Enabling Warping on Stereoscopic Images

Yuzhen Niu Wu-Chi Feng Feng Liu
Portland State University

(a) Input (b) Basic rotation result and disparities (c) Our result

Figure 1: *Stereoscopic image warping. Applying the same warping to the left and right image of a stereoscopic image often introduces vertical disparities, as shown in (b). Our method enables existing image warping algorithms on stereoscopic images. (c) shows that our warping result is free from vertical disparities. In this paper, we show stereoscopic images using the red-cyan anaglyph representation.*

Abstract

Warping is one of the basic image processing techniques. Directly applying existing monocular image warping techniques to stereoscopic images is problematic as it often introduces vertical disparities and damages the original disparity distribution. In this paper, we show that these problems can be solved by appropriately warping both the disparity map and the two images of a stereoscopic image. We accordingly develop a technique for extending existing image warping algorithms to stereoscopic images. This technique divides stereoscopic image warping into three steps. Our method first applies the user-specified warping to one of the two images. Our method then computes the target disparity map according to the user specified warping. The target disparity map is optimized to preserve the perceived 3D shape of image content after image warping. Our method finally warps the other image using a spatially-varying warping method guided by the target disparity map. Our experiments show that our technique enables existing warping methods to be effectively applied to stereoscopic images, ranging from parametric global warping to non-parametric spatially-varying warping.

CR Categories: I.4.m [Image Processing and Computer Vision]: Miscellaneous;

Keywords: Stereoscopic image warping, disparity mapping

Links: DL PDF

1 Introduction

These years we have been observing a tremendous resurgence of interest in stereoscopic 3D. A variety of stereoscopic displays and cameras are available. This brings in the demand for tools for authoring and processing stereoscopic content. However, extending existing tools to stereoscopic content is often non-trivial as stereoscopic content has an extra dimension of disparity that needs to be correctly taken care of to deliver a pleasant viewing experience.

This paper focuses on stereoscopic image warping. Warping is one of the basic image processing techniques and a wide range of image warping methods have been developed [Wolberg 1990]. Applying the same warping to the left and right image of a stereoscopic image can transform them consistently; however this straightforward solution is often problematic. Figure 1 shows a simple 2D image transformation: rotation. If we rotate the left and right image with the same rotation matrix, we introduce vertical disparities, which often bring "3D fatigue" to viewers [Mendiburu 2009]. Moreover, the original horizontal disparity distribution is changed.

This paper presents a technique that extends existing image warping algorithms to stereoscopic images. Our idea is to warp one of the two images of a stereoscopic image using the user-specified warping and warp the other to both follow the user-specified warping and meet the disparity requirement. Our technique consists of three steps. We first apply the user-specified warping to one of the two images. Without loss of generality, we always warp the left image using the user-specified warping. We then compute the target disparity map according to the user-specified warping. We consider that a good target disparity map should be consistent with the warping applied to the input image and maintain the perceived roundness of the image objects. For example, if an object is stretched in the image space, it should also be stretched in depth. If the disparity remains the same, the warped image may suffer from the *cardboarding* artifacts where the perceived object becomes flattened [Mendiburu 2009]. Based on this observation, we develop an automatic disparity mapping technique that scales the local disparity range according to how this region is warped. Thus, the target disparity map is optimized to preserve the perceived 3D shape of image content after warping.

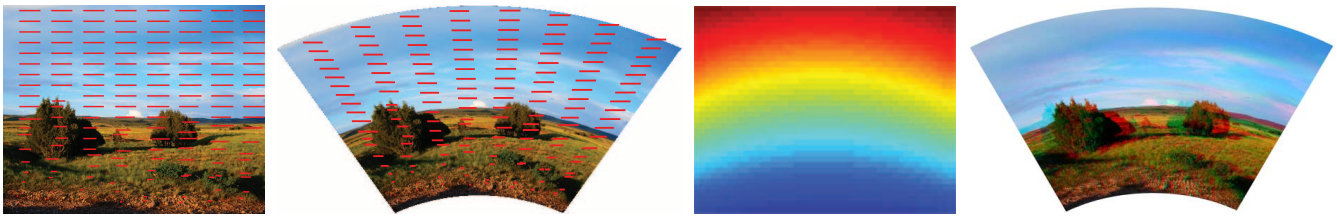

(a) Input left image/disparity (b) Output left image/disparity (c) Scaling map (d) Warping result

Figure 2: *Workflow of stereoscopic image warping. Our method pre-processes the input stereoscopic image and estimates a sparse set of disparities (a). Our method applies the user-specified warping to the left image (b), computes a scaling map that captures the local image stretching (c), and computes the target disparity map (imposed on the left warping result (b)). We use the warm and cool color to encode the large and small scaling factor in (c), respectively. Our method finally warps the right image guided by the target disparity map (d).*

The final step is to warp the right image guided by the target disparity map. This is very similar to the warping step in previous disparity editing research [Lang et al. 2010]. Our warping step, however, is more challenging as the right image needs to be transformed more drastically when the user-specified warping becomes complex. Directly applying the spatially-varying mesh-based warping method from the previous research often cannot follow the user-specified complex warping. We introduce a critical pre-warping step that transforms the warping of the right image into the standard disparity-guided image warping problem which can then be solved using the mesh-based warping method.

