

Order Independent Image Compositing

Conrad Bielski and Pierre Soille

Land Management Unit of the Institute for Environment and Sustainability
Joint Research Centre of the European Commission
T.P. 262, I-21020 Ispra (Va), Italy
{Conrad.Bielski,Pierre.Soille}@jrc.it

Abstract. Image compositing is defined as the assembling of two or more overlapping images into a single image. Recently, a morphological image compositing algorithm was proposed that automatically positions seam lines along salient image structures. This algorithm requires that all images used in the compositing procedure be held in computer memory. Hence, when composing large images such as satellite imagery acquired over a large region, the maximal capacity of random access memory of 32 bit computers is rapidly exceeded. In this paper, we present a parallel algorithm producing the same results whilst requiring only one input image at a time to be held in memory. The algorithm is illustrated for the automatic production of pan-European mosaics of Landsat images.

1 Introduction

Image mosaicing can be defined as the registration of two or more images which are then combined into a single image. This procedure is applied when more than one image or scene was acquired in order to cover the entire study area. This is a routine image processing task in remote sensing, astronomy, microscopy, medical imaging and, more recently, digital photography. Prior to image mosaicing, the scenes themselves must be geometrically registered, i.e. all images must be aligned geometrically so that differences between viewpoints and/or sensors are corrected. After geometric registration, the images are assembled so as to create a single scene having a field of view corresponding to the union of the field of views of the input scenes. The assembling of the individual registered images must address the problem of determining how the pixels belonging to regions visible from more than one scene (i.e., overlapping domains) should be represented. This problem is sometimes referred to as *image compositing* [1]. Seam lines that define the extent of the imagery used in image compositing in many cases are produced manually [2,3] while many remote sensing software solutions produce visible seams given the user-defined choice of which image should be on top, see Fig. 1.

Recently, the second author of this paper proposed a morphological [4] image compositing algorithm to automatically position seam lines along salient image structures [5]. The proposed morphological image compositing algorithm automatically adapts to the morphology of the image objects by following their




Fig. 1. Two co-registered Landsat TM scenes of Ireland are shown on the left. Basic compositing simply chooses which scene to place on top thereby creating a seam line that follows the data extent of the top image.

boundaries. This is achieved through marker controlled segmentation of the gradient intensity where the marker is the internal border of the overlapping regions and the gradient intensity is the point-wise minimum of gradients based on the overlapping input images. The number of overlapping images also drives the seam generation sequence because regions with fewer overlapping images must be processed before the seam line can continue to regions with a greater number of overlapping images.

The order independent image compositing algorithm presented in this paper was developed because the above described compositing technique could not be applied for generating very large mosaics. Indeed, it requires a single image of the size of the final mosaic to be held in memory which is impossible when dimensions exceed the available RAM size (4 GB for our system). When such processing is required in the field or on the fly (e.g., for web mapping services), portable or desktop devices cannot handle such large imagery. This problem is solved by proposing a scheme allowing for the processing of one input image at a time. However, this introduces the problem of order dependence since the final mosaic depends on the processing order of the input images. We address this problem by presenting an order independent algorithm. That is, the same results are obtained whatever sequence the images within the dataset are processed. Consequently, the resulting order independent algorithm is suitable for parallel processing. In Sec. 2, the method details are explained while Sec. 3 presents the application of this algorithm to Landsat Thematic Mapper (TM) imagery.

2 Methodology

Prior to image compositing, the input images must first be co-registered in a common reference coordinate system and resampled at the same spatial resolution. This is a usual procedure when analysing remotely sensed images where scenes are orthorectified, reprojected, and gridded at the node of a given sampling grid defined by a given projection of the spheroid on to the plane (unless geodetic coordinates with a given sampling increment along parallels and meridians are used). Gridding involves the application of an interpolation function such as that defined by the cubic convolution [6].

