

SEAMLESS IMAGE MOSAICKING VIA SYNCHRONIZATION

E. Santellani, E. Maset*, and A. Fusiello

DPIA, University of Udine, Via delle Scienze, 208 - 33100 Udine, Italy
maset.eleonora@spes.uniud.it, andrea.fusiello@uniud.it

Commission II, WG II/1

KEY WORDS: Image mosaicking, Synchronization, Image alignment, Colour correction, Seams optimisation

ABSTRACT:

This paper proposes an innovative method to create high-quality seamless planar mosaics. The developed pipeline ensures good robustness against many common mosaicking problems (e.g., misalignments, colour distortion, moving objects, parallax) and differs from other works in the literature because a global approach, known as *synchronization*, is used for image registration and colour correction. To better conceal the mosaic seamlines, images are cut along specific paths, computed using a Voronoi decomposition of the mosaic area and a shortest path algorithm. Results obtained on challenging real datasets show that the colour correction mitigates significantly the colour variations between the original images and the seams on the final mosaic are not evident.

1. INTRODUCTION

Aligning and stitching together multiple images is a classical problem in photogrammetry (Pan et al., 2009) and computer vision (Uyttendaele et al., 2001, Szeliski, 2006, Brown and Lowe, 2007). Image mosaicking finds application in various scenarios, ranging from satellite or aerial imagery (Du et al., 2008), street-view panoramas (Li et al., 2016) or video stabilization (Hansen et al., 1994), to name a few.

Since there are many technical difficulties in taking a photo with a very large *field of view* (FOV), often the only practical solution is to acquire multiple images with smaller FOV and merge them together. *Image mosaicking* can be therefore defined as the process of stitching different photos of the same scene in a single wide image. The result should be as natural as possible, ideally indistinguishable from a real photo that covers the entire scene. However, deviation from planarity or photometric nuisances between the images can lead to a mosaic in which seams are evident.

The goal of this paper is to develop an innovative procedure to create seamless mosaics exploiting a global approach, known as *synchronization* (Singer, 2011). Starting from the geometric and radiometric information between pairs of overlapping images, the synchronization method is able to simultaneously estimate global homographies and colour corrections for all the images, avoiding the errors that accumulate when adding an image at a time to the mosaic. Finally, in order to minimize and conceal the seams that can still be visible after the global colour correction (e.g., due to parallax and moving objects), the cutting paths are determined using a Voronoi tessellation and optimized with the Dijkstra's algorithm (Dijkstra, 1959).

The experimental validation was conducted on datasets composed by tens of images, acquired by a helicopter or an Unmanned Aerial Vehicle (UAV). Although the scenes are not perfectly planar, there are moving objects and strong illumination and intensity differences, the obtained mosaics appear homogeneous, without artifacts, and the seams are well concealed.

*Corresponding author

The paper is organized as follows. In the next section, the literature on mosaic generation is reviewed. Section 3 introduces the synchronization problem and illustrates the solutions for homography and affinity synchronization, that are used in the mosaicking process for image registration and colour correction, respectively. In Sec. 4 the proposed procedure is described in detail, while Sec. 5 shows the results. Finally, Sec. 6 draws the conclusion.

2. STATE OF THE ART

Aligning and stitching images into seamless mosaics is a procedure usually composed by three main steps: image registration, colour correction and blending. Image mosaicking can be performed independently from (and prior to) the structure-from-motion and dense matching phases, that are instead required to generate orthophotos. The goal of mosaic creation is, in fact, to visualize a wide area on a single image under perspective projection, whereas orthophotos are orthographic projections.

In the last decades several methods for automatic image mosaicking appeared in the literature, proposing a complete pipeline for the final mosaic generation (Davis, 1998, Marzotto et al., 2004, Brown and Lowe, 2007) or focusing the attention on the optimisation of one of the previously cited steps (Schroeder et al., 2011, Oliveira et al., 2011, Li et al., 2016).

