Challenges of Close-Range Underwater Optical Mapping

Ricard Prados, Rafael Garcia
Computer Vision and Robotics Group
University of Girona, Girona, 17001 Spain
Email: {rprados, rafa}@eia.udg.edu

Javier Escartín
Équipe de Géosciences Marines
CNRS/IPGP UMR7154
75238 Paris, France
Email: escartin@ipgp.fr

László Neumann
Computer Vision and Robotics Group
University of Girona, Girona, 17001 Spain
ICREA, Barcelona, 08010 Spain
Email: lneumann@eia.udg.edu

Abstract—Underwater optical mapping often involves the use of image mosaicing techniques. High quality mosaicing requires the application of blending methods to achieve continuous and artifact-free mosaics. Image blending has a dilated history of over three decades in the terrestrial and aerial fields. Unfortunately, the nature of the underwater medium adds additional difficulties to the mosaicing and blending tasks. In this paper a survey of the blending methods is given, focusing the attention on its applicability to underwater mosaicing. Image acquisition is performed by Autonomous Underwater Vehicles (AUVs) or Remotely Operated Vehicles (ROVs) in the deep ocean, a medium with aggressive light absorption and disrupting scattering effects that requires of the use of artificial lighting. A comprehensive comparison of the basic features and limitations of some of the most important existing blending techniques is presented. The goal is the generation of seamless and visually pleasant large area photo-mosaics of the seafloor, free from double contouring, ghosting and other disturbing and common blending artifacts.

I. Introduction

High resolution optical imaging provides scientists (e.g., archeologists, geologists, biologists, among others), with accurate and rich visual information of the ocean seafloor. Object measurement, monitoring and terrain prospection are among the tasks which can be performed based on underwater photographs.

Despite increased resolution of nowadays electronic still cameras, underwater regions of interest are usually too extensive to be covered with the required level of detail on a single shot. Furthermore, due to light attenuation, seafloor images should be acquired at short range in order to minimize its impact. Consequently, several images should be merged in order to obtain a picture that covers the whole surveyed area. Moreover, this approach needs to deal with additional difficulties such as non-uniform illumination, light scattering and moving objects (see Fig. 1). All these effects cause strong visual artifacts when several images are stitched together to form a photo-mosaic.

Image blending techniques are widely used in indoor and outdoor contexts in order to generate seamless panoramas. Most of the conventional algorithms work on the basis of a still camera with rotation and negligible translation between images. In the underwater medium the phenomena described above constrain the acquisition range to a few meters over the seafloor, commonly using artificial light sources, which leads to illumination artifacts but also to geometrical problems such as parallax even when the 3D relief of the seabed is small. Despite the existence of 3D reconstruction methods intended to underwater imagery [1], [2], [3], which allow avoiding parallax issues, those ones cannot be applied when the overlap between images is small. That is the case of the datasets used in this work, acquired by autonomous or teleoperated vehicles equipped with still cameras that are synchronized with stroboscopic flashes as lighting sources. The stroboscopic light requires about 10 seconds to recharge and shot, due to its high level of power, restricting this time interval the maximum acquisition rate.

Quality blending is important inasmuch as a scientist can significantly improve its visual analysis and data interpretation when working with a continuous and uniform high resolution photo-mosaic instead of a simple sequence of stitched images, i.e., a non-blended mosaic, or the individual images.

Besides the image quality of the mosaic, the size of the data to process, concerning picture size and number of pictures, is another challenge by itself. The large photo-mosaic dimensions affects not only its processing time but also the computer memory management during its processing, requiring a tailored strategy for the agile processing and visualization of the results.

The aim of this work is to review the current image blending techniques aimed at mosaicing. We point out their capabilities and weaknesses when applied in the underwater medium. The target of the study is to develop an adequate processing pipeline to deal with huge and challenging underwater image sequences in order to build high quality large scale photo-mosaics of the seafloor.

II. State-of-the-Art Blending Approaches

A. Literature Review

The basic principles of image blending where established four decades ago [4] and include two main concepts which drive to two algorithm groups [5]: *transition smoothing* and *optimal seam finding*. On the one hand, transition smoothing methods [6], [7] fade the images along a common overlapping region in order to minimize the visibility of the seam. On the other hand, optimal seam finding algorithms [8], [9] focus on computing the joining boundary which reduces in
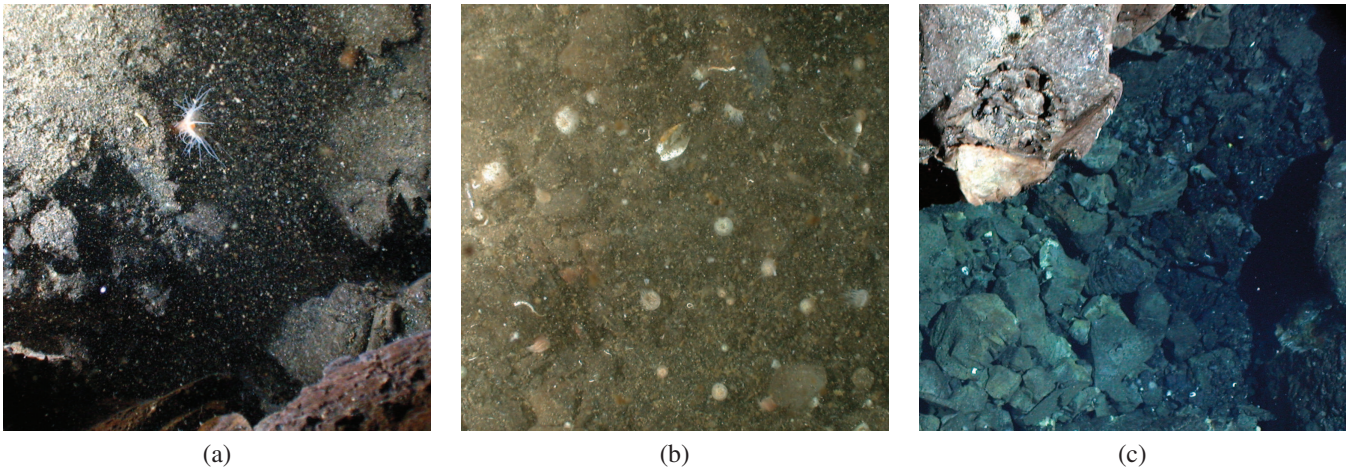
Fig. 1. (a) Example of *backward scattering* due to the reflection of the rays from the light source on the suspended particles. The numerous particles in suspension reflects the light making difficult the identification of the rock structures. (b) Example of *forward scattering* caused by the local inter-reflection of the light on the particles suspended on the water. The inter-reflections hide the terrain behind the particles. (c) Effects produced by the *light absorption* of the water resulting in an evident attenuation of certain color wavelengths which increases with the distance. (Images courtesy of Dan Fornari (Woods-Hole Oceanographic Institution))

a higher degree the photometric differences on its common area to achieve a non-noticeable transition. Actually, several current methods combine the benefits of both approaches, smoothing the transition along an optimally found seam [4], [10]. Unfortunately, most of the state-of-the-art techniques are intended to deal with terrestrial or aerial images, where the camera rotates over its focal point in the case of conventional panorama generation, or it translates when being very far from the scene, thus avoiding parallax problems.