The main contribution of this paper is a technique that enables existing monocular image warping methods for stereoscopic images. Our technique is simple and does not change existing warping methods. It is applicable to a variety of image warping techniques, ranging from parametric global warping to non-parametric warping. This paper provides a novel automatic method to compute a good disparity map for the warped image as a key component of our technique. We also introduce a critical pre-warping step to transform the warping of the right image to follow the user-specified complex warping into a disparity-guided warping problem.

2 Related Work

An overview on image warping techniques for 2D images is beyond the scope of this paper. Good surveys on image warping are available [Wolberg 1990; Gomes et al. 1999; Szeliski 2010]. This section focuses on stereoscopic image authoring and processing.

Disparity control is important for creating high-quality stereoscopic images and videos. The disparity needs to be properly distributed so that the scene content exists in the stereoscopic comfort zone [Howard and Rogers 2002; Hoffman et al. 2008]. Algorithms have been developed to determine the camera parameters [Jones et al. 2001; Mueller et al. 2008; Koppal et al. 2011; Zilly et al. 2010] before capture. Recently, Heinzle *et al.* [2011] built a computational stereo camera system that closes the control loop from capture and analysis to automatic setting of these parameters. Oskam *et al.* [2011] developed a system for stereoscopic camera control in interactive 3D applications.

Disparity editing tools have also been developed for post-production. Wang and Sawchuk [2008] presented a framework for disparity editing that either directly works on the dense disparity map or assumes known camera parameters and applies image-based rendering methods to novel view synthesis. Lang *et al.* developed a set of disparity mapping tools to control the disparity distribution in a nonlinear and locally adaptive fashion [Lang et al. 2010; Smolic et al. 2011]. Koppal *et al.* [2011] developed a viewer-centric editor for stereoscopic movies that provides tools for both shot planning and disparity editing. Didyk *et al.* [2011] introduced a perceptual model of disparity for computer graphics and applied to a number of applications. Disparity editing software is also available[1][2][3]. Our work builds on these efforts, as the last step of our method eventually reduces the warping of the right image into the warping of a pre-warped image, and then applies a similar mesh-based warping method to Lang *et al.* [2010].

Some 2D image editing tools have been extended to stereoscopic images. Wang *et al.* [2008a] presented an algorithm for simultaneous color and depth in-painting for stereoscopic images. Lo *et al.* [2010] extended 2D object copy and paste to stereoscopic images. Chang *et al.* [2011], Basha *et al.* [2011], and Lee *et al.* [2012] developed methods for stereoscopic image retargeting. Liu *et al.* [2011] developed a stereoscopic video stabilization technique.

Interesting methods have also been developed for 2D to 3D image and video conversion [Moustakas et al. 2008; Knorr and Sikora 2007; Saxena et al. 2009; Guttmann et al. 2009]. These methods recover depth from the input images and videos and generate stereo content by image-based rendering methods or warping-based methods. Kim *et al.* [2011] recently presented a method to generate stereoscopic views from light fields.

3 Stereoscopic Image Warping

Warping for monocular images can be defined as a function that maps pixel positions from the input to the output. It only works on pixel coordinates in 2D. A stereoscopic image has a third dimension of disparity (depth). One straightforward solution to stereoscopic image warping is to apply the same warping function to both the left and right image. The disparity is implicitly determined. This solution has an advantage that the left and right image are consistently warped, but at the same time has two major limitations. First, vertical disparities can be introduced as the corresponding feature points in the left and right image have different coordinates and are mapped to different locations although the same warping function is used. Second, the disparity map is arbitrarily changed, damaging critical depth perception for viewing.

Our technique allows existing image warping methods to be applied to stereoscopic images. Instead of warping the left and right image directly, our technique works on both the images and the disparity map to properly handle the disparity problems. As shown in Figure 2, our technique for stereoscopic image warping consists of three steps after a pre-processing step of disparity estimation.

1. Warp the left image using the user-specified warping method.
2. Compute the target disparity map.
3. Warp the right image using a spatially-varying warping method guided by the disparity map.

[1] http://stereo.jpn.org/eng/stphmkr/
[2] http://www.thefoundry.co.uk/products/nuke/
[3] http://www.thefoundry.co.uk/products/ocula/

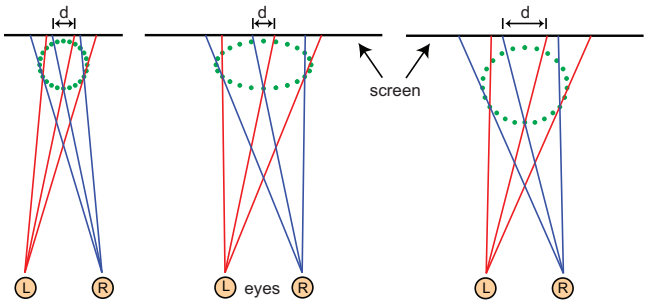

Figure 3: *Disparity and perceived object roundness.* ***Left****: an input image with a ball.* ***Middle****: when the image size is doubled and the disparity remains the same, the ball is perceived flattened.* ***Right****: when the disparity range is doubled (the disparity is doubled in this particular case), the roundness of the ball is roughly preserved.*