Algorithms for image alignment can be divided into two broad categories (Szeliski, 2006): direct (pixel-based) and feature-based. Direct methods exploit the entire image data, thus providing very accurate registration but requiring at the same time a close initialization. Feature-based algorithms, instead, do not require initialization and can be computationally less expensive. Moreover, since the introduction of invariant features (e.g. SIFT, (Lowe, 2004)) and robust feature matching, feature-based methods have gained increasing attention and are nowadays widely used. (Brown and Lowe, 2007) proved that, formulating stitching as a multi-image matching problem and using invariant local features to find matching between the images, lead to a method insensitive to the ordering, orientation, scale and illumination of the input images.

To obtain a clean, pleasant looking mosaic, a robust alignment process must be followed by colour correction. Neighbouring images can indeed show colour and appearance differences due to exposure level variation, changes in lighting condition and different camera settings. Colour correction methods proposed in the literature can be divided into model-based parametric approaches and non parametric ones (Xu and Mulligan, 2010). The former assume that the relation between two images can be described by a colour transfer function, whereas the latter consider no particular parametric format of the colour mapping function and typically use a look-up table to directly record the mapping of the colour levels. (Xu and Mulligan, 2010) evaluated the performance of various colour correction approaches, showing how the gain compensation method by (Brown and Lowe, 2007) and the local colour transfer approach by (Tai et al., 2005) are fast, effective and general (applicable in various scenarios).

Even after colour correction, seams and artifacts can be visible in the mosaic. Image blending techniques are able to conceal the colour differences along the seamlines but cannot handle residual geometric misalignment deriving from parallax and moving objects. For these reasons, it is mandatory to compute optimal seamlines that avoid crossing overlap regions with high image discrepancies. (Davis, 1998) used the Dijkstra's algorithm (Dijkstra, 1959) to compute the best cutting path dividing overlapping regions, segmenting the mosaic into disjoint regions and sampling pixels in each region from a single source image. (Uyttendaele et al., 2001) proposed a weighted vertex cover algorithm in order to remove effects caused by moving objects and (Li et al., 2016) formulated the seamline optimisation as a unified graph cuts energy minimization problem, concealing the image parallax in the resulting mosaic.

As described in the following sections, we propose an alternative method to image mosaicking, based on synchronization (Singer, 2011), to globally align images and compensate for their colour differences, in order to achieve homogeneous, high-quality planar mosaics.

3. SYNCHRONIZATION

Consider a network of nodes where each node is characterized by an unknown state. Suppose that pairs of nodes can measure the ratio (or difference) between their states, that are typically represented by elements of a group Σ . The goal of *synchronization* (Singer, 2011) is to estimate the unknown states from the pairwise measures. Several instances of synchronization have been studied in the literature, which correspond to different instantiations of Σ . Among them, it is worth citing $SE(d)$ for *rigid-motion synchronization* (Arrigoni et al., 2016b), $SL(d)$ for *homography synchronization* (Schroeder et al., 2011) and $Aff(d)$ for *affine matrix synchronization*. Please note that $SE(d)$, $SL(d)$ and $Aff(d)$ are all subgroups of $GL(d)$. In this paper the attention is focused on synchronization over $SL(3)$, that will be applied for image registration, and over $Aff(1)$ for colour correction.

In order to formally define the problem and its solution, let Σ be a group and let $*$ denote its operation. Suppose that the pairwise relations between the index pairs $(i, j) \subseteq \{1, \dots, n\} \times \{1, \dots, n\}$ are known, and refer to them as z_{ij} . *Synchronization* can be formulated as the problem of recovering $x_i \in \Sigma$ for $i = 1, \dots, n$ such that the following *consistency constraint* is satisfied

$$z_{ij} = x_i * x_j^{-1} \quad (1)$$

The solution is defined up to a global (right) product with any group element, i.e., if $x_i \in \Sigma$ satisfies (1) then also $x_i * y$ satisfies (1) for any (fixed) $y \in \Sigma$.

If the known pairwise measures are noisy, the consistency constraint cannot be satisfied exactly. Thus, as shown in Fig. 1, the searched solution is the one that minimizes the *consistency error*:

$$\epsilon(x_1, x_2, \dots, x_n) = \sum_{(i,j)} \delta(z_{ij}, x_i * x_j^{-1}) \quad (2)$$

where $\delta : \Sigma \times \Sigma \rightarrow \mathbb{R}^+$ is a metric function for Σ (Arrigoni et al., 2016a).