Among the several approaches listed in Fig. 2 a subset of them, based on their main features, has been selected to perform the evaluation and comparison experiments.

In the early 70's, D. Milgram [4] established the basic concepts to achieve a seamlessly transition between two images, which were applied by most of the methods that arise in the next decades. The author proposed a "zero order" adjustment to compensate illumination differences, a non-optimal seam definition strategy (improved in [10]) to compute the image boundaries on the overlapping region (*i.e. optimal seam finding*) and a pixel-row based weighted average to smooth the transition along the boundary (*transition smoothing*). The method was intended to deal with a single overlapping region between two registered images and the transition smoothing was performed only at pixel-row level.

The limitation of two overlapping images present in [4] was firstly addressed instance by Peleg [11], introducing the concept of *Seam-Eliminating Function* (*SEF*). The *SEF* is based on a luminance weighting map, computed using an iterative relaxation algorithm, which is intended to smooth the transition from any number of overlapping images, setting at zero the intensity differences along the seams.

In the context of the transition smoothing methods, Burt and Adelson [12] introduced in 1983 the concept of *Image spline*. This one was used by the authors to refer the transformation needed to obtain a smooth transition among several images.

The main idea of the approach lies in the application of a different transition smoothing function (*e.g.* a weighted average) along an arbitrary boundary on a given transition region width T depending on the scale of the image features (*i.e.* the size of the salient features and structures). The images to be merged are decomposed in a set of band-pass component images, in order to apply a different spline with an appropriate transition width on each band. This band-pass approach ensures that features with the same scale are fused at the same filtering level along an appropriate transition region width T. The computed band-pass components are then recombined into the final mosaic image using a simple addition. The method suppresses the visibility of the seams and reduces the noticeability of the misalignments when registration is not perfect, having as a counterpart double contouring and ghosting effects when the image misalignment is significant (see Fig. 3). In 1996, Hsu and Wu [13] extended the idea of Burt and Adelson [12] by applying the method in wavelet subspaces, although obtaining similar results but with a higher computational cost.

Ghosting and double contouring artifacts in the overlapping area due to moving objects can be reduced by analyzing the pixels belonging to the transition region. One of the first approaches facing moving objects was proposed by Davis [8] in 1998, who found an optimal seam using Dijkstra's algorithm [14] through the photometric differences computed between two registered images. The path traverses the overlapping region through the pixels having the smallest difference in value between images and assumedly belonging to the same scene structure, avoiding the moving objects, which rests inside or outside the final fused mosaic. Unfortunately, if the object is not fully contained in one of the overlapping images it might be bisected. Still based on the same idea of using image differences to localize moving objects in image panoramas, Uyttendaele *et al.* [7] proposed in 2001
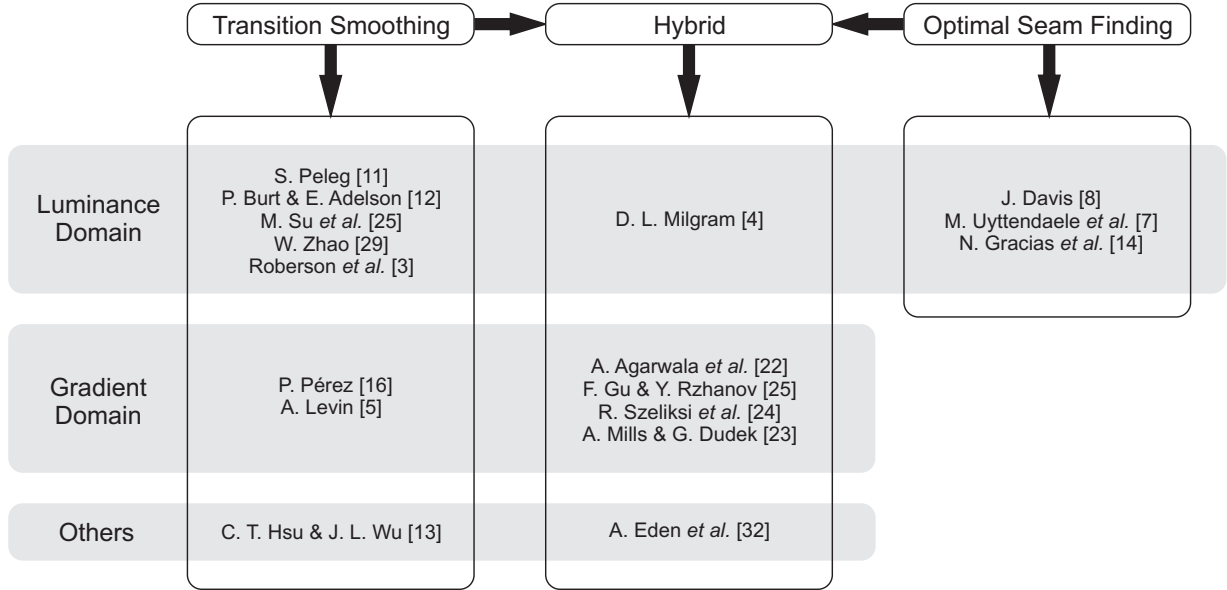
State-of-the-Art Blending Approaches

Transition Smoothing → Hybrid ← Optimal Seam Finding

Luminance Domain

Transition Smoothing:
S. Peleg [11]
P. Burt & E. Adelson [12]
M. Su *et al.* [25]
W. Zhao [29]
Roberson *et al.* [3]

Hybrid:
D. L. Milgram [4]

Optimal Seam Finding:
J. Davis [8]
M. Uyttendaele *et al.* [7]
N. Gracias *et al.* [14]

Gradient Domain

Transition Smoothing:
P. Pérez [16]
A. Levin [5]

Hybrid:
A. Agarwala *et al.* [22]
F. Gu & Y. Rzhanov [25]
R. Szeliksi *et al.* [24]
A. Mills & G. Dudek [23]

Others

Transition Smoothing:
C. T. Hsu & J. L. Wu [13]

Hybrid:
A. Eden *et al.* [32]

