

Seamless Stitching of Stereo Images for Generating Infinite Panoramas

Tao Yan^{1,2} Zhe Huang² Rynson W.H. Lau² Yun Xu¹

¹University of Science and Technology of China, China

²City University of Hong Kong, Hong Kong

Abstract

A stereo infinite panorama is a panoramic image that may be infinitely extended by continuously stitching together stereo images that depict similar scenes, but are taken from different geographic locations. It can be used to create interesting walkthrough environment. An important issue underlying this application is to seamlessly stitch two stereo images together. Although many methods have been proposed for stitching 2D images, they may not work well on stereo images, due to the difficulty in ensuring disparity consistency. In this paper, we propose a novel method to stitch two stereo images seamlessly. We first apply the graph cut algorithm to compute a seam for stitching, with a novel disparity-aware energy function to both ensure disparity continuity and suppress visual artifacts around the seam. We then apply a modified warping-based disparity scaling algorithm to suppress the seam in depth domain. Experiments show that our stitching method is capable of producing high quality stereo infinite panoramas.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual Reality I.4.9 [IMAGE PROCESSING AND COMPUTER VISION]: Applications—Stereo Panorama;

Keywords: Stereo Panorama, Stereo Image Stitching, Stereo Image Editing, Graph Cut, Mesh Warping

1 Introduction

With the popularity of stereo (or 3D) images/videos, techniques for producing and editing stereo media are attracting a lot of research attention in recent years. It can be expected that stereo images will be widely available in the near future, and it would be interesting to take advantage of these contents to develop virtual reality applications, e.g., photo tourism [Snaveley et al. 2006], photorealistic virtual space [Kaneva et al. 2010] and model-based photograph [Kopf et al. 2008]. However, many of these applications rely on existing 2D image processing techniques, which may not always produce good results on stereo images, mainly due to issues related to the additional depth information arising from stereo images.

Panoramic images may be used in IBR applications for virtual walkthroughs. In this paper, we focus on a special kind of panorama: *infinite panorama* [Chan and Efros 2007]. Infinite panorama is created by stitching together images taken from different sites

but with similar contents. It has two advantages over traditional panoramas: (1) it can utilize the vast amount of images available in the Internet, which are taken from diverse geographic locations, and (2) novel panoramas can be created, which can be extended indefinitely, making it possible to create an imaginary environment for exploration.

To produce an infinite panorama given an image database, three issues must be addressed: image retrieval, perspective transformation and image stitching. Kaneva et al. [Kaneva et al. 2010] propose a complete framework to address all three issues. However, the method is primarily designed for 2D images. To apply the framework to stereo images, the techniques used in these stages must be updated. While Feng et al. [Feng et al. 2011] propose an image retrieval technique for stereo images, Du et al. [Du et al. 2013] propose a method to adjust the perspective of stereo images. However, stitching of stereo images remains an issue.

Although there are stitching methods available for 2D images, extending them to stereo images will need to address some new challenges: (1) the resulting image should be smooth around the stitching boundary in both color and disparity so that no obvious artifacts could be observed, and (2) the overall disparity of the resulting image should be consistent so that it will appear to depict a unified scene, which is particularly important as the original images are not taken from the same location using the same camera setting.

In this paper, we propose a method for stitching stereo images together to form a stereo panorama. Unlike traditional works for generating panoramas [Szeliski and Shum 1997][Shum and Szeliski 1998], the stereo images in our case are assumed to have been taken from different geographic locations, and hence perfect alignment is unlikely. Our objective is to stitch them seamlessly so as to synthesize an infinite panorama for use in a walkthrough environment. Unlike [Du et al. 2013], we do not know the camera's intrinsic and extrinsic parameters, since our images are assumed to be obtained from the Internet.

Our method contains three steps. First, we use graph cuts to find an optimal seam to stitch a pair of stereo images. Second, we adjust the disparity range of one stereo image so as to keep it consistent with the other. Finally, we blend the two stereo images together using gradient domain optimization [Pérez et al. 2003]. The main contributions of this paper can be summarized as follows:

- We propose a disparity-aware energy function to be used with the graph cut algorithm, which not only ensures disparity continuity, but also suppresses the visual artifacts around the seam through depth constraint.
- We propose a revised warping-based disparity scaling technique, which ensures that the overall disparity of the resulting image is consistent.