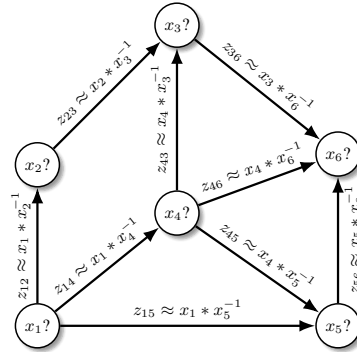


Figure 1. The synchronization problem. Each node is characterized by an unknown state and measures on the edges are ratios of states. The goal is to compute the states that best agree with the measures.

3.1 Synchronization over $(GL(d), \cdot)$

In this section we consider the synchronization problem over the General Linear Group $GL(d)$, which is the set of all $d \times d$ invertible matrices, where the group operation $*$ is matrix multiplication and $1_\Sigma = I_d$. Let $X_i \in \mathbb{R}^{d \times d}$ and $Z_{ij} \in \mathbb{R}^{d \times d}$ denote the matrix representations of $x_i \in \Sigma$ and $z_{ij} \in \Sigma$, respectively. Using this notation, Eq. (1) rewrites $Z_{ij} = X_i X_j^{-1}$.

Let us collect the unknown group elements and all the measures in two matrices $X \in \mathbb{R}^{dn \times d}$ and $Z \in \mathbb{R}^{dn \times dn}$ respectively, which are composed of $d \times d$ blocks, namely

$$X = \begin{bmatrix} X_1 \\ X_2 \\ \dots \\ X_n \end{bmatrix}, \quad Z = \begin{bmatrix} I_d & Z_{12} & \dots & Z_{1n} \\ Z_{21} & I_d & \dots & Z_{2n} \\ \dots & \dots & \dots & \dots \\ Z_{n1} & Z_{n2} & \dots & I_d \end{bmatrix} \quad (3)$$

If not all the pairwise measures Z_{ij} are available, the input matrix becomes $Z_A := Z \circ (A \otimes \mathbf{1}_{d \times d})$, where \circ denotes the Hadamard product, A is the adjacency matrix and the Kronecker product with $\mathbf{1}_{d \times d}$ is required to match the block structure of the measures. The $n \times n$ adjacency matrix is constructed as follows: $A_{ij} = 1$ if the pairwise measure Z_{ij} exists, $A_{ij} = 0$ otherwise. Accordingly, the consistency constraint writes

$$Z_A = (X X^{-b}) \circ (A \otimes \mathbf{1}_{d \times d}) \quad (4)$$

where $X^{-b} \in \mathbb{R}^{d \times dn}$ denotes the block-matrix containing the inverse of each $d \times d$ block of X .

It can be shown (Arie-Nachimson et al., 2012) that

$$Z_A X = (D \otimes I_d) X \quad (5)$$

thus an estimate of X is represented by the eigenvectors of $(D \otimes I_d)^{-1} Z_A$ corresponding to the d largest eigenvalues, where D is the degree matrix defined as $D = \text{diag}(A \mathbf{1}_{n \times 1})$. This is also called the spectral solution.

3.2 Synchronization over $SL(d)$

Consider now the Special Linear Group $SL(d)$, that is the set of $d \times d$ matrices with unit determinant

$$SL(d) = \{R \in \mathbb{R}^{d \times d} \text{ s.t. } \det(R) = 1\}. \quad (6)$$

Synchronization over $SL(3)$ corresponds to the *homography synchronization* problem. Since $SL(d)$ is a subgroup of $GL(d)$, the problem can be addressed via the spectral solution, which computes the top d eigenvectors of $(D \otimes I_d)^{-1} Z_A$, that are collected in a $dn \times d$ matrix U . In order to obtain elements of $SL(d)$ from U , each $d \times d$ block in U , denoted by U_i , must be scaled to unit determinant (Schroeder et al., 2011), which can be done by dividing U_i by $\sqrt[d]{\det(U_i)}$. However, if $\det(U_i)$ is negative and d is even, real roots do not exist; in this case the determinant can be always made positive by exchanging two columns of U .