We assume that n input co-registered images denoted by f_1, \dots, f_n are available and that neither of them is a subset of the union of the others: for all $i \in \{1, \dots, n\}$, $\mathcal{D}_i \not\subseteq \bigcup_{j \mid j \neq i} \mathcal{D}_j$, where \mathcal{D}_i denotes the definition domain of image f_i . Indeed, in case the definition domain of an image falls within the union of all the other image definition domains, the former image will not contribute to the final result so that its inclusion in the input dataset is useless. We aim at defining a unique value for each pixel of the composed image f whose definition domain equals the union of all input definition domains: $\mathcal{D}_f = \bigcup_{i=1}^n \mathcal{D}_i$. The principle of the order independent compositing algorithm is first presented and then pseudo-code describing its implementation is provided.

2.1 Principle

Each input image f_i contains an arbitrary number of lines and columns corresponding to a rectangular frame. Unfortunately, the scene data extent does not necessarily correspond to the frame extent due to the acquisition mode and the orthorectification process, see example in Fig. 1. Therefore, \mathcal{D}_i refers to the domain of definition of the i th image in the sense of the subset of the image frame where data values are present. Figure 2 (left side) shows a case with 4 input images where, for ease of representation, the rectangular frames match the image definition domains. We then construct an *overlap matrix*. This is a symmetrical $n \times n$ indicator matrix indicating whether the definition domains of an arbitrary image pair overlap or not:

$$m_{i,j} = \begin{cases} 1, & \text{if } \mathcal{D}_i \cap \mathcal{D}_j \neq \emptyset, \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

For example, the overlap matrix corresponding to Fig. 2 (left side) is as follows:

$$\begin{pmatrix} 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 \end{pmatrix}.$$




Fig. 2. Left: Four overlapping images. Right: the overlap levels ranging from 1 (light grey for no overlap) to 3 (black for 3 images overlapping).

In general, owing to the arbitrary shape of the definition domains \mathcal{D}_i , the overlap matrix cannot be constructed on the sole knowledge of the frame size and positioning of the upper left corner (or any other reference point) in the reference coordinate system. Hence, for all pairs of images whose frames intersect, we need to compute the intersection between their actual definition domains to assess whether these intersect or not.

Beyond the information available through the overlap matrix, the number of overlapping images (called *overlap level* hereafter) for any given pixel must be known. Indeed, seam lines are detected in an ordered fashion, starting with regions with the least overlap level greater than 1 and proceeding to the subsequent level until the maximum number of overlap is reached. For example, Fig. 2 (right side) shows the domains of equal overlapping levels with a specific grey shade: light grey for overlap level equal to 1 (no overlap), dark grey for overlap level 2, and black for overlap level 3.

We then proceed as follows. The overlap matrix is scanned row by row. The index of the row defines the current *anchor image*. Actual processing is restricted to the region defined by its frame (we assume that the definition domain is buffered by the image frame). Within this region, the least overlap level greater than 1 defines the *current overlap level* denoted by h_{crt} . We also track whether an overlap level higher than the current level occurs. This is used to determine whether the overlap matrix needs to be scanned again later on. The morphological compositing routine [5] is then called while restricting its effect to the processing of those regions whose overlap level is equal to h_{crt} . The routine assigns each pixel of these regions to a unique source image so that the definition domain of the anchor image and those of the images intersecting it can be updated accordingly. This update can only remove some parts of these input definition domains since it concerns regions where more than one image was competing for the same domain. Order independence is achieved because identical results are obtained whatever order the anchor images are processed.

2.2 Implementation

In essence, the above principle is not too complex, however its implementation requires data to be carefully updated on the fly. Information with regards to the anchor image overlap level and changes to the definition domain with increasing overlap level needs to be taken into account because these changes affect the seam line generation.

Code implementation was divided into two steps: generating the overlap matrix and processing the overlap matrix. The overlap matrix is generated based on the list of available images to be mosaiced which can be of any length. List order determines the location of the anchor image within the overlap matrix, i.e., image list item 2 will be anchor image f_2 within the overlap matrix.