Fig. 2. State-of-the-Art Blending Approaches Scheme summarizing some of the most relevant methods in the literature. Concerning the technique principles, *Transition Smoothing* and *Optimal Seam Finding* are the two main groups. Additionally, *Hybrid* methods combine the benefits of both to improve the quality of the results. There are two principal (but not unique) domains in which the methods work, which are *Luminance* and *Gradient*, despite *Wavelet* and *Radiance* have been also used in the literature.

a method to suppress the ghosting effect in mosaic images due to moving objects, along with a procedure to adjust the exposure across multiple images, in order to eliminate visible shifts in brightness and hue. This approach was intended to deal with multiple overlaps, and the authors proposed the search for regions of difference (RODs) on the overlapping areas, only using information from one image per ROD. Hence RODs are defined in different images to be corresponding, *i.e.* to belong to the same scene object, if they have any overlap at all. RODs are then used to build a graph, where the minimum weight vertex cover must be computed. The authors stated that some conflictive situations may appear where a wrong elimination of RODs causes holes in the mosaic, but also that these artifacts are rare in practice. In reference to the exposure artifacts, a block-based adjustment technique is applied on the overlapping region, which iteratively computes a different quadratic transfer function for each block along all the images stack. The blocks are also blended with their neighbors using bilinear or biquadratic interpolation. Gracias *et al.* [15] proposed a method based in graph-cuts on a watershed segmentation, intender to achieve computational and memory cost reduction of Disjtra's based optimal seam finding.

Concerning the transition smoothing field and in the context of image editing with seamless object insertion and cloning, Pérez *et al.* [16] proposed in 2003 a generic interpolation machinery based on Poisson equations. This approach smoothes the transition along the seam between a given object and the background where it is placed. The method represents the first important approach to image mosaicing in the gradient domain and allows to convincingly fuse image regions (or objects) with transparent appearance. The framework is based on the partial differential equation with Dirichlet boundary conditions which specifies the Laplacian of an unknown function over the domain of interest, along with the unknown function values over the boundary of the domain. As an extension of the technique presented in [17], Pérez *et al.* proposed to modify the problem of image interpolation through Poisson equation by introducing further constraints in the form of a guidance field. In the same context, Levin *et al.* [5] proposed a method based in several cost functions for the evaluation of the quality of the stitching defined on the gradient domain. The authors named GIST (Gradient-domain Image STitching) the developed framework based on this method, providing two main approaches to image stitching. First, images are combined in the gradient domain, reducing global inconsistences between the stitched parts due to illumination changes and variations in the camera photometric response. Second, the mosaic image is inferred by optimization over image gradient, reducing seam artifacts and edge duplications. In 2006 Jia *et. al* [18] extended Pérez approach [16] with an objective function to compute optimized boundary conditions. Pérez method was applied in 2007 by Hays & Efros [19] to perform semantically valid image completion.
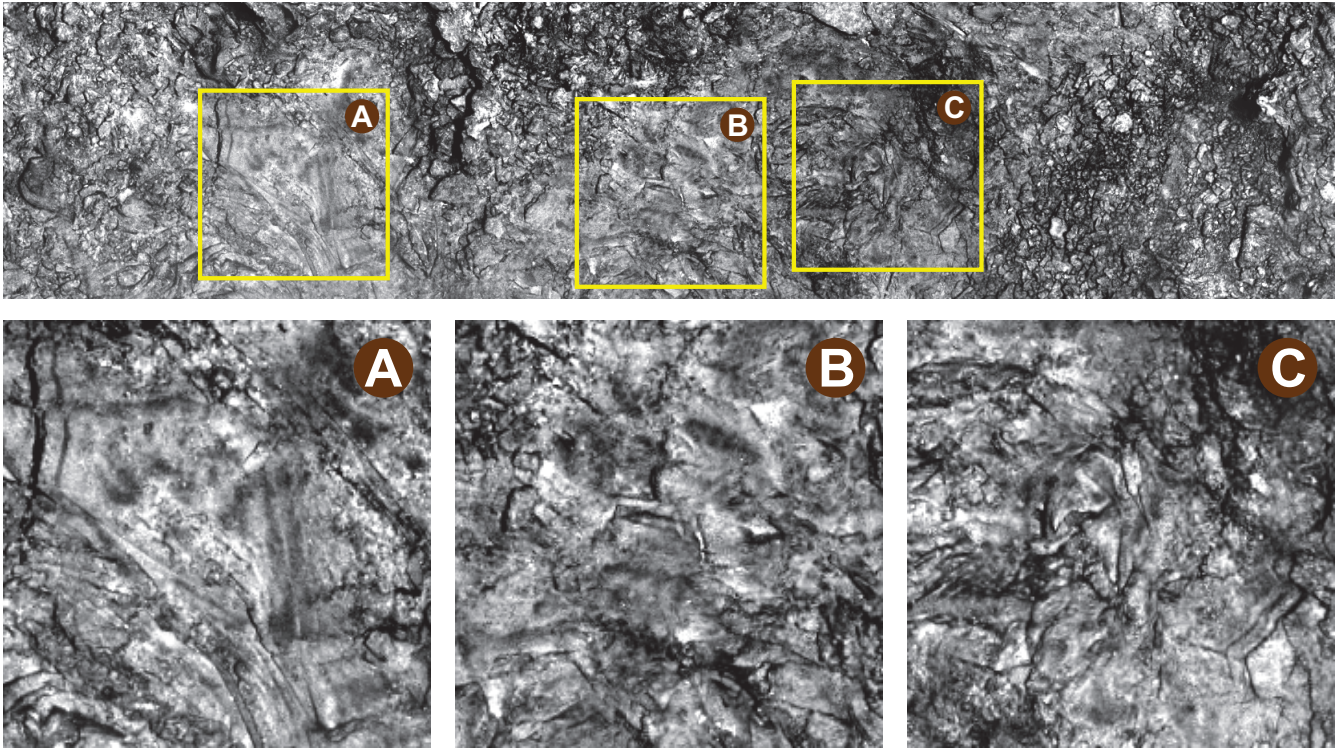
Fig. 3. Small region of a blended mosaic mosaic suffering from double contouring due to missregistration and 3D relief. (A, B, C) Zoomed detail of the artifacts, showing double contours on the most prominent seabed rock structures. (Dataset images are from the CNRS/IFREMER MoMAR'08-Leg1 cruise)

Agarwala *et al.* proposed in 2004 a technique based in the advantages of the two main classes of blending algorithms [20], *i.e. transition smoothing* and *optimal seam finding*. Firstly, a *graph-cut optimization* [21] is performed on the overlapping region pixels, in order to find the pixel-path with the lowest cost in terms of photometric differences between images. Secondly, a *gradient-domain fusion* [16] is applied around the computed path (or seam) in order to reduce or remove any remaining visible artifacts after the image seam joining. In 2007, Agarwala [22] presented a hierarchical approach to improve the efficiency of gradient-domain compositing. The efficiency improvement is achieved by observing that the difference between a given image composite and its associated gradient-domain composite is largely smooth, and the pattern of this smoothness can be predicted *a priori*. This difference is solved by adaptively subdividing the domain using a *quadtree* hierarchical structure [23]. Szeliski *et al.* [24] presented in 2008 a technique for fast Poisson blending and gradient domain compositing which associates to each input image a separate low-resolution offset map, that can be represented using a low-dimensional spline. The resulting linear system is much smaller than either the original Poisson system or the quadtree spline approximation of a single offset map. Since each of the offset fields is represented using a low-dimensional spline, the result representation is called *multi-spline*.