The rest of this paper is organized as follows. Section 2 summarizes related work on stereo image editing and 2D image stitching. Section 3 introduces our stereo image stitching method. Section 4 presents some experimental results to demonstrate the effectiveness of the proposed method. Finally, Section 5 briefly concludes this work and discusses possible future work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

VRST'13, October 06 - 09 2013, Singapore, Singapore
Copyright 2013 ACM 978-1-4503-2379-6/13/10...\$15.00.
<http://dx.doi.org/10.1145/2503713.2503729>




Figure 1: A panorama produced by the proposed method, consisting of five separate stereo images.

2 Related Work

In this section, we first discuss recent stereo image/video editing and processing methods. We then summarize 2D image stitching methods. Finally, we briefly discuss works on creating panoramas.

2.1 Stereo Image Editing

Lang et al. [Lang et al. 2010] propose a disparity mapping method with four simple operators for adjusting the disparity range of stereo videos. However, their method does not consider the relationship between depth perception and viewing parameters, such as display distance and size. To overcome the display adaptation problem, Yan et al. [Yan et al. 2013b] propose a depth mapping method that applies 3D image analysis techniques to remap the depth range of stereo videos according to the viewing parameters. Their method also attempts to preserve 3D scene structures and enforce depth and temporal coherences through a global optimization algorithm. Later, Yan et al. [Yan et al. 2013a] propose to extend the shift-map method to support stereo image editing, such as depth mapping and object depth adjustment. This method processes both the left and right images simultaneously to preserve photo consistence and minimize distortion in image editing. It can handle occlusion and dis-occlusion introduced by the editing operations.

There are also some works that develop specific stereo image editing techniques. Luo et al. [Luo et al. 2012] present a method for seamless stereo image cloning, which involves shape adjustment of the target object according to the depth and the color of the background stereo image. Tong et al. [Tong et al. 2013] propose a depth-consistent stereo image composition method, which allows interactive blending of a 2D image object on a stereo image. Niu et al. [Niu et al. 2012] extend the traditional 2D image warping technique for stereo images. Their objective is to preserve prominent objects and 3D scene structure during the warping operation. Basha et al. [Basha et al. 2013] extend the seam carving method for stereo image resizing. Du et al. [Du et al. 2013] addresses the problem of manipulating perspective in stereoscopic pairs. However, none of these works consider stitching of stereo images.

2.2 2D Image Stitching

Image stitching has been widely studied and has many applications. In general, it has two main steps. First, the two input images are registered to determine their overlapping region. Second, an algorithm is applied on the overlapping region to combine the two images. Image registration is a long standing topic and attracted decades of research. The reader may refer to [Zitova and Flusser 2003] for a survey on relevant techniques. In this paper, we focus on the second step, which is to combine the two registered images into a single panoramic image. The simplest approach to this problem is to alpha-blend the two images in the overlapping region, which

however may produce various artifacts [Levin et al. 2004], such as ghosting. A better approach is to find a one-pixel width seam within the overlapping region to stitch the images together, while minimizing some objective function to suppress artifacts. In general, there are two ways to determine the seam: dynamic programming and graph cuts.

Dynamic programming. It was first used in seam finding by Milgram [Milgram 1977]. The method finds a path that cuts through the overlapping region, minimizing pixel differences of the two images along the seam. It is efficient and guarantees to find an optimal solution. It has been applied in texture synthesis [Efros and Freeman 2001], panorama photography [Davis 1998][Summa et al. 2012] and exploration of image collection [Kaneva et al. 2010]. However, dynamic programming has restrictions on the shape of the seam, which cannot make a 90-degree turn in the middle. It also imposes a grid structure on the pixels, forbidding arbitrary graph topology [Kwatra et al. 2003].

Graph cuts. The graph cut technique [Boykov et al. 2001] assigns a label from a set of k labels to each node of a graph in order to minimize an energy function. Though the k -labeling problem is in general NP-hard [Boykov et al. 2001], the graph cut technique computes an approximate solution that is usually good enough. The energy function to be minimized involves a data term and a pairwise term, which can be used to formulate lots of problems in computer vision. Agarwala et al. [Agarwala et al. 2004] and Kwatra et al. [Kwatra et al. 2003] use graph cuts to synthesize a new image through stitching multiple images together. Gracias et al. [Gracias et al. 2009] propose to use watershed transform to segment the input images into region and then apply graph cuts to label the region boundaries. The advantage of graph cuts over dynamic programming is that it accepts a unary data term in the energy function, which is used in our method to incorporate depth constraint.