Our technique pre-processes the input stereoscopic image to compute its disparity map. Per-pixel disparities can be estimated using dense stereo matching, which is still a challenging computer vision problem [Scharstein and Szeliski 2002; Hirschmuller and Scharstein 2009]. Our method estimates and uses the disparities of a sparse set of feature points using a SIFT-based method [Lowe 2004]. The problem with using the sparse set of feature correspondences is the lack of feature points in textureless regions. Like Lang *et al.* [2010], we estimate dense correspondences from the downsampled version of the input images using an optical flow method [Sun et al. 2010]. We then scale up the optical flow of each pixel in the downsampled image as the disparity of the corresponding pixel in the original image. Since optical flow is typically not as accurate as SIFT feature matches, we weigh SIFT feature matches more than optical flow, as described in Section 3.2. We remove outliers using the fundamental matrix constraint [Hartley and Zisserman 2004].

Our method starts by applying the user-specified image warping method to the left image I^l and creating the warped left image $\hat{I}^l$. This step also applies the user-specified warping to the feature points in the left image and computes their output feature point positions $\hat{p}^l$ in $\hat{I}^l$, which will be combined with the target disparities to compute the output feature positions $\hat{p}^r$ in the right image. Below we describe how our method computes the target disparities as the second step and warps the right image as the third step.

3.1 Optimal Target Disparity Estimation

Warping a stereoscopic image requires to transform its 2D geometric shape and handle its disparity map. While it is typically clear how a user wants to change the image shape, it is less clear for the disparity map. A simple solution is to keep the original disparity map, which often cannot work. As shown in Figure 3, when an image is uniformly stretched, the ball will be perceived as flattened relative to the original if the disparities remain the same. To maintain the roundness of the ball, the perceived depth range of the ball needs to be stretched accordingly. For simple image transformation like scaling, we could use the stereoscopic viewing geometry described in [Guttmann et al. 2009; Oskam et al. 2011] to compute the "correct" target disparity distribution for all the objects in the image and to optimize the target disparity map that best maintains the roundness of these objects. However, this requires the viewing parameters, which are often unknown.

We seek a solution that works well with a wide variety of image warpings and best preserves the perceived 3D shape of image content after warping. In general, the perceived depth does not linearly depend on the depth. But when objects are close to the screen (with small disparities) or the screen size is small, the perceived depth nearly linearly depends on the disparity. These are two common scenarios. To avoid the vergence-accommodation conflict, professional film-makers often give content of interest small disparities. The increasing availability of small consumer stereoscopic devices enables viewing stereoscopic content on small screens. Our idea is then to estimate an optimal target disparity map where the local disparity range changes according to the corresponding local image warping. We estimate this optimal disparity map by minimizing the following energy function that aims to match the scaling of the disparity range with that of the local image warping.

$$\sum_{d_i} \sum_{d_j \in N(d_i)} \|(\hat{d}_i - \hat{d}_j) - s_i(d_i - d_j)\|^2, \quad (1)$$

where d_i and $\hat{d}_i$ are the original disparity and target disparity of feature (pixel) i, respectively. d_j is the disparity of feature j in $N(d_i)$, the neighborhood of feature i. $N(d_i)$ includes all the feature points that are within a small square centered at feature i (typically of size 30×30). s_i is the disparity scaling coefficient, computed according to how the neighborhood around feature i is scaled, as described later on. This energy minimization problem alone is under-determined and requires a boundary condition. To keep the objects with small disparities in the comfort zone after warping, we apply the boundary condition that the disparity with the minimal magnitude is scaled by the corresponding s.

$$\hat{d}_{min} = s d_{min}, \quad (2)$$

where d_{min} is the input disparity that has the minimal magnitude and s is the scaling coefficient computed around this point. The energy minimization problem together with this boundary condition is solved using a sparse linear solver.

We now describe how the scaling factor s_i is computed. Our method defines the neighborhood of a feature as a small square centered at the feature position (typically of size 3×3). Our method computes the output positions of the four corner points using the user-specified warping and estimates an optimal similarity transformation $\bar{H}_s$ that best fits the mapping of the corner points.

$$\bar{H}_s = \arg\min_{H_s} \sum_c \|H_s c - \hat{c}\|^2, \quad (3)$$

where H_s is a similarity transformation, c is one of the four corner points, and $\hat{c}$ is its corresponding output point after applying the user-specified warping. The scaling factor of the best fitting similarity transformation $\bar{H}_s$ is used to set the scaling coefficient s_i. We show an example of this scaling factor map in Figure 2 (c).