Su *et al.* [25], in the context of image blending in the gradient domain, presented a technique based on the minimization of an energy function, considering a combination of both gradient and luminance on the overlapping region. Within this blending energy function, indented to combine low-level image properties, two variation terms are measured and minimized: *image value variation* and *first derivative variation*. The first term, image value variation, measures the difference between corresponding pixel values of the images to be combined and the mosaic itself. The second term, the first derivative variations, measures the difference between the blended values of each respective first derivative and the first derivative of the mosaic. The resultant image is effectively obtained by minimizing the blending energy function.

With the aim of increasing the time and memory efficiency of gradient blending methods, Szeliski *et al.* [24] presented in 2008 a technique for fast Poisson blending and gradient domain compositing which associates to each input image a separate low-resolution offset map, that is represented by low-dimensional spline. The resulting linear system is considerably smaller than either the original Poisson system or the quadtree spline approximation of a single offset map. Since each of the offset fields is represented using a low-dimensional spline, the result representation is called *multi-spline*.

The approaches in the literature that have specifically faced the problem of mosaicing in the underwater context are not numerous. Gu and Rzhanov [26], similarly to [20], proposed a graph-cut, in order to select the optimal seam between two images, and the application around the found boundary of a gradient domain fusion. The method claims to overcome the defects of single graph-cut techniques, showing noticeable

seams in the case of changing illumination conditions, and the ones of gradient domain fusion, producing blurring in case of misalignment. However, the authors do not define the criteria to select the contributing image in the case of multiple image overlaps on a given region. Thus, in [26] the approach assumed to be limited to "panoramic" mosaics where only two images overlap over the same area. On the other hand, a graph-cut on the intensity differences along the overlapping region may lead to non-optimal seams, particularly with images displaying different exposures or illuminations.

Recently, Roberson *et al.* [3] proposed an efficient three-dimensional approach for underwater robotic surveys, mainly targeted to achieve computation and visualization speed on the reconstructions. Despite the usage of illumination compensation functions to solve some of the effects of wavelength attenuation, the selected method for blending is based on the well known Burt & Adelson [12] multi-spline approach. In case of missregistration this method suffers from the issues mentioned above, *i.e.* ghosting and double contouring.

B. Techniques Summary

The classification scheme shown in Fig. 2 summarizes the most important properties of the main mosaicing and blending methods in the literature which may have an important impact in the underwater imagery context. The strengths and weaknesses of the methods are next highlighted, allowing us to perform a qualitative evaluation of all of them.

Considering the main *principle* of the above presented techniques, the combination of a transition smoothing around an optimally found boundary, *i.e.* the use of *hybrid* methods, seems to be the most adopted approach in the recent bibliography, independently on its application context. The tolerance of the techniques to *moving objects* is strongly tight to the *principle*, being all the optimal seam finding based methods able to deal with moving objects in a certain degree. In the general case, this tolerance is not actively treated, but becomes a side effect of the optimal seam search algorithm, which tends to place the image cut on areas where photometric differences are small. As a consequence, overlapping areas with moving objects are cut. Both luminance and gradient *domains* are widely used, with the gradient domain gaining special importance on the latest publications like [24], [27], [28]. The nature of the domain allows to easily reduce exposure differences between neighboring images without more complex preprocessing, due to the fact that it is not sensitive to differences in exposures. Nevertheless, methods actively applying an *exposure correction* algorithm on the input images obtain more visually pleasant results. The ability to remove *ghosting effects* and *double contouring* are commonly tight, and the methods solving one problem are also able to deal with the other. Concerning the *color* channels treatment, it is similar in most of techniques in the literature. Concretely, the blending tasks are always performed in a single channel, independently of the number of channels of the source images (*i.e.*, color or gray scale images). There is only a single method conceived to work on *real time* [29], which therefore

requires *sequential processing*. However, methods which do not perform a *sequential processing* are better conditioned to deal with problems like the exposure compensation so as to ensure global appearance consistency. The *sequential processing* tends to accumulate drift on the image corrections, which may strongly depends on the first processed and stitched image. Few blending methods are claiming to work with *high dynamic range* images. Nevertheless, gradient based blending methods are able to intrinsically deal with this kind of images, due again to the nature of the domain. Methods able to process *high dynamic range* images require the application of *tone mapping* [30] algorithms on the generated mosaic image in order to generate displayable results. The high dynamic range should be reduced in order to be represented into low dynamic range devices, such as a screen monitors or printers. The *multiresolution* approach, based on the idea of Burt and Adelson [12] is applied by Su *et al.* [25] on the wavelet domain, but is the only variation of this idea in the literature. Finally, the *parallax robustness* is strongly related to the tolerance for moving objects. In fact, the parallax robustness can be considered in practice as the ability of a method to avoid cutting or repeating objects or shapes, like in the case of moving objects.

III. COMPARATIVE RESULTS

The main objective of this paper is to point out the benefits and drawbacks of the most representative state-of-the-art blending methods when applied to underwater imaging. Relevant techniques proposed in the [12], [31], [32] and representing the main blending principles are applied and compared to challenging underwater sequences in order to evaluate their performance and determine their main weaknesses. Non-uniform illumination, light attenuation and scattering are characteristics not present in conventional indoor or outdoor images and the state-of-the-art techniques are not intended to adequately deal with them. The comparison of the different approaches will help to select the most adequate processing pipeline to generate large scale underwater photo-mosaics.

Objective quantification of image registration quality has been faced by several authors in [33], [34], [35]. Unfortunately, in absence of a groundtruth, and taking into account the fact of building 2D mosaics from a seafloor that might significantly violate the planar assumption, performing a quantitative comparison between methods is not feasible. Nevertheless, in order to detect the main issues affecting underwater image mosaicing, a qualitative evaluation of the results is sufficient to draw significant conclusions.