3.3 Synchronization over $Aff(d)$

Let us finally consider the Affine Group $Aff(d)$, that is the set of invertible affine transformations in d -space, which admits a matrix representation through $(d + 1) \times (d + 1)$ matrices

$$Aff(d) = \left\{ \begin{bmatrix} M & \mathbf{v} \\ \mathbf{0}^T & 1 \end{bmatrix}, \text{ s.t. } M \in \mathbb{R}^{d \times d}, \mathbf{v} \in \mathbb{R}^d \right\}. \quad (7)$$

$Aff(d)$ is a subgroup of $GL(d + 1)$, therefore the synchronization problem can be solved by computing the top $d + 1$ eigenvectors of $(D \otimes I_{d+1})^{-1} Z_A$. Since this approach leads to an algebraic solution, it does not enforce constraints that matrices in $Aff(d)$ should satisfy.

Specifically, the output matrix U will not have vector $[\mathbf{0}_{1 \times d} \ 1]$ in rows multiple of $d + 1$. In order to recover X from U it is sufficient to choose a different basis for the resulting eigenvectors that satisfies such constraint, which can be found by taking a suitable linear combination of the columns of U , as explained in (Arrigoni et al., 2016b). More precisely, let $F \in \mathbb{R}^{n \times (d+1)n}$ be the 0/1-matrix such that $FU \in \mathbb{R}^{n \times (d+1)}$ consists of the rows of U with indices multiple of $d + 1$. The coefficients $\mathbf{a}, \mathbf{b} \in \mathbb{R}^{d+1}$ of the linear combination are solution of

$$FU\mathbf{a} = \mathbf{0}_{n \times 1}, \quad FU\mathbf{b} = \mathbf{1}_{n \times 1} \quad (8)$$

where the first equation has a d -dimensional solution space. Let $\mathbf{a}_1, \dots, \mathbf{a}_d$ be a basis for the null-space of FU . Thus X is recovered as $X = U[\mathbf{a}_1, \dots, \mathbf{a}_d, \mathbf{b}]$.

In the presence of noise, Equation (8) is solved in the least-squares sense. Then, such a solution is projected onto $Aff(d)$ by forcing the rows multiple of $d + 1$ to $[\mathbf{0}_{1 \times d} \ 1]$.

4. PROPOSED METHOD

The novel algorithm proposed in this paper tries to overcome some common issues in mosaic generation (e.g. misalignments, colour correction, moving objects) thanks to the use of synchronization and the search for an optimal cutting path between overlapping images. The entire process is summarized in Fig. 2 and described in detail in the following paragraphs.

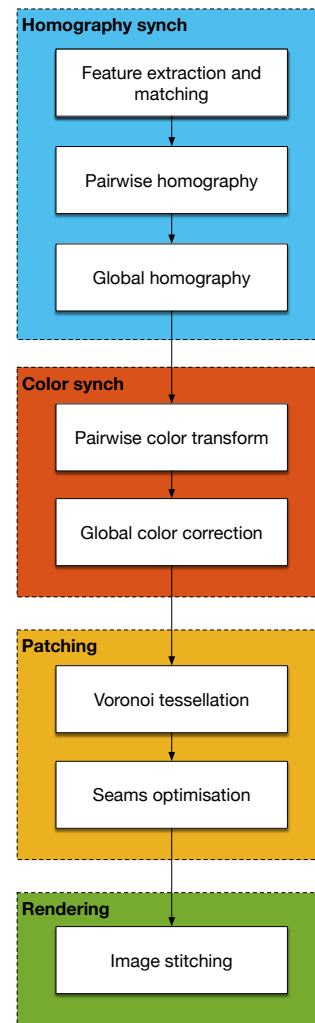


Figure 2. Flowchart of the proposed method.

