	**Shortcut**
Abort	Esc
F-Curve Smooth	+
F-Curve Sharpen	-

Geometry Tracking	**Shortcut**
Object Edit Mode	1
Deform Edit Mode	2
Vertex Edit Mode	3
Mesh Edit Toggle	4
Change Object Render Mode	5
Translate Manipulator	6
Rotate Manipulator	7
Scale Manipulator	8
Fly Manipulator	9
Drag Manipulator	0

*these keyboard shortcuts can be helpful to simulate a multi-button mouse on OSX platforms when you only have a single button mouse available. To use them, hold the appropriate shortcut whilst pressing the left mouse button.

Part 06 - Scripting in PFTrack

Basic Scripting

Introduction

PFTrack script is a simple text-based scripting language which allows you to control the feature tracking, camera solving and optical flow functions of PFTrack without accessing the main graphical user interface. It has been designed to be easy to use by non programmers using logical statement based scripting.

Running script files

Scripts files, typically given the extension ".psc", can be run by specifying their location as an argument to the main PFTrack executable.

On Windows platforms, this is done as follows:

```
C:> pftrack.exe C:\path\to\script.psc
```

On Linux:

```
bash% pftrack /path/to/script.psc
```

On Mac OS X, the full path to the PFTrack executable is needed and since this typically contains spaces its enclosed in double quotes:

```
prompt>"/Applications/ThePixelFarm/PFTrack/PFTrack.app/Contents/MacOS/PFTrack"/path/to/script.psc
```

An example script

An example script is shown below:

```
#PFTrackScript v1.0

// start a new shot
<newShot>

// load the movie file
<importMovie> "/path/to/footage.mov"

// auto track, solve for camera and save
<autoTrack>
<solveCamera>
<saveShot> "/path/to/shot.pts"

<exit>
```

Hopefully the above script is self explanatory - it creates a new shot; imports some footage; auto-tracks the footage; solves for the camera parameters; saves the shot to disk and exits. This shot could be loaded into the GUI version of PFTrack for review.

Basic structure and syntax

All script commands are enclosed by angled brackets (<...>), and their parameters are enclosed with double quotation marks ("..."). Comments are represented using the standard C++ notation of double forward slashes (//). Commands are parsed in a case-independent manner, so "autoTrack" is equivalent to "Autotrack" and "auTotRACK". Filenames may be specified using either full path names, or files relative to the PF_HOME environment variable. The commands closely follow the menu items included in the user interface and perform identically to their GUI equivalents.

A script file must start with the identifier:

```
#PFTrackScript v1.0
```

without it no further processing of the file is performed. Optionally, the use of a specific version of PFTrack can be enforced using the command:

```
<requiresVersion> "4.1"
```

A script must end with the command:

```
<exit>
```

A log of script processing can be generated using the command:

```
<logLocation> “/path/to/log.txt”
```

Shot control

There are 3 command relating to shot control, these are:

```
<newShot>
<loadShot> "/path/to/shot.pts"
<saveShot> "/path/to/shot.pts"
```

Note that scripting works with individual shots, whereas the GUI version of PFTrack works with projects (a project being a collection of shots). However, a shot can be trivial included into an existing project.

Importing Footage

To import movie files such as AVIs and QuickTimes use the command:

```
<importMovie> "/path/to/footage.mov"
```

Importing image sequences is a bit more involved. To use an image sequence in scripting first you must construct what is known as an IMS file. An IMS file, which stands for Image Movie Sequence, is a simple plain text file that details how a series of discreet image frames make up a movie sequence. An example IMS file is given below and consists of the identifier "#IMS", followed by the number of frames, the frame rate and then the path to the image file for each frame:

```
#IMS
4
25.0
"/path/to/frame1.jpg"
"/path/to/frame2.jpg"
"/path/to/frame3.jpg"
"/path/to/frame4.jpg"
```

This IMS file is then used as the parameter to "<importMovie>", eg:

```
<importMovie> "/path/to/footage.ims"
```

Interlacing and 3:2 pulldown

If the footage is interlaced, use:

```
<setParameter> "interlacing" "mode"
```

to specify this where the mode parameter should be one of "None" (the default), "fieldAveraging", "interpolateLower", "interpolateUpper" or "fieldSeparation".

The field dominance used for deinterlacing is controlled with:

```
<setParameter> "fieldDominance" "mode"
```

where mode is either "lowerFirst" or "upperFirst".

To remove 3:2 pulldown from NTSC sequence use:

```
<setParameter> "pulldown" "mode"
```

where mode should be one of "None", "WWWSS", "WWSSW", "WSSWW", "SSWWW" and "SWWWS" where "W" is a whole frame and "S" is a split frame.

Mattes

To restrict auto-feature tracking to particular areas of the footage, a matte image or movie can be specified using the command:

```
<importMask> "/path/to/matte.mov"
```

Where the file is either an AVI/Quicktime or an IMS file as described above.

Tracks

The <importTracks> command will load human-readable 2D tracking data into the shot:

```
<importTracks> "type" "/path/to/file.2dt"
```

The type argument must be either "auto" or "user" to specify if user or auto-features are to be created. The second argument is the filename of the tracking file. This file must be of the same format as the "2D Tracking" files exported by PFTrack.

Camera Parameters

The "cameraPreset" parameter lets you specify a pre-defined camera preset after the footage has been loaded:

```
<setParameter> "cameraPreset" "35mm Full Aperture"
```

The name of the preset ("35mm Full Aperture" in the above example) must match one of the preset names within the PFTrack Camera Parameters window.

To specify the format of the footage, if it cannot be guessed automatically, use:

```
<setParameter> "format" "type"
```

The type argument should be either "PAL", "NTSC" or "Film".