In practice, images may have large differences in illumination or color, where visible seam is inevitable [Levin et al. 2004]. In such a situation, gradient-domain optimization [Pérez et al. 2003] may be used as a post-process to eliminate the seam.

2.3 Digital Panorama

Digital panorama is a popular application of image mosaics. It has attracted a large body of works. Systems for creating panoramas have been presented by Szeliski and Shum [Szeliski and Shum 1997], and Shum and Szeliski [Shum and Szeliski 1998], where individual images are aligned and stitched together. To support its creation, various alignment and stitching methods have been proposed [Szeliski 2004]. The creation of stereo panoramas has also been investigated by Huang and Hung [Huang and Hung 1998], Peleg and Ben-Ezra [Peleg and Ben-Ezra 1999], and Peleg et al. [Peleg et al. 2001]. These works aim at creating traditional panoramas [Wikipedia 2013], where images are taken at the same loca-

tion. To make use of the large collection of images available from the Internet, Chan and Efros [Chan and Efros 2007] propose to create panoramas using images taken from different locations. Kaneva et al. [Kaneva et al. 2010] propose a system to automatically choose and place the images to be stitched. Panoramas created in this way can be infinitely long, and are referred to them as “infinite panoramas”. However, these methods for creating infinite panoramas are developed for 2D images. Extending them for stereo images may not be straightforward. Our work addresses exactly this limitation.

3 Stereo Image Stitching

To present our method, we first denote $S^1 = (I_L^1, I_R^1)$ and $S^2 = (I_L^2, I_R^2)$ as two input stereo image pairs, and (D_{LR}^1, D_{RL}^1) and (D_{LR}^2, D_{RL}^2) as their corresponding disparity maps, respectively. By definition, if $p \in I_L$ and $q \in I_R$ are a pair of corresponding pixels, we have $q - p = D_{LR}(p) = D_{RL}(q)$. Our objective is to stitch S^1 and S^2 together to obtain a stereo panoramic image. Since S^1 and S^2 may have different disparity ranges, we use (D_{min}^1, D_{max}^1) and (D_{min}^2, D_{max}^2) to describe their disparity ranges.

As the two stereo images S^1 and S^2 should overlap with each other after registration, we let O_L be the overlapping region between I_L^1 and I_L^2 as shown in Figure 2, and O_R between I_R^1 and I_R^2 . We assume that the two pairs of images overlap in approximately the same way, such that O_L and O_R have the same width. We use graph cuts to compute a label map U_L for designating pixels in O_L to either I_L^1 or I_L^2 , and then compute a label map U_R in O_R using the projection of U_L (via D_{LR}^1) as a constraint. A warping-based scaling algorithm is then applied to remap the disparity maps so that the overall disparity of the final panorama will be consistent. Finally, a gradient domain optimization algorithm is used to suppress color differences along the two seams, seamlessly blending the two pairs of images into a single pair of panorama.

In the following subsections, we present the details of seam finding and warping-based disparity scaling.

3.1 Seam Finding

To stitch two stereo images together, we need to find a seam in the overlapping region, which is done by assigning binary labels to pixels in the overlapping region. A label map is first obtained for O_L , which will then be used as a constraint on the labeling of O_R . We first describe the labeling procedure for O_L , assuming I_L^1 is to be stitched to the left side of I_L^2 .

3.1.1 Labeling O_L

The labeling for O_L is defined as follows. Pixels labelled with 1 and 2 will be copied from I_L^1 and I_L^2 , respectively, to the stitched image. We use graph cuts to solve this labeling problem. The graph cut algorithm requires the definition of a data cost E_{data} and a smoothness cost E_{smooth} as prior information. It then minimizes the total energy term to obtain a near optimal labeling. Our energy function is defined as follows:

$$\begin{aligned} E_{total} &= E_{data} + w * E_{smooth} \\ &= \sum e_d(p, l_p) + w * \sum_{u, v \in N} e_s(u, v, l_u, l_v) \end{aligned} \quad (1)$$

where $e_d(p, l_p)$ denotes the cost of pixel p being assigned a label l_p , and $e_s(u, v, l_u, l_v)$ denotes the penalty of two neighboring pixels u and v being assigned labels l_u and l_v , respectively. w is a weighting parameter.