When an image is aggressively stretched, our method can create excessive disparities, which can often bring 3D fatigue to viewers [Mendiburu 2009]. This problem can be solved by linearly scaling the disparity map when the disparity range exceeds a predefined threshold. More sophisticated disparity range mapping methods can be further applied [Lang et al. 2010].

3.2 Disparity Guided Right Image Warping

Classic novel view synthesis methods can be used to create the output right image using the input right image and the target disparity map. These methods, however, require a dense disparity map and are sensitive to the disparity accuracy. Inspired by previous research [Krähenbühl et al. 2009; Liu et al. 2009; Lang et al. 2010], we use a spatially-varying warping method instead. Our method divides the input right image into a $m \times n$ uniform grid mesh and formulates the image warping problem as a mesh warping problem. The unknowns are the coordinates of mesh vertices. The mesh

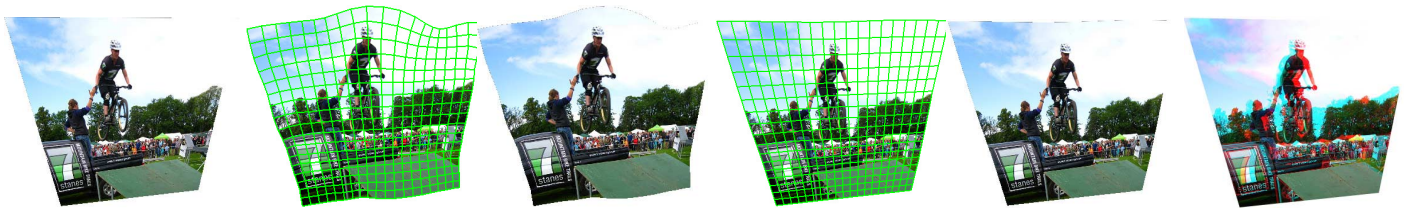

(a) Left warping result (b) Right warping result without pre-warping (c) Right warping result with pre-warping (d) Stereo warping result

Figure 4: *Disparity guided image warping. For this example, a perspective transformation is applied to the left image (a). The spatially varying warping method guided by the disparity map alone cannot warp the right image to be consistent with the left image (b). (c) shows that pre-warping the right image using the same warping applied to the left image solves this problem.*

warping problem is defined as a linear least squares problem that enforces the target disparities on the feature points and minimizes visual distortion. We describe the energy terms below.

Disparity term. Our method encourages each feature point pair $(\hat{\mathbf{p}}_i^l, \hat{\mathbf{p}}_i^r)$ at the warped left and right image to be separated with the target disparity vector $\hat{\mathbf{d}}_i$, where $\hat{\mathbf{d}}_i = [\hat{d}_i \; 0]^T$. Because a feature point is not usually coincident with one of the mesh vertices, our method finds the mesh cell that $\mathbf{p}_i^r$ belongs to and represents $\mathbf{p}_i^r$ with a linear combination of the four vertices of the cell in the original image. The linear combination coefficients are computed using the inverse bilinear interpolation method [Heckbert 1989]. These coefficients are then used to combine the vertices in the output image to compute $\hat{\mathbf{p}}_i^r$. We then define the disparity term as follows.

$$E_d = \sum_{\mathbf{p}_i^r} \kappa_i w_i (\sum \alpha_j \hat{\mathbf{v}}_{i,j} - \hat{\mathbf{p}}_i^l - \hat{\mathbf{d}}_i)^2, \tag{4}$$

where $\hat{\mathbf{v}}_{i,j}$ are the vertices that enclose $\hat{\mathbf{p}}_i^r$ in the right warped image, α_j is the bilinear coefficient, and $\hat{\mathbf{p}}_i^l$ is the feature point's position in the left warped image. κ_i is a parameter that weighs correspondences from SIFT feature matching more than those from optical flow as SIFT feature matching is more robust than optical flow. Our system sets κ_i=1.0 for SIFT feature matches and 0.1 for optical flow matches. w_i is the saliency value of feature i, computed using the graph-based saliency method [Harel et al. 2007]. Our method encourages salient feature points to meet the disparity requirement more than those less salient ones.

Smoothness term. To avoid geometric distortion, our method encourages each cell to undergo a similarity transformation. A similarity transformation that maps (x, y) to (u, v) must satisfy the Cauchy-Riemann equations, namely $\frac{\partial u}{\partial x} = \frac{\partial v}{\partial y}$ and $\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x} = 0$. We use finite differences to compute the partial derivatives and apply this constraint to each vertex $\hat{\mathbf{v}}_{i,j} = (u_{i,j}, v_{i,j})$ in each cell.

$$\begin{aligned} E_s = \quad & \sum_{\hat{\mathbf{v}}_{i,j}} w_{i,j}(u_{i+1,j} - u_{i,j} - v_{i,j+1} + v_{i,j})^2 \\ & + w_{i,j}(u_{i,j+1} - u_{i,j} + v_{i+1,j} - v_{i,j})^2, \end{aligned} \tag{5}$$

where $w_{i,j}$ is the average saliency value of the cell whose top-left corner is vertex $\hat{\mathbf{v}}_{i,j}$.

