Automatic Generation of Seamless Mosaics over Extensive Areas from High Resolution Imagery

Roger Xandri, Fernando Pérez-Aragüés, Vicenç Palà, Roman Arbiol
Remote Sensing Department, Institut Cartogràfic de Catalunya
08036 Barcelona, Spain

ABSTRACT

When covering an extensive area by means of geocoded airborne or satellite imagery there is a need to combine a fairly large number of single images into a big, seamless mosaic. To accomplish this objective there are two main steps to follow. In the first place, some kind of radiometric approximation between contiguous, somewhat overlapping parts must be performed. This approximation can include radiometric slope elimination, removing of local illumination effects (hotspots), dynamic range and histogram balancing. As this is usually not enough, the second step is to find a method to generate a single image in the common area between two overlapping pieces that is better than just pasting one image onto the other. Adaptive feathering is discarded in favor of an approach that searches for an optimal seam that separates the pixels contributing to the mosaic from one image or the other. At a higher computational cost, this approach makes it possible to define “forbidden” zones that should not appear in the final mosaic, thus avoiding virtually any number of problematic areas. Nevertheless, the radiometric preprocessing must be good enough so that the images have small differences on the overlapping area.

Keywords: Seam, Mosaic, Orthophoto, Hotspot, Cost Function.

1. INTRODUCTION

Orthophoto projects use to include several hundreds of photo shots that must be combined to produce a radiometrically and geometrically continuous series of orthoimages [1], [2]. These orthoimages must maintain, at the same time, good radiometric properties regarding contrast and dynamic range. A possible approach is to remove the physical radiometric effects in order to obtain homogeneous images that maximize the probability of obtaining good continuity. After the mosaic is complete, local contrast enhancement and color balancing techniques are used to obtain a high contrast orthoimage.

Usual approaches apply some kind of feathering-like techniques (e.g. [1]). They have been discarded for our purposes due to the blurring originated by numerous small geometrical disparities, especially when working in projects involving large-scale photographs.

This paper deals with operations that lead to the final mosaic before radiometric enhancement, emphasizing the procedures that guarantee the continuity of the result. First, the radiometric preprocessing will be addressed, separating the different physical contributions. Then, the general problem of finding a seam between two images is studied and some enhancements to the initial algorithm are proposed. Next section is devoted to the generalization of these techniques in order to be applied to a large block of images. Finally, practical results are presented and future strategies are outlined.

2. RADIOMETRIC PREPROCESSING

The first step is to eliminate - or, at least, reduce as much as possible - the radiometric differences between adjacent shots. These differences are mainly due to temporal and spatial variations between acquisitions and also to artifacts introduced during the digitization process [3]. The procedure designed at the ICC intends to achieve the following goals:

· the final result should be a homogeneous block of images with smooth radiometric transitions between photos. The contents of a block must not be perceived as a mosaic of several, distinct pieces. Moreover, the appearance of the final image must be approximately the same irrespective of the order in which the pieces have entered the mosaic.
· It must be born in mind that the resultant image will often undergo generalization processes in order to produce smaller scale series. This generalization, usually via some kind of averaging, tends to entail a loss of dynamic range. Thus, the dynamic range of the final mosaic must be high enough. In most cases a trade-off between high dynamic range and saturation control must be found.
· In order to make the process scalable without extra effort, intrinsic information contained in the images must be used whenever possible. External information must be restricted to the minimum.
· If possible, initial radiometric differences must be physically modeled and evaluated.