4.1 Image alignment

The first step of the proposed procedure is to extract features from all the images (e.g. SIFT features, (Lowe, 2004)) and match them. A robust feature matching algorithm should be used in order to avoid wrong matches that can cause strong misalignments between the images. For this reason, the method proposed in (Maset et al., 2017) has been chosen. Starting from the correspondences between pairs of views, it jointly updates them so as to maximize their consistency.

Pairwise homographies are then robustly estimated using RANSAC (Fischler and Bolles, 1981), computing image transformation parameters through the Direct Linear Transformation (DLT) method (Abdel-Aziz and Karara, 1971). A possible solution to project all the images in the same reference system for mosaic generation is to compose relative transformations multiplying the obtained pairwise homographies. However, this approach accumulates error at each successive multiplication. To solve this problem, synchronization over $SL(3)$ (see Sec. 3.2) is applied, converting in this way pairwise homographies into absolute ones. This guarantees that all relative information are considered simultaneously, minimizing misalignment errors among the whole dataset.

To improve the accuracy of the synchronization process, a weighting factor can be assigned to each pairwise homography, that describes its reliability. In practice, the unitary elements of the ad-

gency matrix A contained in Eq. 4 are replaced by the estimated weights. In the proposed procedure, these weights are assumed to be proportional to the area of the convex hull that contains the features matched in each image pairs.

4.2 Colour correction

Changes of the illumination conditions, different camera settings and vignetting are some of the causes that make the seams of the mosaic visible, even when the scene is planar, the images are sharp and the alignment is perfect. Colour variations between overlapping images should be modelled by a non-linear function and often involve the three colour channels simultaneously. However, the simplified approach that considers the RGB channels independently and that models the transformation with an affinity proved to work well. Thus, in the proposed method the relation between the three colour channel of adjacent images (i, j) is assumed to be an affine transformation, that can be written in matrix form as

$$\begin{bmatrix} C \\ 1 \end{bmatrix}_i = \begin{bmatrix} a_c & b_c \\ 0 & 1 \end{bmatrix}_{i,j} \cdot \begin{bmatrix} C \\ 1 \end{bmatrix}_j \quad (9)$$

where C is in turn R,G, or B. Formulating the problem in this way corresponds to estimating the parameters of three affine transformations between each pair of overlapping images, that have to be then composed in order to compute a global colour correction for each single image. It is easy to see that this problem can be solved via the synchronization over the *Affine Group* $Aff(d)$, described in Sec. 3.3.

In the presence of small residual misalignments, a pixel-based method used to estimate the pairwise affine transformations can lead to inaccurate results. An alternative robust approach, adopted in this paper, consists in exploiting the histograms of the overlapping area computed for both images. The parameters of the affine transformation are computed as the angular coefficient and intercept of the straight line that fits the plot of one cumulative histogram versus the other cumulative histogram (Cox et al., 1996).

Once all the relative affine transformations have been computed for each colour channel, the absolute ones can be retrieved via synchronization, as done for the homographies. A weighting matrix can be introduced, where the weights are proportional to the overlapping area size, in order to give more confidence to the most reliable pairwise colour transformation. Please note that synchronization retrieves absolute affine transformation, up to a global one. This degree of freedom can be fixed by choosing one image that does not undergo colour correction. The unaltered image can be identified automatically as the one that has the best colour balance, or it can be defined by the user.

An example of the results achieved with this approach is represented in Fig. 3.

4.3 Voronoi tessellation and seams optimisation

Even if image alignment and colour correction produce optimal results, seams can still be visible on the final mosaic due to, e.g., parallax or moving objects. We therefore propose an approach that first reduces the seamlines total length and then conceals the remaining ones.

First of all, it is advisable to remove redundant images, i.e., images completely covered by the adjacent ones. Using all the images can indeed generate a mosaic composed by many little




Figure 3. Mosaic before (a) and after (b) colour correction.

patches, increasing at the same time the total length of the seamlines. This issue can be faced searching for a subset of images with the property that no one is completely covered by the union of the others. An iterative greedy approach can be followed to discard redundant elements, considering an image at a time and evaluating if its projection in the mosaic reference frame is completely covered by the projection of the other images. If so, the current image is discarded and the dataset is updated. The process is repeated until all the remaining elements are verified. The drawback of this procedure is that it is order-dependent and does not guarantee that the minimum subset is kept. The process can be driven by arranging the dataset in a convenient way, since the first images analysed are more likely to be discarded. A possible choice is to sort the images according to their colour balancing or to the alignment error.