We combine the above disparity term and smoothness term and obtain the following linear least squares problem:

$$E = E_s + \lambda E_d, \tag{6}$$

where λ is weight with a default value 10. We solve this energy minimization problem using a sparse linear solver.

3.2.1 Pre-warping

This approach can generate good results for simple stereoscopic image warping. When the warping becomes complex and the feature points are not evenly distributed, the spatially-varying warping applied to the right image is often not consistent with the user-specified warping applied to the left image. We solve this problem by pre-warping the right image with the same user-specified warping applied to the left image. As the left and right image in a stereoscopic image are similar, this pre-warping result gives a good approximation. We then apply the spatially-varying warping upon the pre-warping result and enforce the target disparities. Pre-warping is critical for our method to handle complex warping. It effectively reduces the warping of the right image into the warping of a pre-warped image according to a disparity map, which is much easier than directly warping the right image to follow the complex user-specified warping. In addition, as the required amount of warping is small after pre-warping, the fold-over artifacts rarely appear. In implementation, the right image is only re-sampled once to create the final result using texturing mapping after the final mesh is estimated. Figure 4 shows an example where pre-warping significantly improves the stereoscopic image warping result.

4 Results

This section shows a variety of stereoscopic image warping examples. All the stereoscopic images are rendered in red-cyan anaglyph and better viewed electronically. The disparity can be perceived by estimating the offset between the red and cyan channels. We also provide the left and right view of each image in our project website[4] for viewing on 3D displays to better assess our results.

4.1 Parametric warping

Similarity transformation. A similarity transformation allows translation, rotation, and uniform scaling. For uniform scaling, our method uniformly scales the disparity map and produces the same result as simultaneously scaling the left and right image of an input stereoscopic image. For rotation, the target disparity map from our method is the same as the input disparity map. We can consider that the scene rotates around the optical axis of the left camera. Since the scene depth distribution does not change, the (horizontal) disparity map remains the same (if the input image is taken by a rectified stereo camera). Figure 6 shows an example of rotating an input image with 30, 60, and 90 degrees. For a general similarity transformation, our method uniformly scales the disparity map with the scaling coefficient in the similarity transformation. We show a similarity transformation result in Figure 5 (a).

Affine transformation. As the local image warping introduced by an affine transformation is location-invariant, the disparity range scaling coefficient computed by Equation 3 is the same across the whole image. Figure 5 (b) shows an affine transformation example.

Perspective transformation. The local image warping introduced by a perspective transformation is location-dependent. Figure 5 (c)

[4] http://graphics.cs.pdx.edu/project/stereo-warp

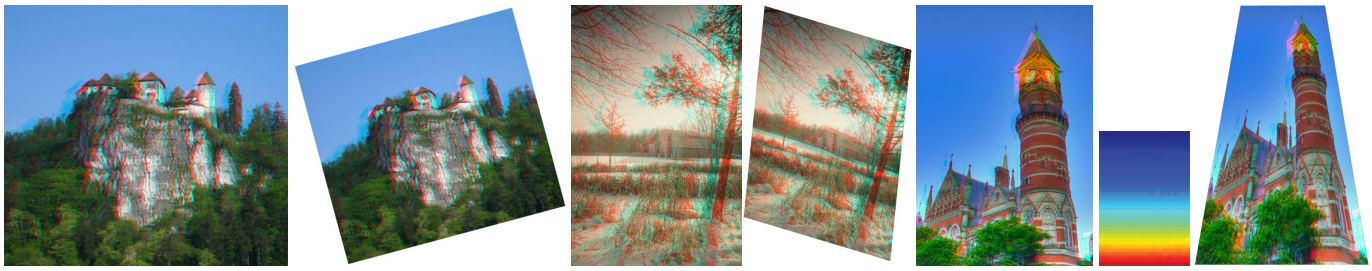

(a) Similarity transformation (b) Affine transformation (c) Perspective transformation (middle: scaling map)

Figure 5: *Parametric stereoscopic image warping. In each example, the left is the input and the right is the output.*

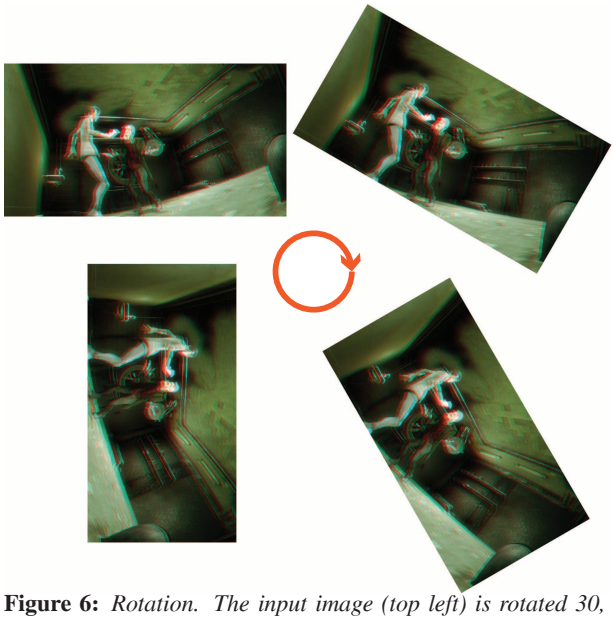

Figure 6: *Rotation. The input image (top left) is rotated 30, 60, and 90 degrees.*

shows a scaling coefficient map of a perspective transformation. This transformation significantly increases the perspective effect of the scene by warping the image and increasing the depth range for the image content near the viewer.