Boundary Constraint: Intuitively, pixels to the left of the seam should come from I_L^1 , and those to the right from I_L^2 . Such a constraint is imposed to encourage pixels on B_L^2 , which is the left boundary of I_L^2 as shown in Figure 2, to be obtained from I_L^1 , and similarly pixels on B_L^1 from I_L^2 :

$$e_{by}(p, l_p) = \begin{cases} 0, & \text{if } p \in B_L^2, l_p = 1, \text{ or } p \in B_L^1, l_p = 2 \\ K_1, & \text{otherwise} \end{cases} \quad (2)$$

where K_1 is a large penalty value.




Figure 2: Overlapping region, O_L , of the two images, I_L^1 and I_L^2 , to be stitched.

Depth Constraint: Normally, discontinuity is unavoidable when stitching two images of different scenes. This is because the objects along the seam are likely different in the two images. With stereo images, we may utilize the depth information to suppress the artifacts by retaining the pixels near to the viewer, instead of those that are further away, in the output image. After two stereo images are stitched together, they will share a common disparity range. So, we first normalize their input disparity maps to the range of $[0, 1]$. (We take the assumption that objects near to the cameras, i.e., with smaller depth values, have smaller normalized disparity values.) We denote D_p^1 and D_p^2 as the normalized disparity values of point p of I_L^1 and I_L^2 , respectively, located within the overlap region O_L . The constraint is formulated as follows:

$$e_{dh}(p, 1) = k_1 e^{k_2 (D_p^1 - D_p^2)} \quad (3)$$

$$e_{dh}(p, 2) = k_1 e^{k_2 (D_p^2 - D_p^1)} \quad (4)$$

where k_1 and k_2 are positive weighting parameters, and the disparity values are normalized to $[0, 1]$ by $D_i = (D_{LR}(i) - D_{min}) / (D_{max} - D_{min})$. Figure 7 demonstrates the effect of depth constraint.

Smoothness Cost: In general, neighboring pixels across the seam should maintain color and disparity similarities to prevent visual artifacts appearing in the stitched image. As such, the smoothness cost is defined to penalize color, gradient and disparity discontinuities around the seam as follows:

$$\begin{aligned} e_s(u, v, l_u, l_v) &= (|I_u^{l_u} - I_v^{l_v}| + |I_v^{l_u} - I_v^{l_v}|) \\ &\quad + w_G (|G_u^{l_u} - G_v^{l_v}| + |G_v^{l_u} - G_v^{l_v}|) \\ &\quad + w_D (|D_u^{l_u} - D_v^{l_v}| + |D_v^{l_u} - D_v^{l_v}|) \end{aligned} \quad (5)$$

where u and v are neighboring pixel pair (4-connected). l_u and l_v are the label values of u and v . They determine which of the two images that we are referring to. The first term tries to ensure color continuity. In the second term, G^1 and G^2 are the gradient maps of the two input images. This term is to ensure gradient continuity. In the third term, D^1 and D^2 are the disparity maps of the two input images, which are normalized to the range of $[0, 255]$. This term is to ensure disparity continuity. w_G and w_D are two weighting parameters.

Combining all the terms discussed above, the total cost of a particular labeling in O_L is then:

$$E_{L,total} = \sum_{p \in O_L} e_{dh}(p, l_p) + \sum_{p \in B_L^1 \text{ or } B_L^2} e_{by}(p, l_p) + w * \sum_{u, v \in N_L} e_s(u, v, l_u, l_v) \quad (6)$$

where N_L is the set of four-connected neighboring pixel pairs in the left image. B_L^1 and B_L^2 are the two image boundaries shown in Figure 2.

3.1.2 Labeling O_R

Given the labeling for O_L , the labeling of O_R has an additional constraint – the corresponding pixels in O_R and O_L should receive the same labels. For a point q in O_R , its corresponding pixel in O_L is computed as $q - D_{RL}(q)$, where D_{RL} can be either D_{RL}^1 or D_{RL}^2 depending on the label assigned to q . The cost function is defined as:

$$e_{gu}(q, l_q) = \begin{cases} 0, & \text{if } O_L(q - D_{RL}^{l_q}(q)) = l_q \\ K_2, & \text{otherwise} \end{cases} \quad (7)$$

where K_2 is a positive penalty value.

