

Training Device Software Enhancement

Stephen J. Dow
Department of Mathematical Sciences
The University of Alabama in Huntsville

May 14, 2003

Final Report

F/DOD/ARMY/AMCOM/Training Device Software Enhancement
DAAH01-97-D-R005 D.O. 37

Period of Performance: 8/28/01 to 12/30/02

The Software Engineering Directorate has developed a series of Basic Skills Trainers which display moving targets on terrain models based on photographic terrain imagery. Each terrain model includes a day image and corresponding IR image constructed from photographs. In order to create images covering the desired field-of-view at the desired resolution, it is necessary to take a series of day and IR photographs from a single location, rotating the camera between shots, and to combine or "stitch" these photos into a panoramic composite. The day images are stitched in a software application called Range Finder in conjunction with building the ground model. This report documents the mathematical methods and operational details used in Range Finder related to stitching.

Several earlier reports included documentation of certain aspects of the Range Finder application, and provide background for the present report. In particular, a 1998 report entitled *Basic Skills Trainer Software Development* introduces the concepts of ground triangulation, stereo models, and relative orientation, and a 2001 report entitled *Missile Trainer Workstation Enhancement* concentrated on the block adjustment process. At the time of the 1998 report, the ground triangulation was constructed one stereo pair at a time, and then the individual triangulations combined into the final ground model during the stitching process. Recent versions of Range Finder deal with the complete "block" of photos together, which allows greater flexibility in the geometry of photos which can be combined into a single terrain (wider fields of view as well as both horizontal and vertical overlapping photos).

We begin with some mathematical preliminaries, essentially as presented in the earlier reports. Photographs are taken with a high-end digital camera to produce the input images. Pixels in these input images are indexed in image coordinates, the upper left pixel having coordinates (0, 0) and the lower right pixel coordinates $(n_c - 1, n_r - 1)$, where n_c is the number of horizontal columns, n_r the number of vertical rows. Each photograph also defines a 3D camera coordinate system oriented about the camera as it was situated when that photo was taken, having the camera center of projection at the origin, x-axis to the right, y-axis down, and positive z-axis forward along the line of sight of the camera. Units in this coordinate system are meters. Let (x_c, y_c, z_c) denote the camera coordinates of a physical location in the scene. The line joining this point and the origin passes through the plane

$z_c = 1$ at $(x', y', 1)$, where $x' = x_c/z_c$, $y' = y_c/z_c$. Coordinates (x', y') are called photo coordinates. Thus 3D camera coordinates are converted to 2D photo coordinates by dividing by the z-coordinate.

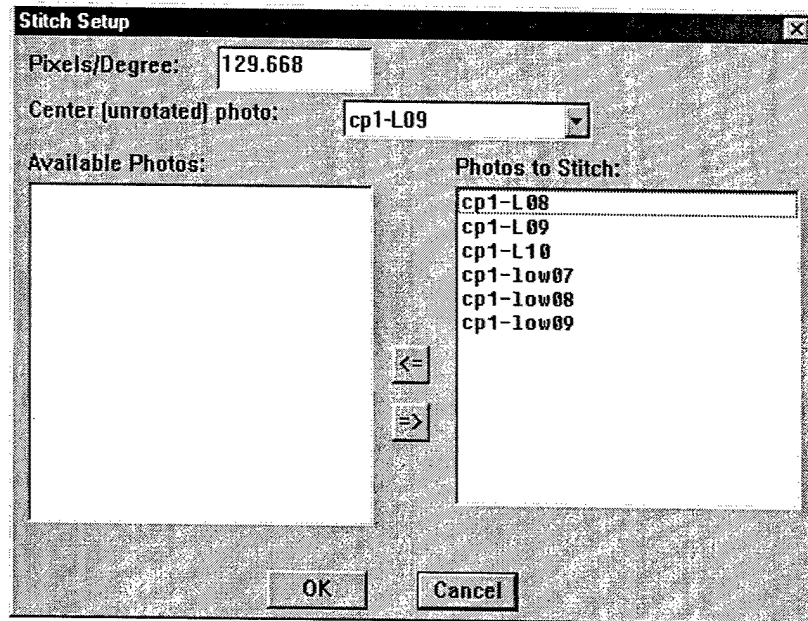
Photo coordinates (x', y') are related to image coordinates (x, y) through the interior orientation (IO), a set of parameters describing the internal geometry of the camera. The IO parameters and equations are given in the earlier reports and omitted here. The block adjustment step computes the IO parameters, along with the translations and rotations between the camera coordinate systems of the photos. The translations represent the movement of the camera from one location to another, multiple locations (called camera stations) being used to allow 3D ground points to be derived by stereo intersection. However, the stitch involves only the photos at a single camera station (called the primary camera station, which is camera station 0 in the software). We call these the stitch photos (or more specifically, if only some of the photos taken from that camera station are used in the stitch, then those are the stitch photos). We are making the assumption that the center of projection of the camera lies on the axis of rotation when the camera is rotated between shots, and so does not move (i.e. this movement if present is negligible). Thus the translations do not directly enter into the stitch computation. Rotations between photos are represented by Euler angles. The initial photo at the primary camera station, called the reference photo, is treated as the basis for the rotations; i.e. its three Euler angles are all zero. The Euler angles of any other photo are the angles which rotate its camera coordinates to those of the reference photo. This means that if a point has camera coordinates (x, y, z) relative to a given stitch photo with Euler angles e_x, e_y, e_z , then that point's camera coordinates relative to the reference photo are (x', y', z') , where $[x' y' z'] = [x y z] R_x R_y R_z$, and R_x, R_y and R_z are 3 x 3 rotation matrices about the x, y, and z axes respectively. The rotation angles e_x, e_y, e_z are displayed on the Block Photos dialog, under the columns labeled rx, ry, rz:

#	name	# meas	c.s.	rx	ry	rz	adj	io
0	cp1-L08	17	0	0.000	0.000	0.000	I	0A
1	cp1-L09	15	0	0.028	-8.918	0.004	A	0A
2	cp1-L10	9	0	0.072	-17.290	0.088	A	0A
3	cp1-R08	8	1	-0.460	1.955	-0.579	A	0A
4	cp1-R09	13	1	-0.408	-6.576	0.586	A	0A
5	cp1-R10	6	1	-0.566	-14.821	0.111	A	0A
6	cp1-low07	5	0	-3.685	6.269	-0.820	A	0A
7	cp1-low08	6	0	-3.604	-2.091	-0.831	A	0A
8	cp1-low09	4	0	-3.696	-9.200	-1.315	A	0A

In this example there are two camera stations, with six photos taken from camera station 0 (the first three and last three in the list) and three photos taken from camera station 1. Looking at the rotation angles, within each group of three successive photos, most of the variation in the angles occurs in the ry column, with ry decreasing by about nine degrees from one photo to the next. This resulted from the camera being rotated on a tripod by about nine degrees between successive shots, with the axis of rotation being nearly vertical. Note that the rx values for the three "low"

photos are significantly different from those of the other photos, the result of the camera being tilted down about 3.7 degrees before taking those three photos. Recognizing these patterns in the angles on the Block Photos dialog is useful in being able to evaluate whether sensible rotations have been produced by block adjustment to begin the stitch process.

Performing a stitch begins with the Stitch Setup command, which brings up the dialog shown below:



When this dialog is displayed for the first time for a block, all block photos at camera station 0 will appear in the list on the left, and the list on the right will be empty. In our example, the six photos at camera station 0 were initially listed as Available Photos, and were then moved to the Photos to Stitch list. The buttons in the middle of the dialog allow photos to be transferred back and forth between the two lists; each list is a “multiselect” list, meaning any number of the photos may be highlighted in either list, those photos then being the ones moved when the appropriate arrow button is hit. The order of the photos listed on the right is significant: it determines the stack order used in the stitch; i.e. in overlap areas the final pixel in the stitched photo originates from the topmost photo in the list which covers that pixel. This stack order can be controlled by taking note of the fact that the middle buttons always transfer photos to the *bottom* of the other list. Thus a photo in the right list can be moved to the bottom of that list by first transferring it to the left list and then back to the right.

To fully explain the two input fields at the top of the dialog (Pixels/Degree, Center photo) it will be useful to discuss what the software does internally in computing a stitch. The stitched image, also called the composite image, is the result of a spherical projection. The process may be visualized as projecting each of the input images onto a sphere centered at the common camera station of the photos. The input images may be visualized as flat rectangular surfaces tangent to the sphere at the points where their lines-of-sight intersect the sphere. The camera coordinate system of the photo specified as the “Center (unrotated) photo” becomes the underlying coordinate system of the composite image (and of the ground model; i.e. ground points have coordinates in this system). The

rows and columns of the composite image are similar to lines of latitude and longitude on the surface of the earth; the portion of the spherical surface covered by the input images is sampled along those lines, with the number of samples (pixels) per degree being specified as the Pixels/Degree parameter.