We show more complex parametric warping examples in Figure 7.

4.2 Non-parametric warping

Content-aware image retargeting. Content-aware image retargeting changes the image size while preserving the important image content [Wang et al. 2008b; Krähenbühl et al. 2009]. Figure 8 (b) shows such a retargeting result where the whole image is squished while the man is free from anisotropic distortion. The scaling map from our method captures this non-uniform warping well, as shown in (c). Our technique enables these retargeting techniques to be applied to stereoscopic images. Our technique warps the right image as shown in (d) and produces a retargetted stereoscopic image shown in (e).

Object resizing. Mesh-based image warping methods, such as feature-aware texturing [Gal et al. 2006], allow objects to be resized differently in an image. Our method extends these methods to stereoscopic images. Figure 8 (f) shows an example where the size of the dog is decreased while the size of the flying disk is increased.

Image shape manipulation. Mesh-based image warping methods can also be used to change the image shape while preserving important image content. Figure 8 (g) shows such an example where the rectangular image is transformed into a fan while the important objects, goats, are not distorted. Our method enables image shape manipulation on stereoscopic images.

4.3 User study

We performed user studies to further evaluate our results. In the user studies, stereoscopic images were displayed on an ASUS VG236H 120 Hz 3D monitor with shuttered glasses. 11 users with normal stereopsis perception participated in the studies and they have no knowledge of the methods used to create the results. We did not record their identities and took notes on their comments on the images for analysis. The first study was to verify whether viewers can easily perceive the roundness distortion illustrated in Figure 3. We tested on four images. For each image, we scale an input stereoscopic image and create two results using the input disparity map and the disparity map estimated using our method described in Section 3.1, respectively. Two images are scaled to half of the original size and the other two are scaled to double of the original size. In each trial, we put the input image in the middle, the two scaled results to the left and right of the input image. The placement of the images is randomized. We asked the participants "*which image has an object whose shape is more similar to the original one, left or right?*". On average participants correctly identified the roundness distortion at a rate of 84.09%.

The second study directly evaluated our warping results. This study included 30 results that cover all the warping cases discussed in Section 4.1 and 4.2. To acquaint the participant with the intended warping, we first showed the left image and its regular 2D warping result side by side. We then showed the input stereoscopic image and our warping result side by side. We asked each participant to rate the viewing experience from 1 to 3, with 3 being the best, according to the following two questions.

1. Can you comfortably obtain the stereopsis perception?
2. Do you like the warping result?

On average, participants rated our results with 2.63/3.0 in the comfort test, which shows that participants can comfortably view our warping results and obtain the stereopsis perception. Participants rated our results 2.65/3.0 in the second test.

4.4 Discussions

Our method produces physically plausible disparity maps for some special cases, such as uniform scaling and rotation. For more general warping, the disparity map in a warped stereoscopic image is often not geometrically valid. This is similar to the warping of image shape, which is often not physically plausible. Our experiments, however, show that viewers can still have robust depth perception. This is consistent with research in stereopsis perception as well as practice in stereoscopic photography and cinematography. An important reason is that other depth cues, such as perspective, relative object size, and occlusion, contribute to the human depth

Figure 7: *More parametric stereoscopic image warping. In each example, the left is the input and the right is the output.*

perception [Lang et al. 2010; Mendiburu 2009; Cutting and Vishton 1995].

One problem with our results is that the boundary of the right image is often not consistent with the left image. We need to slightly crop off some boundaries to create a final stereoscopic image. This is generally a difficult problem to solve especially when an aggressive warping is applied by users. Enforcing the boundary constraint often conflicts with the disparity constraints.

Our experiments also reveal that when a complex warping is applied and a big region in an image does not have feature points, the right image cannot be consistently warped as the left image. Our pre-warping scheme can significantly improve the result; however, occasionally, our result is still not satisfactory. Our method also depends on the feature matching accuracy. A better disparity map will help solve this problem.

Our method tries to match the perceived 3D object size change with the size change in 2D. Our experiments show that our method produces disparity maps that generally lead to comfortable 3D viewing experience. Meanwhile, we emphasize that disparity composition, like image warping, often involves artistic decisions. Experienced users can start from our disparity map and deliberately edit it with the support of disparity mapping tools, such as [Lang et al. 2010].

Finally, we pre-process a stereoscopic image to estimate the sparse disparity map. The cost of transforming the left image using a user-specified warping is independent of our method. The cost for target disparity map estimation depends on the number of feature points. For an image with around 2000 feature points, our method achieves around 60 fps on a desktop machine with Intel Duo Core 3.06GHz CPU and 4GB memory. The warping cost for the right image depends on the mesh size. For a mesh with size 64×36, our method achieves 10 fps.