By adding the above cost, the total cost of a particular labeling in O_R is then:

$$E_{R,total} = \sum_{q \in O_R} e_{dh}(q, l_q) + \sum_{q \in B_R^1 \text{ or } B_R^2} e_{by}(q, l_q) + \sum_{q - D_{RL}^{l_q}(q) \in O_L} e_{gu}(q, l_q) + w * \sum_{u, v \in N_R} e_s(u, v, l_u, l_v) \quad (8)$$

where N_R is the set of four-connected neighboring pixel pairs in the right image. After we have computed a labeling for O_L and O_R , we can obtain the respective seams S_L and S_R .

3.2 Warping-based Disparity Scaling

The stereo images to be stitched may have very different disparity ranges, leading to noticeable discontinuity in depth when viewed in 3D. Figure 6 shows such an example. To suppress this kind of artifact, we linearly scale one of the disparity maps so that it is consistent with the other and then warp the images accordingly. Specifically, we linearly scale the disparity map D_{LR}^2 so that it can be seamlessly stitched with D_{LR}^1 at S_L . As the disparity values in D_{LR}^2 depend on both I_L^2 and I_R^2 , we need to warp them appropriately. Since disparity-based warping is non-trivial, we use a modified version of the method proposed by Yan et al. [Yan et al. 2013b] to accomplish our goal here.

Consider a pair of adjacent pixels u and v along the seam S_L , where $u \in I_L^1$ and $v \in I_L^2$. We first need to determine whether u and v should have the same depth in the stitched stereo image. Pixels with similar color and disparity are likely to belong to the same object in the real scene, and we will set them to the same depth value. As such, if $|I_L^1(u) - I_L^2(v)| < T_c$ and $|D_{LR}^1(u) - D_{LR}^2(v)| < T_d$, u and v are considered to have the same disparity in the stitched image. Hence, we assign the disparity value of u to v . Let the original disparity of v be d_v and the new disparity be $d'_v = d_u$, we assume that they are related by $d'_v = ad_v + b$, where a and b can be determined by fitting all possible (d_v, d'_v) to a linear function. The resulting function $f(d) = ad + b$ is used to scale all other disparity values of D_{LR}^2 , where $d \in D_{LR}^2$.

The linear mapping function transforms D_{LR}^2 into a new disparity map, with which we can warp I_L^2 and I_R^2 to obtain a new stereo image. Let $p \in I_L^2$ and $q \in I_R^2$ be a pair of corresponding pixels whose locations after warping are p' and q' , and the x coordinates of p and q be x_p and x_q . Here, we use two constraints to warp the images based on the new disparity map.

First, pixels along the seam S_L are forced not to shift from their original positions. This constraint is formulated as:

$$E_c = \sum_{p \in S_L} |x_{p'} - x_p|^2 \quad (9)$$

Second, all pairs of corresponding pixels should be scaled to their target disparity, which can be formulated as:

$$E_m = \sum |(x_{q'} - x_{p'}) - f(D_{LR}^2(p))|^2 \quad (10)$$

Other energy terms used to preserve stereo image features are defined similarly as those in [Yan et al. 2013b]. Then, we solve the mesh warping problem like [Yan et al. 2013b] to obtain the new stereo image pair I_L^2 and I_R^2 .

Note that the images may have small perspective differences or orientation variations due to different camera settings. One may use the methods proposed by [Du et al. 2013] and [Luo et al. 2012], respectively, to correct them before the disparity scaling operation.

Finally, after we have obtained the stitched images, the seam may still be visible due to different lightings, contrasts, textures, etc.. We apply gradient domain optimization [Pérez et al. 2003] to smooth the stitched images, achieving a seamless stereo panorama.

4 Experiments and Discussion

We have implemented the proposed method in C with OpenCV and Intel MKL. We tested the program on a PC with an Intel i5 3.1GHz quad-core CPU and 8GB RAM. Parameters utilized in this paper are experimental set as $w = 1$, $K_1 = 10^5$, $K_2 = 10^2$, $k_1 = 1$, $k_2 = 8$, $W_g = 5$, $W_D = 1$, $T_c = 80$ and $T_d = 60$. All stereo images shown in our paper were downloaded from Yahoo flickr album. The stereo camera parameters of these images are not known. All stereo images are shown in red-cyan format in this paper.