The datasets shown on the following comparison were collected by the Victor-6000 ROV during the MoMAR'08 cruise (CNRS/IFREMER, France) at the deep-sea Lucky Strike hydrothermal field (Mid-Atlantic Ridge). The Victor-6000 ROV was equipped with a mapping payload, allowing it to undertake high-resolution, georeferenced seafloor survey work using the OTUS gray scale still camera [36]. The installed stroboscopic artificial lighting system allows the system to capture images every 10 seconds or longer, owing to the strobe recharge time.

Fig. 4 shows the result of merging two images acquired at different altitudes obtained by different state-of-the-art approaches. The images suffer from light attenuation at different levels, due to the difference in the distance from the camera to the seafloor. Additionally, the images also present darker corners caused by both light attenuation and non-uniform illumination of the artificial light device. The methods based on the transition smoothing principle evidence ghosting and double contouring, and despite the transition between the images is smooth in some cases (Fig. 4d,f,g,h) the difference on image sharpness makes both images clearly distinguishable.

An illumination compensation function [37] can be applied to the original images in order to correct this phenomenon, as can be seen in Fig. 5. The obtained images present a more homogeneous appearance along the whole surface, despite the details on the corner pixels, which cannot be recovered, since the detail richness of these regions is constrained by the camera sensor sensitivity.

Fig. 7 shows the same experiments performed in Fig. 4 but using the images corrected with the non-uniform illumination compensation function. In that case, the previously described artifacts are still present but the impression of continuity, particularly in the areas belonging to image corners, is higher.

The main blending principles in the literature, and independently of the problems caused by missregistration and parallax, seem to have special difficulties on dealing with images with significantly different appearances for a given common area. This situation is not common in panoramic or aerial imaging, where pictures are typically acquired in a way that avoids the appearance of this kind of issues. Unfortunately, when dealing with underwater images, it is frequent to register images with highly different appearances when a robot crosses a previously mapped area, but at a different altitude. Furthermore, the pairwise image registration is in some cases not possible, due to this difference in appearance, leading to a situation where the images should be stitched only based on the vehicle navigation data. In that extreme case, the geometrical image transformations might not be as accurate as desired, and images known to correspond to the same area might not even seem to belong to the same scene.

Concerning the generation of large photo-mosaics, and considering the strong differences that may present some sequentially acquired images, the usage of techniques based on sequential processing is not recommended. The drift in the application of some exposure or image enhancement corrections in a sequential manner may lead to totally degraded photo-mosaic.

IV. CONCLUSIONS AND FUTURE WORK

The common problem of mosaic generation and blending on indoor and outdoor scenes has been satisfyingly solved by the methods proposed on the literature. Unfortunately, the underwater medium adds new challenges to this task. Depth depending light attenuation, non-uniform illumination, forward and backward scattering and unavoidable parallax are some