5 Conclusion

In this paper, we present a technique for extending existing image warping algorithms to stereoscopic images. Our technique consists of three steps. We first warp the left image using the user-specified warping, then compute an optimal target disparity map, and finally warp the right image using a spatially-varying mesh-based warping technique guided by the target disparity map. Our experiments demonstrate a variety of popular image warping methods enabled by our technique. In the future, we will extend our method to stereoscopic video warping.

Acknowledgements

We would like to thank Rob Crockett and Flickr users, including -ytf-, turbguy, Dan(aka firrs), jaysdesk, tanj3d and fossilmike, for letting us use their photos under a Creative Commons license or with their permissions. The “Elephants Dream” frames are from Blender Foundation / Netherlands Media Art Institute / www.elephantsdream.org, used under a Creative Commons license and the stereo version is used from Youtube user geekboydischead under a Creative Commons license. This work was supported by NSF CNS-1205746 and the Portland State University Faculty Enhancement Grant.

References

BASHA, T., MOSES, Y., AND AVIDAN, S. 2011. Geometrically consistent stereo seam carving. In *IEEE International Conference on Computer Vision*.

CHANG, C.-H., LIANG, C.-K., AND CHUANG, Y.-Y. 2011. Content-aware display adaptation and interactive editing for stereoscopic images. *IEEE Trans. on Multimedia 13*, 589–601.

CUTTING, J. E., AND VISHTON, P. M. 1995. Perceiving layout and knowing distances: the integration, relative potency and contextual use of different information about depth. In *Handbook of perception and Cognition.*, W. Epstein and S. Rogers, Eds., vol. 5: Perception of Space and Motion. 69–117.

DIDYK, P., RITSCHEL, T., EISEMANN, E., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2011. A perceptual model for disparity. *ACM Trans. Graph. 30*, 4, 96:1–96:10.

GAL, R., SORKINE, O., AND COHEN-OR, D. 2006. Feature-aware texturing. In *Proceedings of EUROGRAPHICS Symposium on Rendering*, 297–303.

GOMES, J., DARSA, L., COSTA, B., AND VELHO, L. 1999. *Warping and Morphing of Graphical Objects*. Morgan Kaufmann Publishers.

GUTTMANN, M., WOLF, L., AND COHEN-OR, D. 2009. Semi-automatic stereo extraction from video footage. In *IEEE International Conference on Computer Vision*, 136–142.

HAREL, J., KOCH, C., AND PERONA, P. 2007. Graph-based visual saliency. In *NIPS*, vol. 19, 545–552.

HARTLEY, R. I., AND ZISSERMAN, A. 2004. *Multiple View Geometry in Computer Vision*. Cambridge University Press.

HECKBERT, P. S. 1989. Fundamentals of texture mapping and image warping. Tech. rep., UC Berkeley.

HEINZLE, S., GREISEN, P., GALLUP, D., CHEN, C., SANER, D., SMOLIC, A., BURG, A., MATUSIK, W., AND GROSS, M. 2011. Computational stereo camera system with programmable control loop. *ACM Trans. Graph. 30*, 94:1–94:10.

HIRSCHMULLER, H., AND SCHARSTEIN, D. 2009. Evaluation of stereo matching costs on images with radiometric differences. *IEEE Trans. Pattern Anal. Mach. Intell. 31*, 1582–1599.

HOFFMAN, D., GIRSHICK, A., AKELEY, K., AND BANKS, M. 2008. Vergence–accommodation conflicts hinder visual performance and cause visual fatigue. *Journal of Vision 8*, 3, 1–30.

HOWARD, I. P., AND ROGERS, B. J. 2002. *Seeing in Depth*. Oxford University Press.

JONES, G., LEE, D., HOLLIMAN, N., AND EZRA, D. 2001. Controlling perceived depth in stereoscopic images. In *Stereoscopic Displays and Virtual Reality Systems VIII*, 42–53.

KIM, C., HORNUNG, A., HEINZLE, S., MATUSIK, W., AND GROSS, M. 2011. Multi-perspective stereoscopy from light fields. *ACM Trans. Graph. 30*, 6, 190:1–190:10.

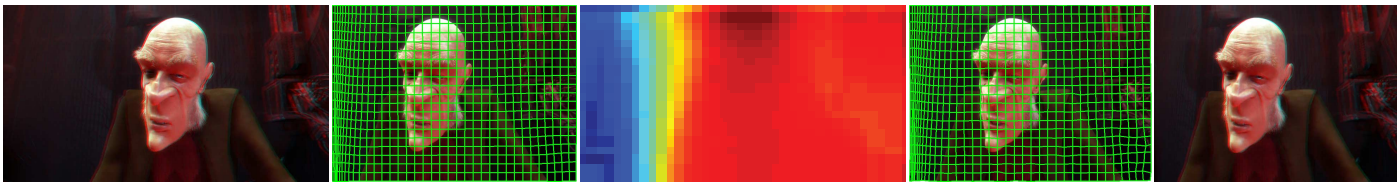

(a) Input (b) Left result (c) Scaling map (d) Right result (e) Retargeting result

(f) Object resizing (g) Content-preserving shape manipulation

Figure 8: *Non-parametric stereoscopic image warping. The first row shows stereoscopic image retargetting. The second row shows object resizing and content-preserving image shape manipulation.*

KNORR, S., AND SIKORA, T. 2007. An image-based rendering (ibr) approach for realistic stereo view synthesis of tv broadcast based on structure from motion. In *IEEE International Conference on Image Processing*, 572 – 575.