The second step involves the actual compositing of the mosaic based on the overlap matrix. Processing of anchor and overlap images follows each matrix row. Each i th matrix row was linked to the i th anchor image thus when dealing with the i th row it should also be understood that one is dealing with f_i as anchor

image. An auxiliary image g whose frame equals that of the current anchor image is used to calculate the current overlap level h_{crt} . We also assume that the actual definition domains are stored on disk as binary images where a pixel value of 1 indicates whether a pixel belongs to it. Pseudo-code for the order independent compositing is summarised hereafter:

1. DO
2. flag \leftarrow false
3. FOR $i \leftarrow 1$ to n // scan the overlap matrix row by row
4. $g = \sum_{\{j \mid m_{i,j}=1\}} \mathcal{D}_j$ // pixel-wise addition (restricted to domain \mathcal{D}_i)
5. $h_{\text{max}} = \max_{\mathbf{x}} g(\mathbf{x})$
6. IF $h_{\text{max}} = 1$ THEN GOTO 3 // no more overlap
7. $h_{\text{crt}} \leftarrow \min_{\{\mathbf{x} \mid g(\mathbf{x}) > 1\}} g(\mathbf{x})$
8. IF $h_{\text{max}} > h_{\text{crt}}$ THEN flag \leftarrow true END IF
9. apply image compositing [5] to anchor image frame for overlap level = h_{crt}
10. update definition domains of all images such that $m_{i,j} = 1$
11. END FOR
12. WHILE flag = true

The implementation recursively goes through all overlap matrix rows until there are no more overlapping images. This is necessary because the current overlap level h_{crt} is unknown *a priori* and therefore must be evaluated each time an overlap matrix row is being processed (line 7). Note that different anchor images will have different overlap levels that are not necessarily sequential.

Success of this order independent strategy depends on the precise updating of the image definition domains. Updated image definition domains are based on two results: the image compositing procedure for the current overlap level and the re-insertion of areas with overlap levels greater than the current level. For the anchor image, the image compositing results always fall within its own definition domain and overlap levels greater than the current overlap level are re-inserted as required. The overlap images, on the other hand, can only be updated within the definition domain of the anchor image because information with respect to image definition domain overlap are only known within the definition domain of the anchor image for any given overlap matrix row. This will be more evident with the example shown in the next section. When the flag is not reset to true when scanning the overlap matrix (line 8), image compositing has been applied to the entire mosaic and the program stops.

3 Application

A sample application of the order independent image compositing procedure is presented based on Landsat Thematic Mapper (TM) imagery over the island of Sardegna, Italy. The imagery was taken from the Image2000 [7] dataset which covers all European Union member states. These images were acquired during the summer season over the 1999 to 2001 time frame. The island of Sardegna was chosen as the sample region to clearly present the results because only

five Landsat TM images are required to cover the entire island. However, the algorithm can be handle any number of scenes. For example, the composition of the 39 Landsat TM images covering Italy by the order independent image compositing procedure was obtained with no additional random access memory requirements.

The overlap matrix for Sardegna is presented in figure 3. Note that Landsat scenes are in no particular spatial order and overlap between definition domains is readily apparent: anchor images f_2 and f_3 have 2 images with overlapping definition domains, anchor images f_1 and f_4 have 3 images with overlapping definition domains and anchor image f_5 has 4 images with overlapping definition domains. Image compositing can only proceed within the image definition domain which is the region where actual remote sensing measurements are available. For the orthorectified Landsat TM scenes, this region is the rotated square within the image. Based on the above overlap matrix, image compositing was

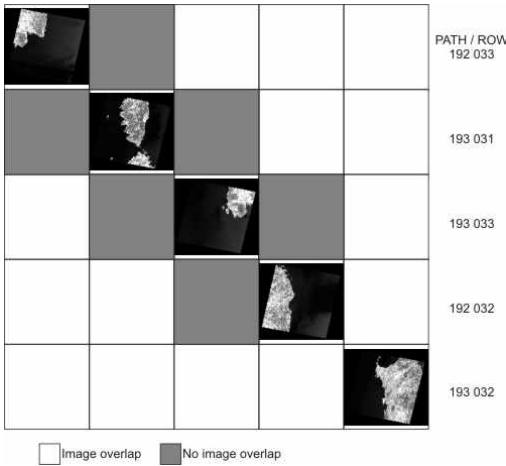


Fig. 3. Overlap matrix for the island of Sardegna, Italy. Five Landsat TM images cover this region (image path and row are presented to the right of the matrix) with band 5 (near infrared) presented in the diagonal.