Since color balancing is unnecessary prior to mosaic, the problem can be restricted to single band images without loss of generality. In this view, a complex radiometric correction is to be applied to the i^{th} shot whose radiometry on every pixel can be described as $F_i(x,y)$. The corrected radiometry for this photo will take the form

$$F_i'(x,y) = R(F_i(x,y)) + G_i(x,y) \tag{1}$$

where R is applied to the radiometric values of the image and G acts on the pixel coordinates. R functions will be referred as contrast corrections and they correct radiometric differences between corresponding points on adjacent images that depend on the pixel radiometric value. They apply when dynamic response of the sensor is changed between acquisitions (aperture change) or when dealing with scanned analog photographs (customized dynamic range for each scan). G functions will be called spatial corrections. They correct the radiometric value of a pixel according to the position of the pixel in the image and are aimed to correct effects like vigneting, atmospheric artifacts or hotspots. These two corrections are applied consecutively.

Contrast Corrections
It would be desirable that, after correction, the histograms of two neighboring images were identical in their overlapping parts. As a measure of this identity, the mean (μ) and standard deviation (σ) of the common areas will be used. Thus, every image $F_i(x,y)$ $i = 1...n$ will be corrected into $\hat{F}_i(x,y)$ $i = 1...n$ through a contrast function $H_i(F_i(x,y))$ that modifies the mean and standard deviation values in the common parts of the resultant image from μ_i, σ_i to $\hat{\mu}_i, \hat{\sigma}_i$ The differences between final values will be minimized in the least squares sense

$$E_1 = \sum \left\| \hat{\mu}_i - \hat{\mu}_j \right\|^2 \tag{2}$$

$$E_2 = \sum \left\| \hat{\sigma}_i - \hat{\sigma}_j \right\|^2 \tag{3}$$

Spatial Corrections
Three distinct effects will be mentioned here whose spatial dependencies are illustrated in figure 1.

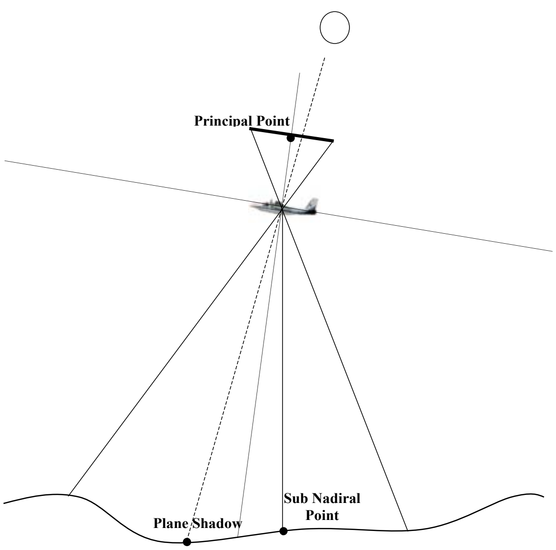

Figure 1. Shot geometry

Vigneting. Almost inevitable in wide-field photographs as in the case of aerial photography. It appears as a change in field illumination as we progress from the optical axis towards the edge of the lens system. It shows a basically radial dependence on the distance from the optical axis in the image space.

Air-mass correction. Depending on the orientation of the camera, the optical path from each pixel target on the ground may change more or less drastically depending on the optical thickness of the atmosphere and the length of the light rays. This phenomenon is especially noticeable in non-vertical aerial photography. It shows a radial dependence on the distance from the sub-nadiral point in object space and can be mostly merged with vigneting when dealing with vertical photography.

Hotspot. In sun illuminated scenes the amount of light captured by the sensor is governed by the so-called Bi-directional Reflectance Distribution Function (BDRF) with contributions from both the ground and the atmosphere [4]. In vertical aerial photography, the backscattering component of this function is dominant and a luminous diffuse spot is noticed around the shadow of the plane on the ground. Thus, it is possible to calculate this point using the position and orientation of the plane, a Digital Elevation Model (DEM) and the acquisition time, applying then a correction around.

