

Real-Time Variational Fisheye Stereo without Rectification and Undistortion

Menandro Roxas¹ and Takeshi Oishi²

Abstract—Dense 3D maps from wide-angle cameras is beneficial to robotics applications such as navigation and autonomous driving. In this work, we propose a real-time dense 3D mapping method for fisheye cameras without explicit rectification and undistortion. We extend the conventional variational stereo method by constraining the correspondence search along the epipolar curve using a trajectory field induced by camera motion. We also propose a fast way of generating the trajectory field without increasing the processing time compared to conventional rectified methods. With our implementation, we were able to achieve real-time processing using modern GPUs. Our results show the advantages of our non-rectified dense mapping approach compared to rectified variational methods and non-rectified discrete stereo matching methods.

I. INTRODUCTION

Wide-angle (fisheye) cameras have seen significant usage in robotics applications. Because of the wider field-of-view (FOV) compared to pinhole camera model, fisheye cameras pack more information in the same sensor area which are advantageous especially for object detection, visual odometry, and 3D reconstruction.

Real-time dense 3D mapping using fisheye cameras have several advantages especially in navigation and autonomous driving. For example, the wide field-of-view allows simultaneous visualization and observation of objects in different directions.

Several methods have addressed the 3D mapping problem for fisheye cameras. The most common approach performs rectification of the images to perspective projection which essentially removes the main advantage of such cameras - wide FOV. Moreover, information closer to the edge of the image are highly distorted while objects close the center are highly compressed, not to mention adding unnecessary degradation of image quality due to spatial sampling. Other rectification that retains the fisheye’s wide FOV involves reprojecting on a sphere, which suffers from similar degradation especially around the poles. We address these issues by directly processing the distorted images without rectification or undistortion.

We embed our method in a variational framework, which inherently produces smooth dense maps in contrast to discrete stereo matching methods. We propose to use a trajectory field that constrains the search space of corresponding pixels along the epipolar curve. We also propose a fast way of generating the trajectory field that does not require additional processing time compared to conventional variational methods.

Both authors are with The University of Tokyo, Tokyo, Japan ¹roxas,
²oishi@cvl.iis.u-tokyo.ac.jp




Fig. 1. Non-rectified variational stereo method result on a fisheye stereo camera.

The advantage of our proposed method is two-folds. First, without rectification or undistortion, the sensor level image quality is preserved. Second, our method can handle arbitrary camera distortions. While the results in this paper focuses on fisheye cameras, applying our method on other camera models is straightforward.

Our results show additional accurate measurements when compared to conventional rectified methods, and more accurate and dense estimation compared to non-rectified discrete methods. Finally, with our implementation, we were able to achieve real-time processing on a consumer fisheye stereo camera system and modern GPUs.

II. RELATED WORK

Dense stereo estimation in perspective projection consists of a one-dimensional correspondence search along the epipolar lines. In a variational framework, the search is akin to linearizing the brightness constancy constraint along the epipolar lines. In [1], a differential vector field induced by arbitrary camera motion was used for linearization. However, their method, as with other variational stereo methods in perspective projection such as [2], requires undistortion and/or rectification (in case of binocular stereo) to be applicable for fisheye cameras [3].

Instead of perspective rectification, some method reprojects the images to spherical or equirectangular projection [4] [5] [6] [7]. However, this approach suffers greatly from highly distorted images along the poles which makes estimation less accurate especially when using the variational framework. Similar to our brightness constancy linearization approach, [6] generates differential vectors induced by variations on a 2-sphere in which the variational stereo method was applied. However, their graph-based formulation is a solution to the self-induced problem arising from reprojecting the image on a spherical surface. In contrast, our method does

not require reprojection on a 2-sphere and works directly on the distorted images without undistortion, reprojection or rectification. We do this by evaluating the variations directly from the epipolar curve.

Other methods also directly work on the distorted fisheye images. In [8], the unified camera model [9] was used to determine the path of the search space, which are incrementally shifted (akin to differential vectors) from a reference pixel to the maximum disparity. At each point, the projection function is re-evaluated which the authors claim was costly compared to linear search. Their mapping method, while real-time, only produces semi-dense depth maps. In [10], a similar parameterization of the epipolar curve was done, but only applied on window-based stereo matching. Other methods adapt linear matching algorithms to omni-directional cameras such as semi-global matching [11], plane-sweeping [12] and a variant called sphere-sweeping [13]. Sparse methods were also adapted to handle fisheye distortion such as [14] among others.

III. VARIATIONAL FISHEYE STEREO

In this section, we will first introduce the problem of image linearization for fisheye cameras in Sec. III-A. We will then propose our trajectory field and warping techniques in Secs. III-B and III-C. Finally, we will combine our approach with a variational optimization method and summarize the algorithm in Sec. III-D.