applied. The step-by-step results of the order independent procedure are presented in figure 4. On the left is presented the actual number of image overlaps for this dataset. It shows that the greatest overlap level is 3 (white regions). During the order independent procedure, the dark grey regions (overlap level=1) will never be processed because there is only a single image available in those areas. The light grey regions indicate areas where the overlap level is equal to 2. These regions must first be processed before image compositing can occur for regions where there are 3 images overlapping (i.e. overlap level = 3). This figure is only shown as a reference because in practice the overlap level is computed on the fly based on the images found within any given overlap matrix row thus ensuring order independence. The middle diagram in figure 4 presents the updated image definition domains after image compositing has been applied to all regions where the overlap level is 2. From the overlap level mosaic there should be three regions that persist in the updated image definition domains. These




Fig. 4. Left: overlap level map for the Sardegna dataset. Middle: updated image definition domains for image overlap level 2. Right: final image definition domains, i.e., when the maximum overlap level which is equal to 3 is reached.




Fig. 5. The result of applying the order independent image compositing procedure. On the left is the mosaic of the final image definition domains and on the right is the mosaic based on Landsat TM band 5 imagery.

rectangular regions where the overlap level is greater than 2 can be seen in the corners of the images because they persist throughout the processing of overlap level 2. Note however that in the final row, $\mathcal{D}_{4,4}$ has been completely updated. This result was already mentioned in the principle section and occurs because all regions related the last row of the overlap matrix have already been processed. The rightmost diagram presents the final image definition domains. In this case, all definition domains fit perfectly together like a puzzle and therefore there is no more image overlap anywhere within the mosaic: $D_f = \cup_{i=1}^n \mathcal{D}_i$ and for all $i \neq j$, $\mathcal{D}_i \cap \mathcal{D}_j = \emptyset$ since the \mathcal{D}_i 's have been updated through image composition. The final image definition domains are then used to create the image

mosaic presented in figure 5. Inland seam lines are barely visible even though no radiometric equalisation was performed between the different scenes taken at different times. In the sea however, seam lines are noticeable because of the lack of salient structures. Automatic compositing in these homogeneous regions is therefore primarily driven by low frequency fluctuations caused by variable atmospheric and sea conditions at the time of acquisition.

4 Conclusion and Perspectives

Efficient order independent image compositing of large images can be achieved based on the proposed algorithm. In addition, the algorithm enables the user to add further functionalities such as the minimisation of specific structures occurring within the overlap regions. For example, cloud and shadow complexes were removed when compositing the satellite images shown in this paper. For conciseness this was not presented here but will be detailed in an extended version. The algorithm also directly applies to 3-dimensional imagery such as those produced by imaging tomography devices. Given the size of the 3-D input images, order independent processing on an image by image basis is an asset.

References

1. Porter, T., Duff, T.: Compositing digital images. *Computer Graphics* 18 (1984)
2. Hummel-Miller, S.: A digital mosaicking algorithm allowing for an irregular join line. *Photogrammetric Engineering and Remote Sensing* 55 (1989) 43–47
3. Nunes de Lima, V., ed.: *IMAGE 2000 and CORINE Land Cover 2000 —Products and Methods—*. European Commission, Joint Research Centre (2005) In Press.
4. Soille, P.: *Morphological Image Analysis: Principles and Applications*. corrected 2nd printing of the 2nd edn. Springer-Verlag, Berlin and New York (2004)
5. Soille, P.: Morphological compositing. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2005) [Submitted].
6. Keys, R.: Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 29 (1981) 1153–1160
7. Nunes de Lima, V., Peedell, S.: IMAGE2000-The European spatial reference to support environmental information. In: *Proceedings of 18th International Conference Informatics for Environmental Protection (EnviroInfo)*, Geneva, 21–23 October. Volume II., Genève, Editions du Tricorne (2004) 268–276 [Invited session].