Let G be the compound function that takes in account all these corrections and is applied on the resulting image $F_i'(x,y)$ $i = 1...n$ after contrast correction

$$F_i'(x,y) = \hat{F}_i(x,y) + G_i(x,y) \tag{4}$$

The objective is again to minimize differences in the overlapping parts of the corrected images

$$E_3 = \sum \left\| F_i'(x,y) - F_j'(x,y) \right\|^2 \tag{5}$$

Absolute Reference
It can be easily noticed that all corrections that have been postulated are bounded to relative values between images i.e. the summatories described before are applied only on the overlapping areas between images. The result is then likely to have the correction functions corresponding to one of the photos badly resolved. This drawback can be easily solved posing a restriction on, let's say, the mean radiometric value of all the pixels in the block. In general, however, a big project is composed of a number of blocks that have different radiometric properties and are not processed simultaneously. Hence the preferred solution is to use an external low resolution image as a reference that helps to give continuity to the different blocks and also allows some radiometric coherence in time when updating cartographic series.

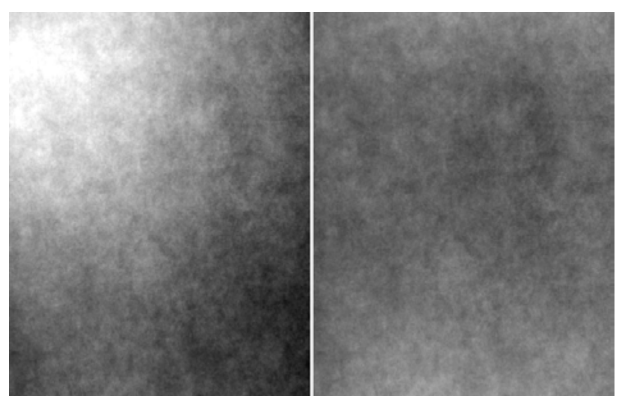

Figure 2. Radiometric Preprocessing

Figure 2 shows the mean of approximately 200 photos before and after radiometric preprocessing. The dynamic range in the left image is approximately 50 gray values out of 256 while in the right image is about 15.

3. SEAM GENERATION

Once both overlapping images have been brought radiometrically as close as possible it is time to define a seam path that separates the contributions from both images to the mosaic. The best kind of seam to search for is the one that is not perceptible by the Human Visual System (HVS). In this view, straight lines are much more recognizable by the human eye than crooked curves. In figure 3, both a straight and a crooked seam are depicted in the central image. In the lateral images the resultant mosaic between adjacent photos using one or the other seam are shown.

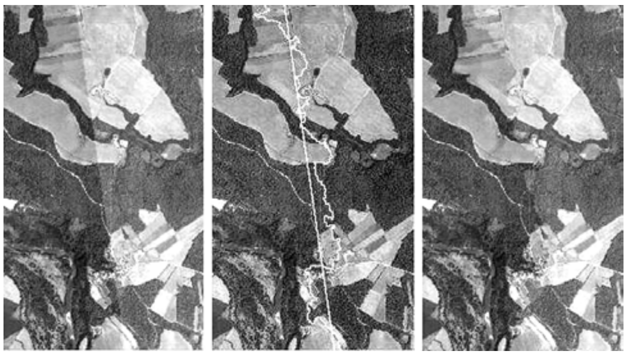

Figure 3. Straight vs. crooked seams

The search of such a seam involves some consecutive steps. First, a cost function must be defined that penalizes some paths and favor others. Then, the minimum cost that allows us to find a path crossing the image must be calculated [2]. At last, a path must be selected that reaches this cost in as few pixels as possible.

Path

A path is defined as a connected, ordered sequence of points in an image. Hereafter, 4-neighbor connectivity is assumed. A path can connect two points but also two areas.

Cost Function

A cost must be defined for every path traversing the image. This cost must quantify the idea that high-cost paths originate bad, visible seams whereas low-cost paths generate good ones.

Cumulative costs must be avoided since they tend to favor straight paths and the HVS is highly sensitive to straight discontinuities. Moreover, good seams separate *areas* with low radiometric differences rather than *pixels* with low radiometric differences. Then some kind of low pass filtering on the differences is advisable.