A. Image Constancy Constraint

Classical variational stereo method consists of solving a dense disparity map between a pair of images that minimizes a convex energy function. Given two images, I_0 and I_1 , with known camera transformation and intrinsic parameters, the one-dimensional disparity u can be solved by minimizing:

$$E(u) = E_{data}(u) + E_{smooth}(u) \quad (1)$$

The above functional consists of a data term and a smoothness term. Building upon perspective camera stereo methods [1][2], we only need to modify the formulation of the data term in order to accommodate the distortion effects in fisheye cameras.

In general, the data term penalizes the residual, ρ , which measures the constancy between corresponding pixels in I_0 and I_1 . I can be any value such as brightness, intensity gradient, non-local transforms [15], etc. For fisheye cameras, these correspondences are constrained along the epipolar curve, $\gamma : \mathbb{R} \rightarrow \mathbb{R}^2$ and finding them constitutes a one-dimensional search [8][11][12] along γ . In our case, we will solve the correspondences using a variational framework.

Let \mathbf{x} be a point on the epipolar curve. The corresponding point at a distance u along the curve can be expressed as $\pi(\exp(\hat{\xi}_1) \cdot \mathbf{X}(\mathbf{x}, u))$, where $\pi : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ is the projection of the 3D point \mathbf{X} on the image plane Ω_1 of I_1 , and $\hat{\xi}_1$ is the camera pose of I_1 relative to I_0 . We can then express the residual as:

$$\rho(\mathbf{x}, u) = I_1 \left(\pi \left(\exp(\hat{\xi}_1) \cdot \mathbf{X}(\mathbf{x}, u) \right) \right) - I_0(\mathbf{x}) \quad (2)$$

Assuming that I_0 and I_1 is linear along the curve, we can approximate (2) using the first-order Taylor expansion, and using a simplified notation $I_1(\mathbf{x}, u)$ as in [1]

$$\bar{\rho}(\mathbf{x}, u) = I_1(\mathbf{x}, u_\omega) + (u - u_\omega) \frac{d}{du} I_1(\mathbf{x}, u) \Big|_{u_\omega} - I_0(\mathbf{x}) \quad (3)$$

So far, our formulation of the data term still follows that of [1].

Formally, the derivative $\frac{d}{du} I_1(\mathbf{x}, u)$ can be expressed as the dot product of the gradient of $I_1(\mathbf{x}, u)$ and a differential vector at \mathbf{x} :

$$\frac{d}{du} I_1(\mathbf{x}, u) = \nabla I_1(\mathbf{x}, u) \cdot \underbrace{\frac{d}{du} \pi \left(\exp(\hat{\xi}_1) \cdot \mathbf{X}(\mathbf{x}, u) \right)}_{\text{differential vector}} \quad (4)$$

However, in practice, we directly solve for the variations of I along the epipolar curve. In discrete form, we have:

$$\frac{d}{du} I_1(\mathbf{x}, u) = I_1(\mathbf{x} + \gamma') - I_1(\mathbf{x}) \quad (5)$$

where γ' is the differential vector. (Note that Eqs. (4) and (5) result in the pole-stretching problem when using spherical projections [6]).

Minimizing Eq. (3) results in the incremental disparity $(u - u_\omega)$ which we will designate from here on as δu_ω . For small δu_ω , the differential vector in (4) is equal the tangential vector of the epipolar curve $\gamma' = \nabla \gamma$.

Moreover, since the linearity assumption of I is only valid for a small disparity, (3) is usually embedded in an iterative warping framework [16] around a known disparity u_ω (hence, the term $I_1(\mathbf{x}, u_\omega)$). That is, for every warping iteration ω , we update $u_{\omega+1} = u_\omega + \delta u_\omega$.

This formulation raises two issues when used in a fisheye camera system.

- First, the warping technique requires a re-evaluation of γ at every iteration to find the tangential vectors at u_ω which is tedious and time consuming.
- Second, even if we assume that the image is perfectly linear along the epipolar curve, ∇I is only valid along the direction of the tangential vectors. In a perspective projection, this is not a problem since the tangential vectors indicates the exact direction of the epipolar lines. In our case, the gradient will need to be evaluated exactly along the curve.

In the following sections, we will elaborate on our approach to solve these two issues.

B. Trajectory Field Representation for Epipolar Curves

To avoid the re-evaluation of the epipolar curve for every warping step, we generate a trajectory field image that represents the tangential vectors γ' at every pixel \mathbf{x} . As a result, γ' at the next iteration step can be simply solved using bicubic interpolation.

Moreover, instead of solving for the parameterized curve function for every pixel [17], we programmatically generate the trajectory field. We first assume a known transformation $\hat{\xi}_1$ between two camera positions with non-zero translation




Fig. 2. Calibration (left) and trajectory (right) field for a binocular fisheye stereo.

($|t| \neq 0$) and known projection π . Our method is not restricted on any type of camera model [9] [18] [19] as long as the projection function π is defined.