The previously described step reduces the number of tiles that compose the mosaic. To find a method that further reduces the overall length of the seams, one can start from the following considerations. The *honeycomb conjecture* (Hales, 2001) states that the hexagonal tiling is the best way to divide a surface into

regions of equal area with the least total perimeter. Moreover, the Voronoi tessellation of 2D lattices of points gives an irregular honeycomb tessellation. For these reasons, the Voronoi decomposition can be applied to approximately minimize the total length of the seams. The seed points for the Voronoi tessellation are assumed to be the centroid of the images projected in the mosaic plane; for each seed the method determines a corresponding region consisting of all points closer to that seed than to any other.




Figure 4. (a) Image frames projected onto the mosaic reference.
 (b) Voronoi tessellation.

Figure 4(a) shows the area covered by each image after projecting it on the mosaic surface, whereas Figure 4(b) represents the areas assigned to the images using Voronoi decomposition.

Please note that we assume that the Voronoi region is contained in the corresponding image; this happens in most cases, but it cannot be formally guaranteed unless a constrained tessellation is used.

Finally, the best cutting paths between overlapping images are computed, in order to avoid the creation of seamlines that pass through regions where there are significant differences between adjacent images. For each seamline, the following procedure is used:

1. *Costmap computation*: a costmap that provides the information for the search of the best cutting path is constructed as the squared difference (pixel by pixel) between the images in the overlap area.

2. *Costmap to graph conversion*: the 2D costmap is converted into a graph assuming that each pixel becomes a node and that each adjacent nodes pair is connected by two oriented edges. The edge takes the weight from the value of the end pixel.
3. *Best path identification*: Dijkstra's algorithm (Dijkstra, 1959) is used over the generated graph to find the path of minimum cost.

When multiple images overlap on the same area, a costmap is computed for each image pairs and then averaged.

The procedure described requires as input the starting and end points of the seamlines. Since the Voronoi vertices computed before are independent from the image contents (they depend only on the position of the image centroids), they can fall in an area where the differences between adjacent images are high, thus negatively influencing the search of the best cutting path. To maintain the Voronoi polygons and at the same time determine optimal cutting paths, starting and end points are chosen as the ones that have the lowest cost in a neighbourhood of each Voronoi vertex, as shown in Fig. 5.




Figure 5. Costmap for the vertex optimisation. The chosen vertices are represented by a red cross.

Figure 6 reports an example of the seamlines detected through the proposed procedure. First the Voronoi tessellation is performed (Fig. 6(a)), then straight cuts between adjacent images are substituted by the optimal seamlines obtained through the Dijkstra's algorithm (Fig. 6(b)). One can notice that the cutting paths follow the space between the trees and rarely cross them. This is due to the fact that the ground between the trees is less textured and the cost corresponding to this zone is low. Thus, during the best cutting path computation, this area is preferred over the ones covered by trees.

The whole procedure described in this section determines a segmentation of the mosaic plane into disjoint regions. The final mosaic is obtained by filling each region with pixels sampled from a single source image.

5. EXPERIMENTAL VALIDATION

The proposed method was validated on two challenging dataset of aerial images. The results were compared to the ones obtained with *AutoStitch*¹ (Brown and Lowe, 2007), a popular software for image stitching.

Dataset 1 is composed by 29 images of size 1600×1200 pixels, acquired by a helicopter during a four strips flight. The images are characterized by small overlap areas and varying camera settings were used for each strip acquisition, thus the images present

¹available at <http://matthewalunbrown.com/autostitch/autostitch.html>




Figure 6. Seams produces by Voronoi tessellation (top) and after optimisation.

significant colour differences that have to be corrected during the mosaic generation. In particular, some images have a preponderant unrealistic red colour. The scene is almost planar (parallax effects are limited) but there are moving vehicles.