With all these restrictions in mind, a difference image is created from every overlapping pair of images in the form

$$\mathrm{I}_{ij}(x,y) = \underset{k=1..n}{MAX} \left| \mathrm{F}'^{k}_{i}(x,y) - \mathrm{F}'^{k}_{j}(x,y) \right| \qquad (6)$$

taking the maximum of the absolute differences in each band k. For practical reasons, the values in this image are bounded to a maximum of 127. After a mean filtering using a 5x5 window, the cost assigned to a seam path π is just the maximum value of I_{ij} over the pixels that belong to that path

$$\mathrm{E}(\pi) = \underset{(x,y)\in\pi}{MAX} \left(\mathrm{I}_{ij}(x,y)\right) \qquad (7)$$

excluding the initial and final pixels. As a consequence, any two-pixel-long path will have zero cost.

Minimum cost

The issue now is to find the minimum cost N needed to traverse the difference image from a set of possible departure pixels A to a set of possible arrival pixels B. When looking for a vertical seam, sets A and B would typically be the first and last row in the difference image. A fast algorithm has been chosen that just states the existence or not of a path with cost N without providing the pixels that form the path. A bisection approach will be used so that a result can be achieved in a maximum of $\log_2 127 \approx 7$ tries. Once the minimum cost N has been stated, a path with this cost can be obtained using a somewhat slower algorithm that generates the pixel coordinates. Usually, this kind of algorithms can provide paths that are not simple, i.e. some pixels can have more than two neighbors. If this is the case, an additional simplification algorithm must be used.

Segmentation

It is immediately obvious that this approach does not optimize the path for pixels with cost below N. Actually, when N is relatively high, bad seams are going to be found since the path traverse areas in a straightforward manner over pixels with cost below N but high enough for the seam to be noticed.

A recursive segmentation algorithm has been set up to overcome this drawback in the following way:

1. Find the minimum cost N for a path from A to B.
2. Find a particular path with this cost N.
3. Divide this path into segments using the points where cost N is attained.
4. Apply step 1 to the new segments

This algorithm comes to an end when all segments are two-pixel paths.

Shortcuts

This algorithm is refined in order to avoid that newly calculated subpaths cross old ones and eliminate, if possible, some of the pixels with maximum cost. The procedure is as follows:

- An ordered list is created containing the points in a path (subpath) with maximum cost, N.
- A new minimum cost M (M < N) is sought for a path connecting the first point with any other in the list, starting the tries with the farthest point.
- If such a cost is found, an instance of this path is obtained that substitutes the old one and the maximum-cost points in between are eliminated from the list.

As an example of the performance of the algorithm, Table 1 shows the percent of pixels with a given cost in the seam paths found for two different images before (clear) and after (shaded) the segmentation algorithm has been applied.

As desired, the algorithm tends to give priority to low cost pixels and minimize the number of times the maximum cost value is attained.

Cost	Image 1 (cost 4)		Image 2 (cost 9)	
0	**12.8 %**	**39.5 %**	**5.0 %**	**19.9 %**
1	**24.7 %**	**49.4 %**	**9.9 %**	**35.8 %**
2	**23.6 %**	**9.1 %**	**10.7 %**	**25.0 %**
3	**22.0 %**	**1.8 %**	**10.7 %**	**11.0 %**
4	**16.9 %**	**0.2 %**	**9.7 %**	**4.2 %**
5	**0.0 %**	**0.0 %**	**10.8 %**	**2.2 %**
6	**0.0 %**	**0.0 %**	**10.4 %**	**1.0 %**
7	**0.0 %**	**0.0 %**	**11.3 %**	**0.6 %**
8	**0.0 %**	**0.0 %**	**10.7 %**	**0.2 %**
9	**0.0 %**	**0.0 %**	**10.9 %**	**0.1 %**

Table 1. Cost distribution in the seam

4. ENHANCEMENTS

Several enhancements have been implemented over this general scheme that extend the possibilities of the method.