Fig. 5. Image suffering from light attenuation on the image areas further from the artificial lighting source before and after the application of a non-uniform illumination correction function (stretched representation of a value range [0.5286..1]). (Dataset images are from the CNRS/IFREMER MoMAR'08-Leg1 cruise)
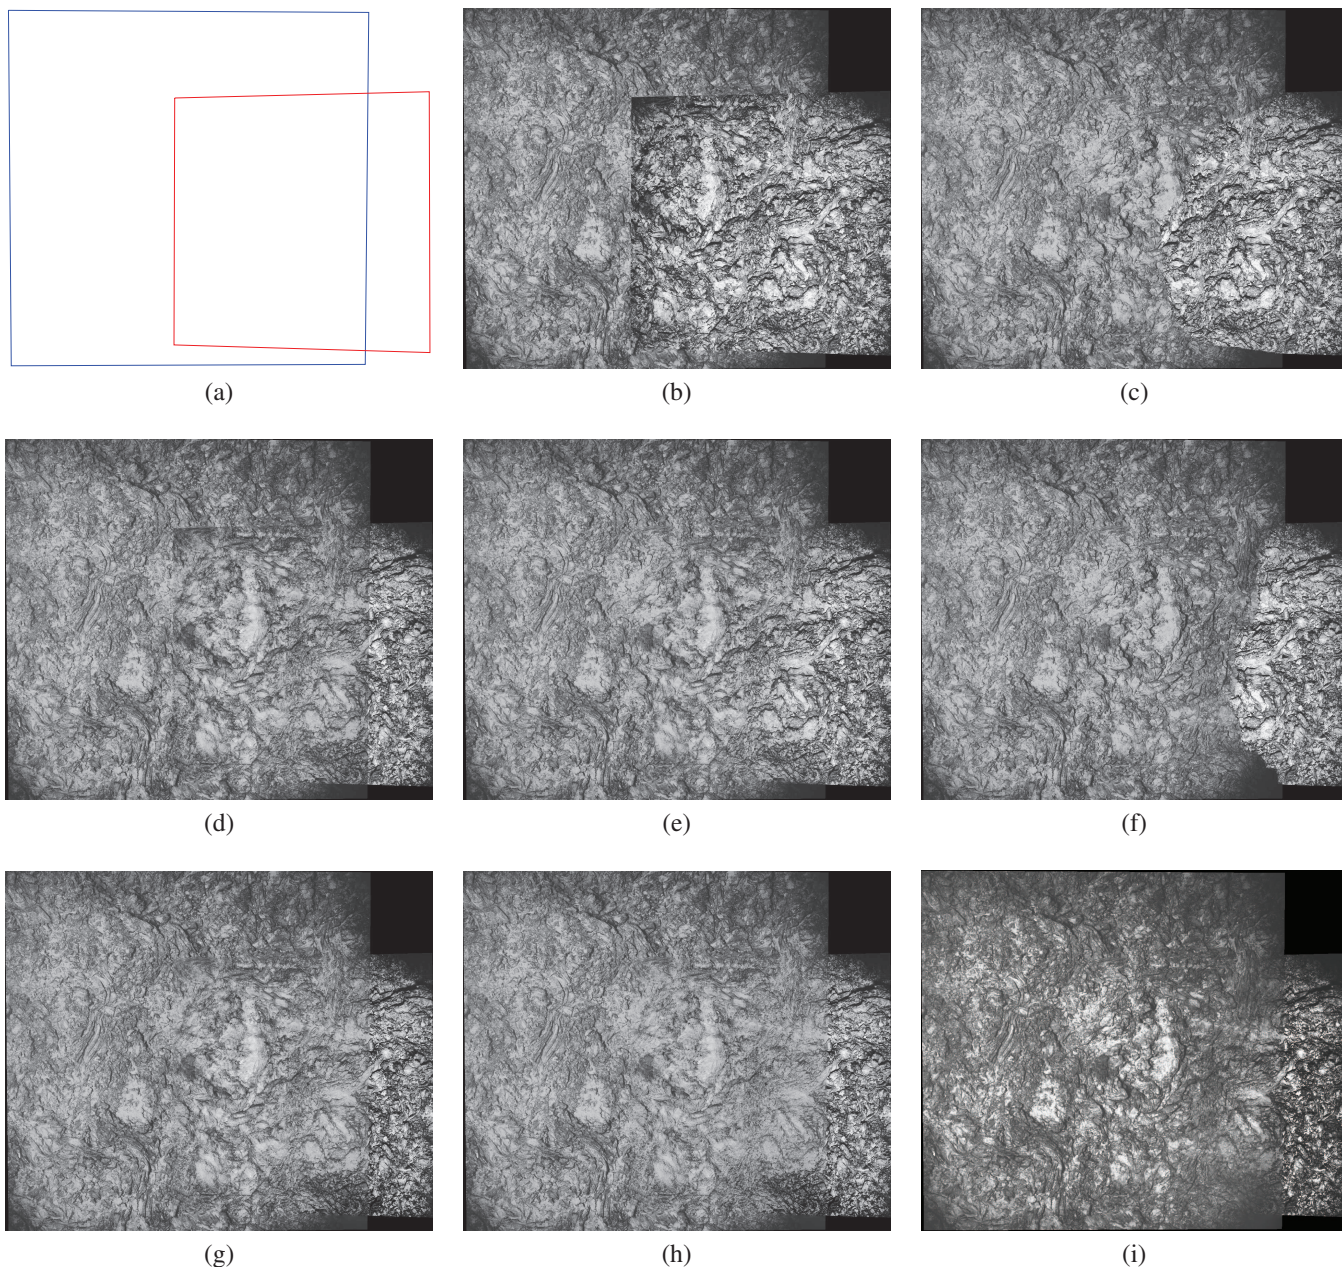
Fig. 4. Results of merging two images of the same scene acquired at different altitudes, suffering from light attenuation close to the borders and showing different visual appearances, using basic state-of-the-art techniques. (a) Registered images without blending. (b) Registered images selecting selecting the closest to the image centers pixels. (c) Linear average of the intensity values. (d) Linear average of the intensity values selecting only the pixels closer to the image centers. (e) Graph cut on the intensity domain. (f) Burt & Adelson multiband blending on the intensity domain. (g) Burt & Adelson multiband blending on the intensity domain giving more priority to the pixels closer to the image centers. (h) Gradient blending performed on a wide region around a graph cut seam computed on the intensity domain. None of the obtained results achieves a convincing image merging with continuous appearance. (Dataset images are from the CNRS/IFREMER MoMAR'08-Leg1 cruise)
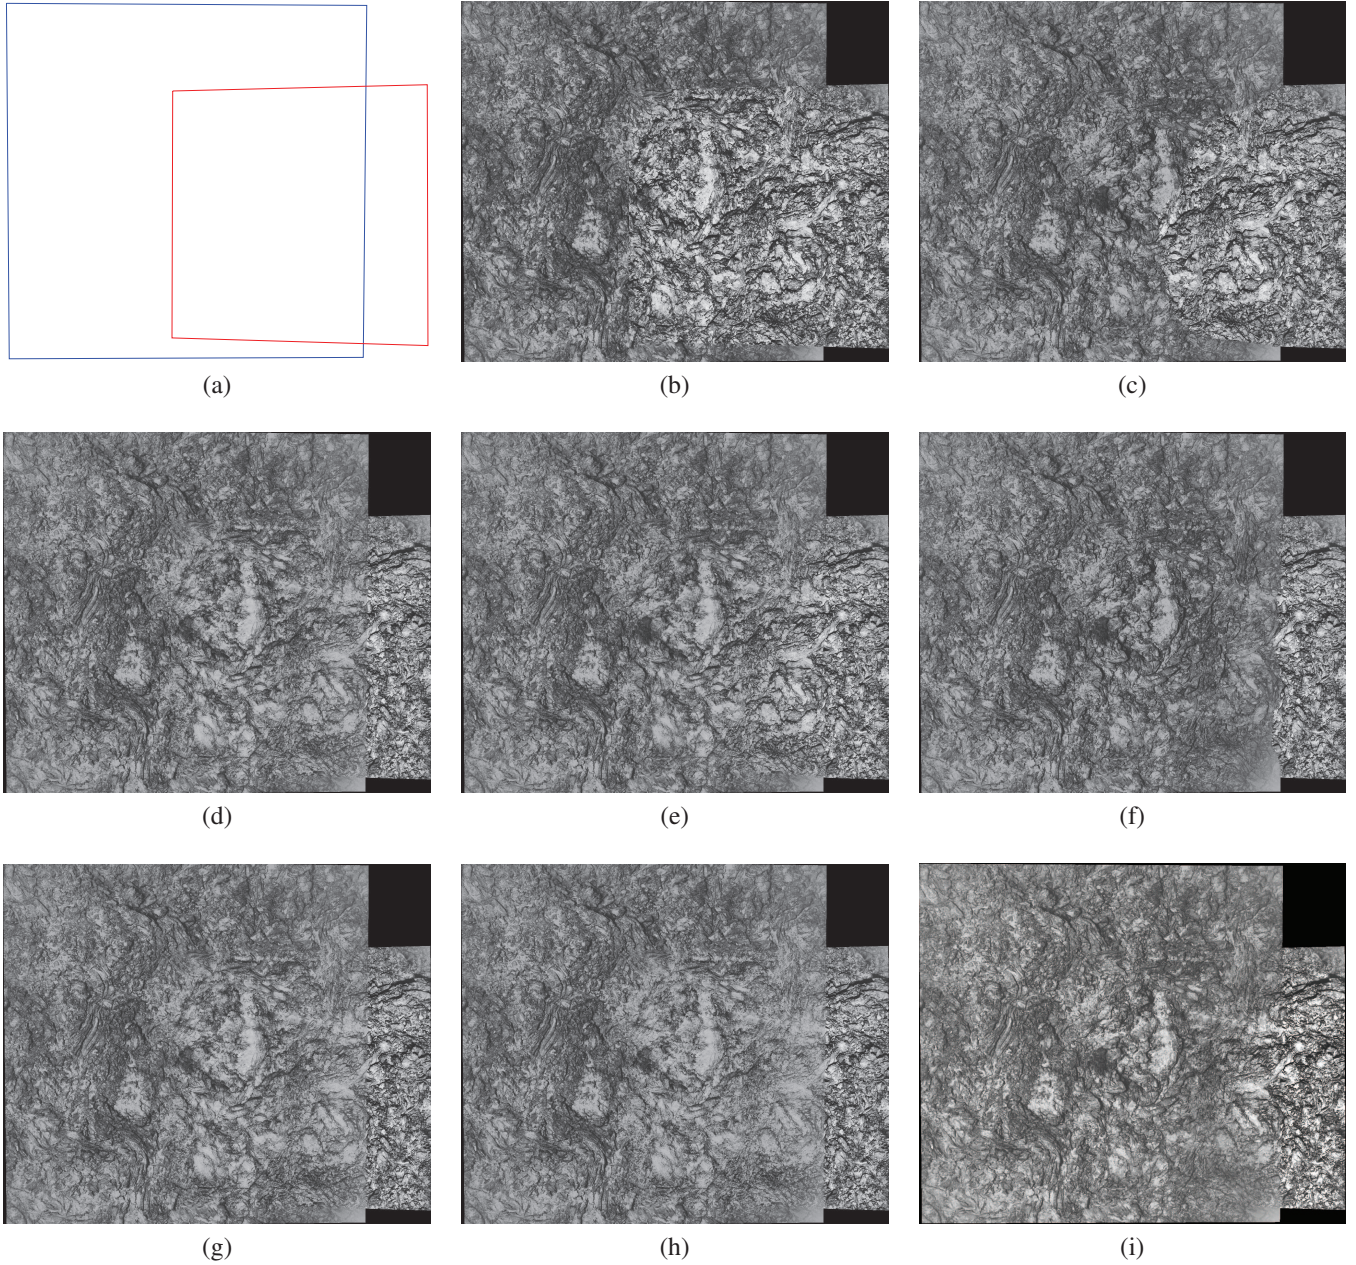
Fig. 6. Results of merging two images of the same scene acquired at different altitudes, after the application of an illumination correction function, using basic state-of-the-art techniques. (a) Registered images without blending. (b) Registered images selecting selecting the closest to the image centers pixels. (c) Linear average of the intensity values. (d) Linear average of the intensity values selecting only the pixels closer to the image centers. (e) Graph cut on the intensity domain. (f) Burt & Adelson multiband blending on the intensity domain. (g) Burt & Adelson multiband blending on the intensity domain giving more priority to the pixels closer to the image centers. (h) Gradient blending performed on a wide region around a graph cut seam computed on the intensity domain. None of the obtained results achieves a convincing image merging with continuous appearance. (Dataset images are from the CNRS/IFREMER MoMAR'08-Leg1 cruise)