Figures 1, 3 and 4 show three stereo infinite panoramas. They are produced by stitching the images along each row of Figure 8. Figure 3 shows a 3D natural scene with rivers, forests, mountains, bridge and rocks. It is synthesized from the input stereo images shown in the second row of Figure 8. The difficulty with this set of images comes from the presence of different objects near the boundaries of individual input images, and discontinuity is inevitable. Our depth constraint successfully suppresses some of the artifacts illustrated in Figure 7. Although color and texture differences still exist along the seam, they are weakened to some extent through gradient domain optimization [Pérez et al. 2003].

In Figure 1, a stereo panorama with canyon, rivers, plants, and a bridge is shown. It is synthesized from the input stereo images shown in the first row of Figure 8. These input images have diverse depth ranges. For example, the bridge in the first image is very close to the viewer and the rocks in the background are far away. However, the second image has a relatively much smaller depth range. If these two images are stitched directly, depth discrepancy can be observed as shown in the right image of Figure 6. By applying warping-based disparity scaling, such problem can be avoided as shown in the left image of Figure 6.




Figure 3: A panorama of rivers, forests and mountains.



Figure 4: A panorama of lawn, house, monument, flagpole and railway

Figures 4 and 5 show additional results produced by our method. We can see that our method can effectively stitch fairly complex scenes together, even though the contents are very different.

5 Conclusion and Future Work

In this paper, we have proposed a seamless stitching method for stereo images. Our method contains two main steps: seam finding and warping-based disparity scaling. We have addressed issues specific to stereo images in both steps. Experiments show that our method is effective in joining stereo images, while suppressing most visual artifacts.

However, our stitching method is not without its shortcomings. For example, our warping-based disparity scaling technique simply scales the disparity maps of two images to be stitched to the same depth range. This is only an approximation. If the depth ranges of the two images differ significantly, they usually represent two completely different scenes and may not be stitched together. Currently, we simply do not allow the images to be stitched if their content, perspective or orientation differ too much. A better method may involve analysing the contents of stereo images and changing their perspective, which may be worth looking at as a future work.

By repeatedly applying our method, a stereo infinite panorama can be produced from a set of stereo images. This may be a way for creating a virtual walkthrough environment [Kaneva et al. 2010] for stereoscopic display. We have shown that our method can produce good results even if the input images have very different contents.

Acknowledgements

The work presented in this paper was partially supported by a S-RG grant from City University of Hong Kong (Project Number: 7002768) and two National Natural Science Foundations of China (Project Numbers: 61033009 and 60970085).

References

- AGARWALA, A., DONTCHEVA, M., AGRAWALA, M., DRUCKER, S., COLBURN, A., CURLESS, B., SALESIN, D., AND COHEN, M. 2004. Interactive digital photomontage. *ACM Trans. on Graphics* 23, 3 (Aug.), 294–302.
- BASHA, T., MOSES, Y., AND AVIDAN, S. 2013. Stereo seam carving a geometrically consistent approach. *IEEE Trans. on Pattern Analysis and Machine Intelligence (to appear)*.
- BOYKOV, Y., VEKSLER, O., AND ZABIH, R. 2001. Fast approximate energy minimization via graph cuts. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 23, 11, 1222–1239.
- CHAN, L., AND EFROS, A. 2007. Automatic generation of an infinite panorama. Tech. rep., Carnegie Mellon University, Dec.
- DAVIS, J. 1998. Mosaics of scenes with moving objects. In *Proc. IEEE CVPR*, 354–360.
- DU, S., HU, S., AND MARTIN, R. 2013. Changing perspective in stereoscopic images. *IEEE Trans. on Visualization and Computer Graphics*.
- EFROS, A., AND FREEMAN, W. 2001. Image quilting for texture synthesis and transfer. In *Proc. ACM SIGGRAPH*, 341–346.
- FENG, Y., REN, J., AND JIANG, J. 2011. Generic framework for content-based stereo image/video retrieval. *Electronics Letters* 47, 2, 97–98.
- GRACIAS, N., MAHOOR, M., NEGAHDARIPOUR, S., AND GLEASON, A. 2009. Fast image blending using watersheds and graph cuts. *Image and Vision Computing* 27, 5 (Apr.), 597–607.
- HUANG, H., AND HUNG, Y. 1998. Panoramic stereo imaging system with automatic disparity warping and seaming. *Graphical Models and Image Processing* 60, 3 (May), 196–208.
- KANEVA, B., SIVIC, J., TORRALBA, A., AVIDAN, S., AND FREEMAN, W. 2010. Infinite images: Creating and exploring a




Figure 5: A panorama of buildings and walls.