Image Selection
Sometimes, high cost areas must not only be avoided in the search of the seam but also in the final mosaic.
The most familiar example would be a photogrammetric flight over sparsely clouded terrain. *Cloud* pixels will appear immediately as a high cost area in the difference image, but this is not enough. The seam path selection must ensure that the final mosaic will not contain the cloud. The most obvious solution is to add a high cost segment that connects the cloud to the appropriate edge of the image in which appears.
A second typical example is the set of usually small stretching areas generated during the geocoding process of images over high relief areas. In this case, the scenario is often composed of two "stretching maps" – one per image – containing small but numerous zones that must be eliminated from its image. A simple graph must be created for each map that connects all "same side" stretches and these two graphs must be connected to the correct side of the difference image leaving, at the same time, space for the seam path to be traced.
As an example, figure 6 depicts an across-track seam found after detection and interconnection of stretched areas in each shot.

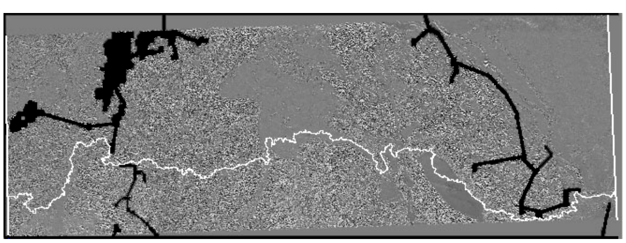

Figure 4. Across-track seam

Start and arrival areas are printed in light gray. The black areas over the seam correspond to stretched areas in the lower photo while the ones below the seam correspond to stretched areas in the upper shot. The seam is forced to pass between them so that all wrong pixels lie out of the mosaic.

Multiple seams
In a scenario as the one described before, finding a single path traversing the image from side to side can be extremely difficult if a low cost is to be maintained. Then, a way must be found to isolate a few number of these forbidden areas and generate a linear seam plus one or several closed seams to cut out these areas from the final mosaic.

5. BLOCK MANAGEMENT

Once the seam for a pair of images is found, the same method can be applied to find a seam path for every pair of consecutive images in a large block. However, the problem of connecting the individual seams to generate the polygon that defines the image pixels contributing to the final mosaic is not trivial.

In general, a block of images is composed of several flight lines (strips), each of them containing several shots. The along-strip overlapping areas suffer from perspective differences (corrected to a great extent in geocoding) but temporal decorrelation is low. So, good, easy seams are expected. Across-strip overlaps, on the other hand, can have significant changes in illumination and also temporal differences resulting in more frequent high-cost areas. Since a medium-sized project can contain a good number of blocks with large temporal gaps and even acquisition system changes, problems are magnified at block borders.

There are two main strategies for the mosaic. One is to define the useful part of a photo and consider the mosaic as a puzzle of images. In this case, 4 connected seams per photo will be necessary. Another approach is a mosaic made as a collage where the last photo to enter the mosaic block the view of the previous ones in the overlapping parts. Then, at least 2 wisely chosen, connected seams per photo are needed.

The procedure to manage and connect the different seams obtained, allowing at the same time some parallel processing, is described below. Note that the block need not have a horizontal & vertical disposition and thus the terms left, right, upper, lower are just relative to the strips irrespective of their real orientation on the ground.

Along-strip Seams
The along-strip seam for a photo is defined as the border between itself and the shot to its left. Thus, all images in a block have an along-strip seam except the leftmost one in each strip. In order to avoid seam crossings in high-overlap projects, areas to the left of its seam are eliminated from the difference image before calculating next seam. All strips can start independently at the same time.

Across-strip Seams
Every image, except the ones belonging to the first strip, has an across-strip seam that connects its along-strip seam to the right neighbor's one, defining at the same time the border with the upper photo. In addition, the common seam separating two consecutive strips must not make use of the along-strip seams. If a collage strategy is going to be applied these seams must be calculated in the same order as the corresponding images are going to enter the mosaic.