of the problems for which state-of-the-art approaches do not provide a satisfactory solution.

State-of-the-art methods based on the *transition smoothing* principle suffer from ghosting and double contouring when missregistration, D relief or moving objects are present. When fusing images acquired at different altitudes, and consequently showing different levels of light attenuation and scattering, the blending using this kind of methods, even if smooth, does not provide a mosaic with convincing continuous appearance. This fact is due to the lack of adequate image enhancement and non-uniform illumination compensation mechanisms. Concerning *optimal seam finding* methods, in absence of the above mentioned image preprocessing strategies, lead to highly noticeable seams and joining. Furthermore, the estimated seam might not be optimal inasmuch as several methods work on the luminance differences within the overlapping regions, while these may have strong differences in illumination and detail richness. Finally, *hybrid* methods also fail to achieve really convincing continuous and homogeneous mosaics. Consequently, a new specific processing pipeline seems to be necessary to correctly deal with the challenges described above.

Our work in process is targeted to solve the specific underwater imaging issues prioritizing image quality / appearance. An adequate image preprocessing to reduce the illumination problems caused by both inhomogeneous artificial lighting and light attenuation, and an improved graph cut strategy to find the optimal image seams on underwater imagery are some of the solutions applied by the under development. From a technical point of view, an efficient implementation based on the decomposition of the problem in –limited by available memory– smaller sub-problems allows the processing of gigapixel photo-mosaics.

References

[1] T. Nicosevici and R. Garcia. On-line visual vocabularies for robot navigation and mapping. In *IROS'09: Proceedings of the 2009 IEEE/RSJ international conference on Intelligent robots and systems*, pages 205–212, Piscataway, NJ, USA, 2009. IEEE Press.

[2] T. Nicosevici, N. Gracias, S. Negahdaripour, and R. Garcia. Efficient three-dimensional scene modeling and mosaicing. *J. Field Robot.*, 26(10):759–788, 2009.

[3] M. Johnson-Roberson, O. Pizarro, S.B. Williams, and I. Mahon. Generation and visualization of large-scale three-dimensional reconstructions from underwater robotic surveys. *J. Field Robotics*, 27(1):21–51, 2010.

[4] D.L. Milgram. Computer methods for creating photomosaics. *IEEE Transactions on Computers*, 24(11):1113–1119, 1975.
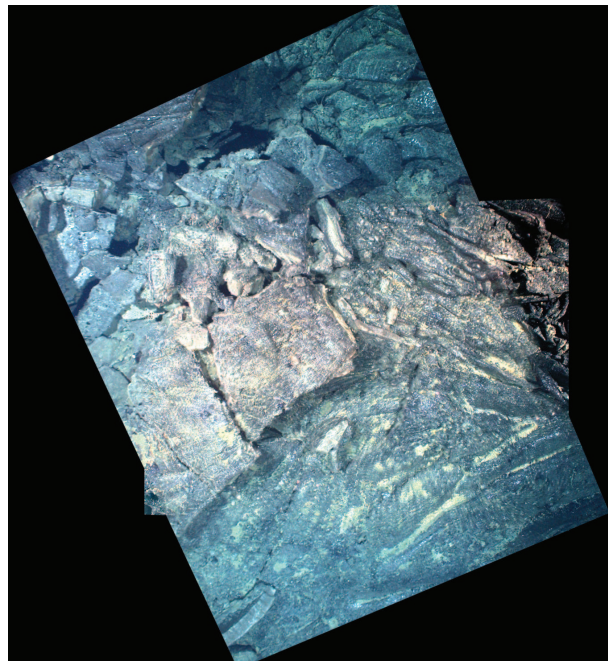
(a)
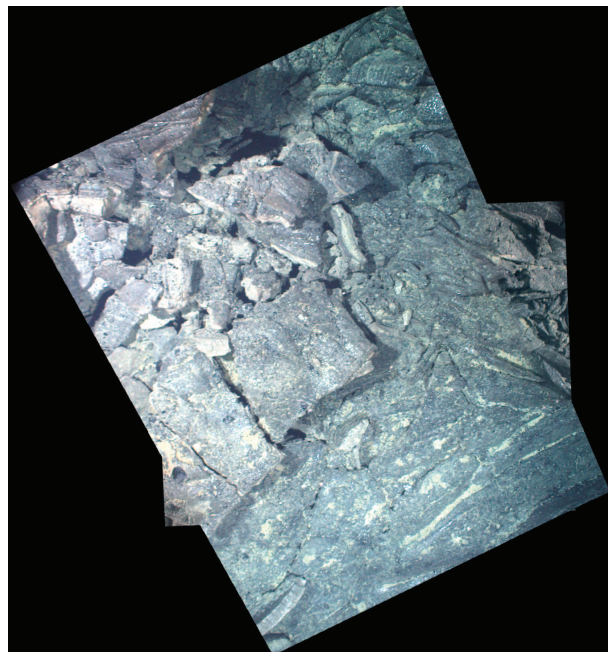
(b)