Figure 6: Stitching with (left) and without (right) disparity scaling. Note that as the two input images have very different disparity ranges, the stitched image shown on the right has noticeable sudden change of depth across the seam.



Figure 7: Stitching with (left) and without (right) depth constraint. On the left, more pixels of the foreground (trees) are shown in the stitched image, where discontinuity may be interpreted as occlusion. On the right, such discontinuity is perceived as artifact.

large photorealistic virtual space. *Proceedings of the IEEE* 98, 8, 1391–1407.

KOPF, J., NEUBERT, B., CHEN, B., COHEN, M., COHEN-OR, D., DEUSSEN, O., UYTENDAELE, M., AND LISCHINSKI, D. 2008. Deep photo: model-based photograph enhancement and viewing. In *Proc. ACM SIGGRAPH Asia*.

KWATRA, V., SCHODL, A., ESSA, I., TURK, G., AND BOBICK, A. 2003. Graphcut textures: image and video synthesis using graph cuts. *ACM Trans. on Graphics* 22, 3 (July), 277–286.

LANG, M., HORNUNG, A., WANG, O., POULAKOS, S., SMOLIC, A., AND GROSS, M. 2010. Nonlinear disparity mapping for stereoscopic 3D. *ACM Trans. on Graphics* 29, 3.

LEVIN, A., ZOMET, A., PELEG, S., AND WEISS, Y. 2004. Seamless image stitching in the gradient domain. In *Proc. ECCV*, no. 3024. Jan., 377–389.

LUO, S., SHEN, I., CHEN, B., CHENG, W., AND CHUANG, Y. 2012. Perspective-aware warping for seamless stereoscopic image cloning. *ACM Trans. on Graphics* 31, 6 (Nov.).

MILGRAM, D. 1977. Adaptive techniques for photomosaicking. *IEEE Trans. on Computers* C-26, 11, 1175–1180.

NIU, Y., FENG, W., AND LIU, F. 2012. Enabling warping on stereoscopic images. *ACM Trans. on Graphics* 31, 6 (Nov.).

PELEG, S., AND BEN-EZRA, M. 1999. Stereo panorama with a single camera. In *Proc. IEEE CVPR*, vol. 1.

PELEG, S., BEN-EZRA, M., AND PRITCH, Y. 2001. Omnistereo: panoramic stereo imaging. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 23, 3, 279–290.

PÉREZ, P., GANGNET, M., AND BLAKE, A. 2003. Poisson image editing. *ACM Trans. on Graphics* 22, 3 (Jul.), 313–318.



Figure 8: *Input stereo images utilized to produce the stereo panoramas shown in this paper.*

- SHUM, H., AND SZELISKI, R. 1998. Construction and refinement of panoramic mosaics with global and local alignment. In *Proc. ICCV*, 953–956.
- SNAVELY, N., SEITZ, S., AND SZELISKI, R. 2006. Photo tourism: exploring photo collections in 3D. *ACM Trans. on Graphics* 25, 3 (July), 835–846.
- SUMMA, B., TIERNY, J., AND PASCUCCHI, V. 2012. Panorama weaving: fast and flexible seam processing. *ACM Trans. on Graphics* 31, 4 (July).
- SZELISKI, R., AND SHUM, H. 1997. Creating full view panoramic image mosaics and environment maps. In *Proc. ACM SIGGRAPH*, 251–258.
- SZELISKI, R. 2004. Image alignment and stitching: A tutorial. Tech. rep., MSR-TR-2004-92, Microsoft Research.
- TONG, R., ZHANG, Y., AND CHENG, K. 2013. Stereopasting: Interactive composition in stereoscopic images. *IEEE. Trans. on Visualization and Computer Graphics* 19, 8 (Aug), 1375–1385.
- WIKIPEDIA, 2013. Panoramic photography.
- YAN, T., HE, S., LAU, R., AND XU, Y. 2013. Consistent stereo image editing. In *Proc. ACM Multimedia*.
- YAN, T., LAU, R., XU, Y., AND HUANG, L. 2013. Depth mapping for stereoscopic videos. *International Journal of Computer Vision* 102, 1-3 (Mar.), 293–307.
- ZITOVA, B., AND FLUSSER, J. 2003. Image registration methods: a survey. *Image and Vision Computing* 21, 11 (Oct.), 977–1000.