KOPPAL, S., ZITNICK, C., COHEN, M., KANG, S. B., RESSLER, B., AND COLBURN, A. 2011. A viewer-centric editor for 3d movies. *IEEE Computer Graphics and Applications 31*, 20 –35.

KRÄHENBÜHL, P., LANG, M., HORNUNG, A., AND GROSS, M. 2009. A system for retargeting of streaming video. *ACM Trans. Graph. 28*, 5, 126:1–126:10.

LANG, M., HORNUNG, A., WANG, O., POULAKOS, S., SMOLIC, A., AND GROSS, M. 2010. Nonlinear disparity mapping for stereoscopic 3d. *ACM Transactions on Graphics 29*, 4.

LEE, K.-Y., CHUNG, C.-D., AND CHUANG, Y.-Y. 2012. Scene warping: Layer-based stereoscopic image resizing. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 49 – 56.

LIU, F., GLEICHER, M., JIN, H., AND AGARWALA, A. 2009. Content-preserving warps for 3d video stabilization. *ACM Transactions on Graphics 28*, 3, 44.

LIU, C.-W., HUANG, T.-H., CHANG, M.-H., LEE, K.-Y., LIANG, C.-K., AND CHUANG, Y.-Y. 2011. 3d cinematography principles and their applications to stereoscopic media processing. In *ACM international conference on Multimedia*, 253–262.

LO, W.-Y., VAN BAAR, J., KNAUS, C., ZWICKER, M., AND GROSS, M. 2010. Stereoscopic 3d copy & paste. *ACM Transactions on Graphics 29*, 6, 147:1–147:10.

LOWE, D. G. 2004. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision 60*, 2, 91–110.

MENDIBURU, B. 2009. *3D Movie Making: Stereoscopic Digital Cinema from Script to Screen*. Focal Press.

MOUSTAKAS, K., TZOVARAS, D., AND STRINTZIS, M. 2008. Stereoscopic video generation based on efficient layered structure and motion estimation from a monoscopic image sequence. *IEEE Trans. Circuits Syst. Video Technol. 15*, 8, 1065 – 1073.

MUELLER, R., WARD, C., AND HUŠÁK, M. 2008. A systematized wysiwyg pipeline for digital stereoscopic 3d filmmaking. In *Proceedings of SPIE*, vol. 6803.

OSKAM, T., HORNUNG, A., BOWLES, H., MITCHELL, K., AND GROSS, M. 2011. Oscam - optimized stereoscopic camera control for interactive 3d. *ACM Trans. Graph. 30*, 6, 189:1–189:8.

SAXENA, A., SUN, M., AND NG, A. 2009. Make3d: Learning 3d scene structure from a single still image. *IEEE Transactions on Pattern Analysis and Machine Intelligence 31*, 5, 824 – 840.

SCHARSTEIN, D., AND SZELISKI, R. 2002. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Int. J. Comput. Vision 47* (April), 7–42.

SMOLIC, A., POULAKOS, S., HEINZLE, S., GREISEN, P., LANG, M., HORNUNG, A., FARRE, M., STEFANOSKI, N., WANG, O., SCHNYDER, L., MONROY, R., AND GROSS, M. 2011. Disparity-aware stereo 3d production tools. In *European Conference on Visual Media Production*.

SUN, D., ROTH, S., AND BLACK, M. J. 2010. Secrets of optical flow estimation and their principles. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2432–2439.

SZELISKI, R. 2010. *Computer Vision: Algorithms and Applications*. Springer.

WANG, C., AND SAWCHUK, A. A. 2008. Disparity manipulation for stereo images and video. In *Proc. SPIE, Vol. 6803*.

WANG, L., JIN, H., YANG, R., AND GONG, M. 2008. Stereoscopic inpainting: Joint color and depth completion from stereo images. In *IEEE Conf. on Comput. Vision and Pattern Recogn.*

WANG, Y.-S., TAI, C.-L., SORKINE, O., AND LEE, T.-Y. 2008. Optimized scale-and-stretch for image resizing. *ACM Transactions on Graphics 27*, 5.

WOLBERG, G. 1990. *Digital image warping*. Wiley-IEEE Computer Society Press.

ZILLY, F., MULLER, M., EISERT, P., AND KAUFF, P. 2010. The stereoscopic analyzer-an image-based assistance tool for stereo shooting and 3d production. In *IEEE International Conference on Image Processing*, 4029–4032.