Figure 7 shows the mosaic created by aligning the images via homography synchronization (see Sec 4.1) and stitching them together according to a simple painter's algorithm (newer images overwrites older ones). Small misalignments can be noticed and strong colour distortions are evident between adjacent strips.

The result obtained with the complete method proposed in Sec. 4 is represented in Fig. 8(a). Please note how the colours have been corrected with the affinity synchronization step (Sec. 4.2) and differences between adjacent strips have been removed (reddish images visible in Fig. 7 are now indistinguishable). Moreover, misalignments have been concealed and the seamline optimisation carried out with the Dijkstra's algorithm (Sec. 4.3) preserved the integrity of the moving objects and avoided ghosting effects. A limited number of seams are still visible on the street and they could be easily eliminated by applying a blending algorithm at the end of the proposed procedure. The mosaic generated with *AutoStitch* is shown in Fig. 8(b). The software uses gain compensation and blending to conceal seams between overlapping images. The obtained mosaic is good, no misalignment is visible and the colour correction algorithm worked quite well. However, images with a strong red component are still visible and the blending in correspondence of moving objects caused ghosting effects.

Dataset 2 is composed by 27 images of size 4000×3000 pixels, acquired with a UAV with a forward and a side overlap of 80%.




Figure 7. Dataset 1: Mosaic obtained after image alignment, with no colour correction and seamline optimisation.

The scene is not planar (there is a large building in the middle and high trees in the bottom left part) and this determines strong misalignments after the homography computation. There are no evident colour differences except for weak illumination changes. No moving objects can be noticed in the images. Figure 9 shows the mosaic generated after homography synchronization, stitching images together according to the painter's algorithm, whereas Fig. 10(a) represents the final mosaic, in which colour variations and visible cuts are perfectly concealed. Due to the high overlap, 13 images were removed because completely covered by the others. Minimizing the number of images used to cover the entire scene results in seams of shorter length and reduces the misalignment and parallax visible effects. For this dataset, the mosaic generated with *AutoStitch* (Fig. 10(b)) is not satisfactory. The blending algorithm in presence of strong parallax caused blurred areas and other artifacts.

6. CONCLUSION

In this paper we proposed a novel procedure for image mosaicking, based on the approach known as *synchronization*. The image alignment and colour correction steps that usually characterize mosaic generation were addressed as a homography synchronization and an affine synchronization problem, respectively. Synchronization can be seen as upgrading from relative information, which involves two overlapping images at a time, onto absolute information, which involves all the images simultane-




Figure 8. Dataset 1. (a) Mosaic obtained with the proposed method. (b) Mosaic obtained with *AutoStitch*.



Figure 9. Dataset 2: Mosaic obtained after image alignment, with no colour correction and seamline optimisation.

ously. Thanks to this approach, misalignments and colour differences are globally minimized, thus leading to the generation of homogeneous, visually appealing mosaics. In order to conceal the seams that can still be visible after image registration and colour correction, a final step is applied, that is composed by Voronoi tessellation and seams optimisation via the Dijkstra's algorithm (Dijkstra, 1959). Results shown in Sec. 5 proved that

this novel procedure generates pleasant looking mosaic, without resorting to blending algorithms. As a matter of fact, each pixel of the resulting mosaic comes from a *single* image.

ACKNOWLEDGEMENTS

The authors wish to thank Helica s.r.l. and Airmap for providing dataset 1 and 2, respectively.

REFERENCES

- Abdel-Aziz, Y. and Karara, H., 1971. Direct linear transformation into object space coordinates in close-range photogrammetry, in *proc. symp. close-range photogrammetry*. In: *Proc. Symp. close-range Photogrammetry, Univ. Illinois at Urbana-Champaign*, pp. 1–18.
- Arie-Nachimson, M., Kovalsky, S. Z., Kemelmacher-Shlizerman, I., Singer, A. and Basri, R., 2012. Global motion estimation from point matches. *Proceedings of the Joint 3DIM/3DPVT Conference: 3D Imaging, Modeling, Processing, Visualization and Transmission*.
- Arrigoni, F., Fusiello, A. and Rossi, B., 2016a. Camera motion from group synchronization. In: *Proceedings of the International Conference on 3D Vision (3DV)*, pp. 546–555.
- Arrigoni, F., Rossi, B. and Fusiello, A., 2016b. Spectral synchronization of multiple views in SE(3). *SIAM Journal on Imaging Sciences* 9(4), pp. 1963 – 1990.