Combined Seams
It is easy to see that the restrictions above mentioned impose that the across-strip seam for a photo starts on the point where the previous one ended. But still the new seam could cross the old ones. This is avoided with the creation of the combined seam formed by the across- and along-strip seams of a photo

and the along-strip seam of the photo to the right. This seam defines an area that can be subtracted from the difference image before calculating next across-strip seam in the very same manner as we did in the along-strip case.

At the end of this process a set of combined seams is obtained that, when entered in the correct order, allows the generation of the final mosaic via a *collage*.
If intensive edition is foreseen the *puzzle* approach is preferred since automatic update of the images is greatly improved. In this case, the common across-strip seam for the lower strip is used to complete the combined seam into one or more closed polygons defining the *useful* pixels in the image.

6. RESULTS

The results presented here correspond to a real project in Catalunya, a region in northeastern Spain. 32000 km^2 (12500 mi^2) are covered in several flights at 1:30000 scale with overlapping between 30 and 50% across-track and 60-80% along-track depending on terrain ruggedness (absolute differences in height about 3000 m or 10000 feet). The whole project consists of more than 6000 photographs divided into 18 aerial triangulation blocks that will result in 1121 orthoimages containing a mean of 12 seam paths each.

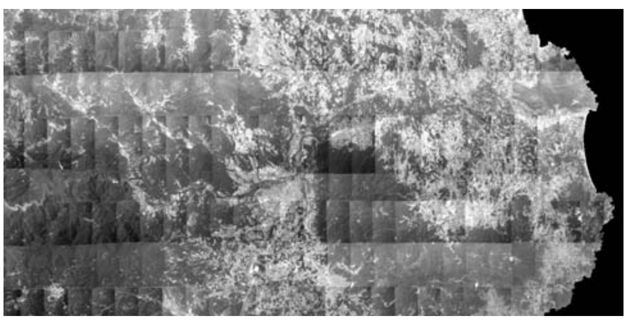

Figure 5. Original Block

Figures 5 & 6 show one of the aerial triangulation blocks containing approximately 215 photographs before and after radiometric preprocessing. Even though some radiometric artifacts still remain, it is easy to realize that the quality of the seam paths found is much higher in the second image.

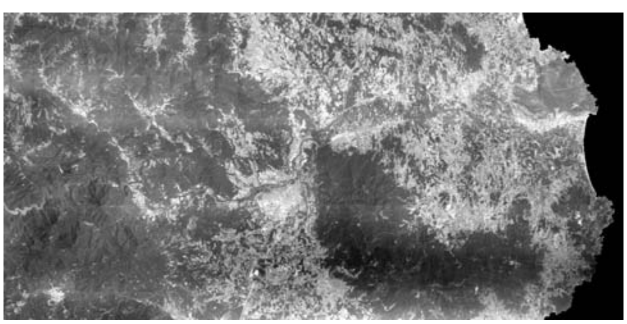

Figure 6. Block after radiometric processing

There is an extra benefit that cannot be perceived in grayscale images. Due to adaptation of the scanner to the density ranges in the film, tonal dominants are different for each photography. Since radiometric corrections deal with each band separately, in the final result these tonal differences are compensated for, rendering a continuous output.

Eight photos from the previous block have been selected to illustrate the mosaicking process. In figure 7 all the original photos have been just overlapped and the calculated seams are drawn on the mosaic.