To specify a frame rate if it cannot be automatically guessed, use:

```
<setParameter> “frameRate” “value”
```

The value argument should be a numerical frame rate, such as “25.0” or “29.97”.
The type of camera motion in the footage is specified with:

```
<setParameter> “cameraMotion” “value”
```

where the value argument should be either “freeMotion” or “rotationOnly”.

Focal length type is set using:

```
<setParameter> “focalType” “value”
```

The value argument should be either “constant” or “variable”.

For a constant focal length the following command:

```
<setParameter> “focalFix” “mode” “value” “units”
```

can be used to specify known or approximately known focal lengths. The mode arguments should be either “unknown”, “approximate” or “known”. In the case of “known” or “approximate” focals, the value argument should be the numerical focal length, and the units argument should be either “mm” or “pixels” depending on the units of the numerical focal length. If the mode argument is “unknown”, the value and units arguments should be omitted.

Auto Feature Tracking

To perform an auto feature track use the command:

```
<autoTrack>
```

There are many optional parameters that can be changed which effects auto tracking, although the default settings are usually sufficient:

```
<setParameter> “startFrame” “value”
<setParameter> “endFrame” “value”
```

These commands will set the start and end frame for feature tracking and camera solving. The value argument should be a whole number representing the desired frame.

```
<setParameter> “numFeatures” “value”
```

This command will specify the number of auto-features that will be tracked. The value argument should be a whole number, such as “300” or “1000”.

```
<setParameter> “selectSeparation” “value”
<setParameter> “selectThreshold” “value”
<setParameter> “minSeparation” “value”
<setParameter> “maxResidual” “value”
```

These commands will set the Selection Separation, Selection Threshold, Min Tracking Separation and Max Tracking Residual parameters. The value argument should be a floating-point number, such as “1.0” or “3.64”.

```
<setParameter> “minLength” “value”
<setParameter> “searchRange” “value”
```

These commands specify the Min Tracking Length and Search Range parameters. The value argument should be a whole number, such as “5” or “100”.

```
<setParameter> “trackColour” “value”
<setParameter> “trackRed”, “value”
<setParameter> “trackGreen”, “value”
<setParameter> “trackBlue”, “value”
<setParameter> “backTrack” “value”
<setParameter> “trackAffine” “value”
<setParameter> “trackDeflicker” “value”
<setParameter> “trackGuided” “value”
<setParameter> “trackConsistency” “value”
<setParameter> “areaFeatureLimit” “value”
<setParameter> “spreadFeatures” “value”
```

These commands will control different aspects of auto-feature tracking, as seen in the Tracking Parameters window. The value argument should be either “true” or “false”.

Camera Solver

To start the camera solver on a feature tracked shot, use:

```
<solveCamera>
```

Like feature tracking there are many optional parameters that can be changed which effect the camera solver although their default settings is usually sufficient.

```
<setParameter> “startFrame” “value”
<setParameter> “endFrame” “value”
```

These commands will set the start and end frame for feature tracking and camera solving. The value argument should be a whole number representing the desired frame.

```
<setParameter> “inlierThreshold” “value”
```

This command will set the inlier threshold for the camera solver. The value argument should be a floating-point distance measured in units of pixels.

```
<setParameter> “percentageOutliers” “value”
```

This command sets the estimated percentage of outliers (i.e. Bad feature tracks). The value argument should be a floating-point number between 0.0 and 100.0.

```
<setParameter> “autoKey” “value”
<setParameter> “singleFrameStep” “value”
<setParameter> “user” “value”
<setParameter> “auto” value”
```

These commands control various aspects of the camera solver, as described in the Solver Controls window. The value argument should be “true” or “false” to switch each control on or off. The “user” and “auto” parameters specify whether user or auto-features should be included in the solution.

```
<setParameter> “keyframe” “value”
```

This command can be used to specify a keyframe position for the camera solver. The value argument should be a frame number. To use manual keyframe mode, use the <setParameter> “autoKey” “false” command, as described above.

```
<setParameter> “similarity” “mode”
```

This command will set the similarity measure used to track auto-features. The mode argument should be either “NCC” for normalized-cross-correlation, or “RMS” for root-mean-squared error.

Improve Solution

After the camera has been solved, the following commands control how the solution can be improved if the solution is not satisfactory:

```
<setParameter> “focalWeight” “value”
```

This command controls the weight of the focal length constraint, and is only available for cameras with a constant focal length. The value argument should be a floating-point number greater than zero.

```
<setParameter> “improveIterations” “value”
```

This command sets the number of iterations that will be used when trying to improve the solution. The value argument should be a whole number.

```
<setParameter> “improveConvergence” “value”
```

This command sets the convergence threshold used when improving the solution. The value argument should be a floating-point number, and can be represented using scientific notation such as “1.0e-10”.

```
<improveSolution>
```

This command starts the solution improvement process.

Export

To export the current shot to a third-party application, use:

```
<export> “format” “/path/to/file”
```

where format argument must match the name of one of your PFTrack export plugins, such as “Maya Ascii 4.5” or “Softimage XSI”. These names can be found by looking at the export options within PFTrack itself.

Optical Flow

To start optical flow use the command:

```
<opticalFlow>
```

The following optional commands control the optical flow calculation:

```
<setParameter> “flowRange” “mode”
```

This command sets the range of frames for which optical flow will be calculated. The argument should be one of “all”, “current” or “range”.

```
<setParameter> “flowFrom” “value”
<setParameter> “flowTo” “value”
```

These commands set the from and to frame numbers when using a flowRange value of “range”.

```
<setParameter> “flowQuality” “mode”
```

This command sets the quality of the optical flow calculation. The mode argument should be either “standard” or “high”.

```
<setParameter> “blockSize” “value”
```

This command sets the block size for the optical flow calculation. The value argument should be a whole number.