Fig. 7. Comparison between the results obtained by (a) a multiband blending with exposure fusion and (b) the approach being developed, which applies algorithm improvements allowing to reduce the effects of non-uniform illumination, light attenuation and missregistration. (Images courtesy of Dan Fornari (Woods-Hole Oceanographic Institution))

[5] A. Levin, A. Zomet, S. Peleg, and Y. Weiss. Seamless image stitching in the gradient domain. In *Proc. of the European Conference on Computer Vision (ECCV04)*, Prague, Czech Republic, May 2004.

[6] T. Porter and T. Duff. Compositing digital images. In *SIGGRAPH '84: Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, pages 253–259, New York, NY, USA, 1984. ACM.

[7] M. Uyttendaele, A. Eden, and R. Szeliski. Eliminating ghosting and exposure artifacts in image mosaics. In *Proc. Int. Conf. on Comp. Vision and Patt. Recog. (CVPR2001)*, pages 509–516, 2001.

[8] J. Davis. Mosaics of scenes with moving objects. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, Santa Barbara, CA, USA, June 1998.

[9] A. Efros and W. Freeman. Image quilting for texture synthesis and transfer. *Proceedings of SIGGRAPH 2001*, pages 341–346, August 2001.

[10] D.L. Milgram. Adaptive techniques for photomosaicking. *IEEE Transactions on Computers*, C-26(11):1175–1180, Nov. 1977.

[11] S. Peleg. Elimination of seams from photomosaics. In *Pattern Recognition and Image Processing*, volume 16, pages 90–94, 1981.

[12] P. Burt and E. Adelson. A multiresolution spline with application to image mosaics. *ACM Trans. Graph.*, 2(4):217–236, 1983.

[13] C.T. Hsu and J.L. Wu. Multiresolution mosaic. *Consumer*, 42(4):981–990, November 1996.

[14] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.

[15] N. Gracias, M. Mahoor, S. Negahdaripour, and A. Gleason. Fast image blending using watersheds and graph cuts. *Image and Vision Computing*, 27:597–607, April 2009.

[16] P. Pérez, M. Gangnet, and A. Blake. Poisson image editing. *ACM Transactions on Graphics (SIGGRAPH'03)*, 22(3):313–318, 2003.

[17] A. Bugeau, Ma. Bertalmío, V. Caselles, and G. Sapiro. A comprehensive framework for image inpainting. *IEEE Transactions on Image Processing*, 19(10):2634–2645, 2010.

[18] Jiaya Jia, Jian Sun, Chi-Keung Tang, and Heung-Yeung Shum. Drag-and-drop pasting. *ACM Transactions on Graphics (SIGGRAPH)*, 2006.

[19] James Hays and Alexei A Efros. Scene completion using millions of photographs. *ACM Transactions on Graphics (SIGGRAPH 2007)*, 26(3), 2007.

[20] A. Agarwala, M. Dontcheva, M. Agrawala, S. Drucker, A. Colburn, B. Curless, D. Salesin, and M. Cohen. Interactive digital photomontage. In *Proc. SIGGRAPH04*, August 2004.

[21] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *Proc. IEEE Trans. Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, Nov. 2001.

[22] A. Agarwala. Efficient gradient-domain compositing using quadtrees. *ACM Trans. Graph.*, 26(3):94, 2007.

[23] S. Hanan. *Applications of spatial data structures: Computer graphics, image processing, and GIS*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1990.

[24] R. Szeliski, Uyttendaele, M., and D. Steedly. Fast poisson blending using multi-splines. Technical report, Interactive Visual Media, April 2008.

[25] M. Su, W. Hwang, and K. Cheng. Analysis on multiresolution mosaic images. *IEEE Transactions on Image Processing*, 13(7):952–959, July 2004.

[26] F. Gu and Y. Rzhanov. Optimal image blending for underwater mosaics. In *OCEANS 2006*, pages 1–5, Sept. 2006.

[27] Y. Xiong and K. Pulli. Color correction based image blending for creating high resolution panoramic images on mobile devices. In *ACM SIGGRAPH ASIA 2009 Posters*, SIGGRAPH ASIA '09, pages 47:1–47:1, New York, NY, USA, 2009. ACM.

[28] A. Mills and G. Dudek. Image stitching with dynamic elements. *Image Vision Comput.*, 27(10):1593–1602, 2009.

[29] W. Zhao. Flexible image blending for image mosaicing with reduced artifacts. *International Journal of Pattern Recognition and Artificial Intelligence*, 20(4):609–628, 2006.

[30] L. Neumann, K. Matkovic, and W. Purgathofer. Automatic exposure in computer graphics based on the minimum information loss principle. *Computer Graphics International Conference*, 0:666–667, 1998.

[31] M. Brown and D. G. Lowe. Recognising panoramas. In *ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, page 1218, Washington, DC, USA, 2003. IEEE Computer Society.

[32] A. Eden, M. Uyttendaele, and R. Szeliski. Seamless image stitching of scenes with large motions and exposure differences. In *CVPR '06: Proc. of the 2006 IEEE Computer Vision and Pattern Recognition*, pages 2498–2505, Washington, DC, USA, 2006. IEEE Computer Society.

[33] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: From error visibility to structural similarity, 2004.

[34] R. García B. Möller. Towards objective quality assessment of image registration results. In *Proc. of VISAPP 2007*, 2007.

[35] L. Korhonen J. Boutellier, O. Silvén and M. Tico. Evaluating stitching quality. In *Proc. of VISAPP 2007*, 2007.

[36] V. Rigaud, J. Opderbecke, P. Siméoni, and Cédric Pitout. Real time operational mosaics with the deep sea ROV VICTOR. In *1st International Workshop on Underwater Robotics for Sea Exploitation and Environmental Monitoring*, Rio de Janeiro, Brasil, October 2001.

[37] R. Garcia, T. Nicosevici, and X. Cufi. On the way to solve lighting problems in underwater imaging. In *OCEANS '02 MTS/IEEE*, volume 2, pages 1018 – 1024 vol.2, oct. 2002.