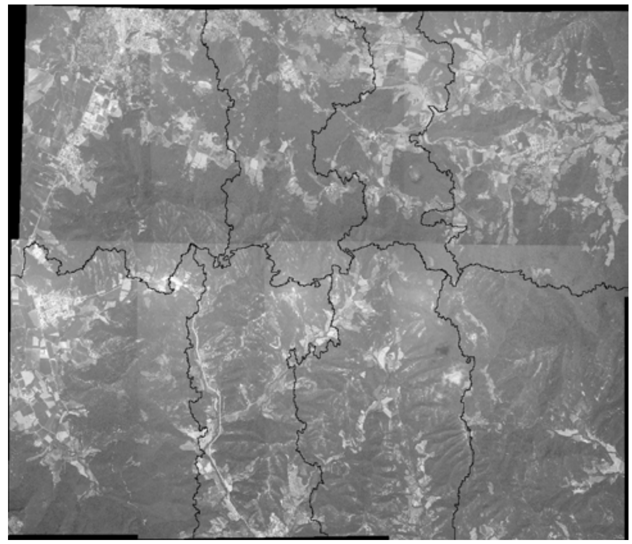

Figure 7. Eight photos with seams

Figure 8 shows the final mosaic obtained using the seams in the previous figure. The final result looks continuous and a careful inspection would reveal that some clouds have been automatically eliminated without manual intervention.

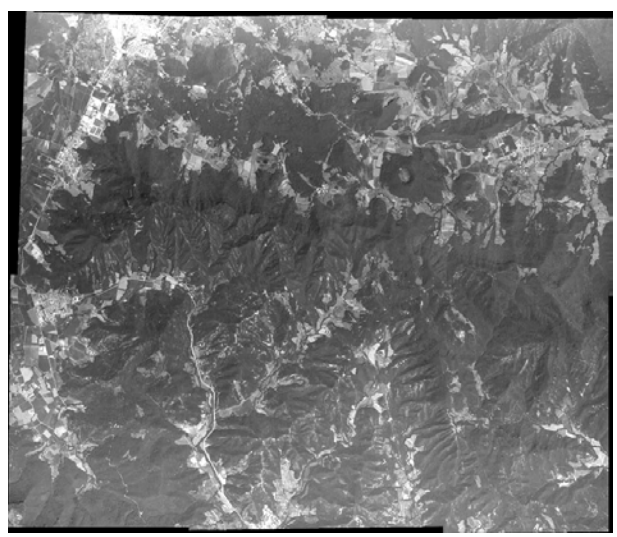

Figure 8. Final mosaic for eight photos

The mosaic in this case uses the *collage* approach. In case the number of clouds made advisable some edition processes on the final mosaic, the *puzzle* option would make subsequent operations easier.

7. FUTURE WORK

There is an additional approach to be explored in cost assignment on the difference image. Up to now, cost has been modified only to penalize areas to be avoided. However, the opposite way also looks promising. For instance, it is well known from operator's experience that seams traced following natural gradients in the image are much less visible than paths traversing flat areas, no matter how low the difference. So, if an edge detection algorithm is applied to both images, common gradients could be used to lower the cost of these pixels in order to favor the paths following these gradients.

Sometimes, a perfect set of seams is not found automatically. Even if it is, later image editions can affect the calculated seams. Then, efforts have to be made in near future to develop interactive, graphic edition software tools designed to handle the polygons describing the profitable part of each photograph. These tools would take into account, for a specific polygon, the neighbor ones in order to maintain the puzzle consistency.

8. REFERENCES

[1] Knopp,D., Gerhard,J., 1998. Automating Image Mosaic Production. Earth Observation Magazine 7(2). http://www.eomonline.com/Common/Archives/1998feb/98feb_knopp.html (accessed April 21, 2005).

[2] Fernández, E., Garfinkel,R., Arbiol,,R., 1998. Mosaicking of Aerial Photographic Maps via Seams Defined by Bottleneck Shortest Paths. Operations Research 46(3).

[3] Kerschner, M., 2001. Seamline detection in colour orthoimage mosaicking by use of twin snakes. ISPRS Journal of Photogrammetry & Remote Sensing 56, pp. 53-64.

[4] Iqbal, M., 1983. An Introduction to Solar Radiation. Academic Press, Toronto, Ontario, Canada, 390 pp.