Figure 10. Dataset 2. (a) Mosaic obtained with the proposed method. (b) Mosaic obtained with *AutoStitch*.

Brown, M. and Lowe, D. G., 2007. Automatic panoramic image stitching using invariant features. *International journal of computer vision* 74(1), pp. 59–73.

Cox, I. J., Hingorani, S., Maggs, B. M. and Rao, S. B., 1996. A maximum likelihood stereo algorithm. *Computer Vision and Image Understanding* 63(3), pp. 542–567.

Davis, J., 1998. Mosaics of scenes with moving objects. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 354–360.

Dijkstra, E. W., 1959. A note on two problems in connexion with graphs. *Numerische mathematik* 1(1), pp. 269–271.

Du, Q., Raksuntorn, N., Orduyilmaz, A. and Bruce, L. M., 2008. Automatic registration and mosaicking for airborne multispectral image sequences. *Photogrammetric Engineering & Remote Sensing* 74(2), pp. 169–181.

Fischler, M. A. and Bolles, R. C., 1981. Random Sample Consensus: a paradigm model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24(6), pp. 381–395.

Hales, T. C., 2001. The honeycomb conjecture. *Discrete & Computational Geometry* 25(1), pp. 1–22.

Hansen, M., Anandan, P., Data, K., Wal, G. and Burt, P., 1994. Real-time scene stabilization and mosaic construction. In: *Proceedings of IEEE Workshop on Applications of Computer Vision*.

Li, L., Yao, J., Lu, X., Tu, J. and Shan, J., 2016. Optimal seamline detection for multiple image mosaicking via graph cuts. *ISPRS Journal of Photogrammetry and Remote Sensing* 113, pp. 1–16.

Lowe, D. G., 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60(2), pp. 91–110.

Marzotto, R., Fusiello, A. and Murino, V., 2004. High resolution video mosaicing with global alignment. In: *Proceedings of the*

IEEE Conference on Computer Vision and Pattern Recognition, Vol. I, IEEE Computer Society, pp. 692–698.

Maset, E., Arrigoni, F. and Fusiello, A., 2017. Practical and efficient multi-view matching. In: *Proceedings of IEEE International Conference on Computer Vision*, Vol. 2 number 4, p. 5.

Oliveira, M., Sappa, A. D. and Santos, V., 2011. Unsupervised local color correction for coarsely registered images. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, IEEE, pp. 201–208.

Pan, J., Wang, M., Li, D. and Li, J., 2009. Automatic generation of seamline network using area voronoi diagrams with overlap. *IEEE Transactions on Geoscience and Remote Sensing* 47(6), pp. 1737–1744.

Schroeder, P., Bartoli, A., Georgel, P. and Navab, N., 2011. Closed-form solutions to multiple-view homography estimation. In: *Applications of Computer Vision (WACV), 2011 IEEE Workshop on*, pp. 650–657.

Singer, A., 2011. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis* 30(1), pp. 20 – 36.

Szeliski, R., 2006. Image alignment and stitching: a tutorial. *Found. Trends. Comput. Graph. Vis.* 2(1), pp. 1–104.

Tai, Y.-W., Jia, J. and Tang, C.-K., 2005. Local color transfer via probabilistic segmentation by expectation-maximization. In: *IEEE Conference on Computer Vision and Pattern Recognition*, Vol. 1, IEEE, pp. 747–754.

Uyttendaele, M., Eden, A. and Skeliski, R., 2001. Eliminating ghosting and exposure artifacts in image mosaics. In: *IEEE Conference on Computer Vision and Pattern Recognition*, Vol. 2, pp. 509–516.

Xu, W. and Mulligan, J., 2010. Performance evaluation of color correction approaches for automatic multi-view image and video stitching. In: *IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, pp. 263–270.