When the OK button on the Stitch Setup dialog is hit, computations are performed to determine the angular extents of all the input images, how large the composite image will be, and which input image will be visible at each output pixel. However the actual stitching (resampling to produce an output image) is deferred until either the Stitch Visible or Stitch IR command is selected. It may be useful at this point in the workflow (after Stitch Setup but before the actual stitch) to take a look at the pixel residuals for the stitch photos on the Point List dialog, as shown below:

The screenshot shows a window titled "Measured Points" with a close button (X) in the top right corner. Below the title bar, there is a control area with "Active point: 912", an "Apply" button, and a "New" button. Below that, there are two dropdown menus: "Pixel Residuals" and "Stitch Photos". The main area is a table with the following columns: "id", "cp1-L08", "cp1-L09", "cp1-L10", "cp1-low07", "cp1-low08", and "cp1-low09". The table contains 20 rows of data, with values ranging from 0.00 to -5.43 and 5.28.

id	cp1-L08	cp1-L09	cp1-L10	cp1-low07	cp1-low08	cp1-low09
800	0.00	0.00	-0.00	0.89		
801	0.00	0.00	-0.07	1.27		
802	0.00	0.00	-1.27	-0.32		
803	0.00	0.00	-0.38	-1.18		
804	0.00	0.00	-5.43	2.86		
900		0.00	0.00	-0.51	0.44	
901		0.00	0.00	0.26	-0.26	
902		0.00	0.00	0.43	-1.98	
903		0.00	0.00	-2.54	1.74	
2700	0.00	0.00		-1.84	0.68	0.37
2701	0.00	0.00		3.65	1.07	5.62
2702	0.00	0.00		-1.58	-0.03	0.75
2703	0.00	0.00		-0.30	0.91	
2704	0.00	0.00		4.24	-4.09	
2800	0.00	0.00			-2.70	0.41
2801	0.00	0.00			-6.68	-3.64
2802	0.00	0.00			3.24	1.45
2803	0.00	0.00			3.87	3.03
2900		0.00	0.00			3.60
						5.28

Residuals are displayed for points measured on two or more of the stitch photos. Values of 0.00 appear under the first stitch photo for which there was a measurement (the initial measurement). The residuals under each subsequent photo represent the difference between where that measurement would be projected on the composite image and where the initial measurement would be projected. Large values (either positive or negative) indicate areas of two photos which do not line up well under the current IO and rotation parameters computed by the block adjustment, and will usually be reflected in noticeable seams in the composite image. Bear in mind that the size of these residuals is affected by the Pixels/Degree parameter from the Stitch Setup, as the residuals are given in units of pixels in the composite image.

When a stitch is performed (Stitch Visible or Stitch IR), the resulting composite image is displayed. In general the main Range Finder window displays up to three panes – two side-by-side panes for displaying the input images and another pane for displaying the composite image. The two rightmost buttons on the Range Finder toolbar control which of these panes are displayed; when the “Block Photos” button is depressed the two side-by-side panes are displayed, and when the “Composite Photo” button is depressed the pane with the composite image is

displayed. It is possible to depress both buttons and see the composite photo in a top pane with the two side-by-side panes below. Following the Stitch Visible or Stitch IR command, the composite image is displayed, but it has not yet been saved to a file on disk. That action is performed through the "Save Composite Image" command under the Stitch menu. However, the user typically will want to perform some brightness adjustment on the composite image prior to saving it, in order to "balance" portions of the composite image originating from different input images. Selecting the "Adjust Brightness" item on the Stitch menu brings up an Adjust Brightness dialog box. This dialog has a pulldown field for selecting one of the input images; the portion of the composite image originating from that input image is adjusted when the brightness slider is moved. Typically the user will want to scroll the composite image to an area where a seam is visible, select one of the two photos giving rise to the seam on the Adjust Brightness dialog, and then adjust the brightness to hide the seam as much as possible. If adjusting one seam in this way is causing another seam to get worse, the option for separately adjusting the four corners may be helpful. For example, if one photo's region needs to be brightened to balance it with a photo on the left but darkened to balance it with a photo on the right, this can be accomplished with the four corners option by brightening with the corner set to upper left and again for lower left and then darkening with the corner set to upper right and to lower right. When the composite image is satisfactory, save it to disk. Typically, we use the same base name for the visible composite image as for the block; e.g. cp1.blk would have composite visible image cp1.jpg. The Stitch IR command may be used to create a composite IR provided that IR images corresponding to each input visible image have been created; in practice this method has not been used; instead the composite IR image is created directly from the original IR frames using the composite visible image as the base layer in the Overlay program. In summary, the items under the Stitch Menu – Stitch Setup, Stitch Visible, Stitch IR, Adjust Brightness, Save Composite Image – represent the steps to be performed for stitching in the order the steps are normally performed, with the exception of the Stitch IR command.