Using π , we project a surface of arbitrary depth onto the two cameras: $\mathbf{x}_0 = \pi(\mathbf{X})$, $\mathbf{x}_1 = \pi(\exp(\hat{\xi}_1) \cdot \mathbf{X})$. This gives us the exact correspondence $\mathbf{w}(\mathbf{x}_0, \mathbf{x}_1) = \mathbf{x}_1 - \mathbf{x}_0$. In a perspective projection, this mapping or the optical flow already represents the slope of the epipolar lines. Assuming pre-rotated images, i.e. $R = 0$, the direction of the optical flow, $\frac{\mathbf{w}}{|\mathbf{w}|}$, will be dependent only on the direction of the camera translation t and independent of $|t|$ and the surface depth $|\mathbf{X}|$. However, for fisheye projection, $\frac{\mathbf{w}}{|\mathbf{w}|}$ is still affected by the camera distortion.

To address this, we can represent the optical flow as the sum of the tangential vectors along the path of the epipolar curve between the two corresponding points. Let the parameterization variable for γ be $s = [0, 1]$. In continuous form, we can express $\mathbf{w}(\mathbf{x}_0, \mathbf{x}_1)$ as:

$$\mathbf{w}(\mathbf{x}_0, \mathbf{x}_1) = \int_0^c \gamma'(s) ds \Big|_{c=1} \quad (6)$$

By scaling the camera translation such that $|t| \rightarrow 0$, the left hand side of (6) approaches $\mathbf{0}$. It follows that the right hand side becomes:

$$\lim_{c \rightarrow 0} \int_0^c \gamma'(s) ds = \gamma'(0) \quad (7)$$

which finally allows us to approximate $\gamma'(0) \approx \frac{\mathbf{w}}{|\mathbf{w}|}$. In short, $\frac{\mathbf{w}}{|\mathbf{w}|}$ gives us the normalized trajectory field.

C. Warping Technique

The trajectory field discretizes the epipolar curve by assigning finite vector values for every pixel. We can think of this approach as decomposing the epipolar curve as a piecewise linear function (see Figure 3) which allows us to express the disparity u as:

$$u = \sum_{\omega=0}^N \delta u_{\omega} \quad (8)$$

where N is the total number of warping iterations.

Clearly, we can better approximate the epipolar curve by setting a magnitude limit to the incremental δu_{ω} and increasing the number of iterations N . Moreover, doing so




Fig. 3. Epipolar curve as a piecewise linear function. Large incremental δu_{ω} results in wrong tracked curve.

also prevents missing the correct trajectory of the curve since δu_{ω} is constrained along γ'_{ω} . (see Figure 3).

To solve the final warping vector of \mathbf{x} using the trajectory field, we use:

$$\mathbf{w} = \sum_{\omega=0}^N \delta u_{\omega} \gamma'_{\omega} \quad (9)$$

D. Anisotropic TGV-L1 Optimization

Before we continue, we will complete the energy functional in (1). We will follow the anisotropic tensor-guided total generalized variation (TGV) constraint described in [2] and combine it with data term (2) which results in the following:

$$E(u) = \lambda \int_{\Omega} |\rho(\mathbf{x}, u)| d^2 \mathbf{x} + \alpha_0 \int_{\Omega} |\nabla v| d^2 \mathbf{x} + \alpha_1 \int_{\Omega} |T^{\frac{1}{2}} \nabla u - v| d^2 \mathbf{x} \quad (10)$$

We can minimize (10) using primal-dual algorithm, which consists of a gradient-ascent on the dual variables $p : \mathbb{R}^2$ and $q : \mathbb{R}^4$, followed by a gradient-descent and over-relaxation refinement step on the primal variables u and $v : \mathbb{R}^2$. The algorithm is summarized as:

$$\begin{cases} p_{k+1} = \mathcal{P} \left(p_k + \sigma_p \alpha_1 (T^{\frac{1}{2}} \nabla \bar{u}_k - \bar{v}_k) \right) \\ q_{k+1} = \mathcal{P} (q_k + \sigma_q \alpha_0 (\nabla \bar{v}_k)) \\ u_{k+1} = (I + \tau_u \partial G)^{-1} (u_n + \tau_u \text{div} (T^{\frac{1}{2}} p_{k+1})) \\ v_{k+1} = v_k + \tau_v (\text{div} q_{k+1} + p_{k+1}) \\ \bar{u}_{k+1} = u_{k+1} + \theta (u_{k+1} - \bar{u}_k) \\ \bar{v}_{k+1} = v_{k+1} + \theta (v_{k+1} - \bar{v}_k) \end{cases} \quad (11)$$

where $\mathcal{P}(\phi) = \frac{\phi}{\max(1, \|\phi\|)}$ is a fixed-point projection operator. The step sizes $\tau_u > 0, \tau_v > 0, \sigma_u > 0, \sigma_v > 0$ are solved using a pre-conditioning scheme following [20] while the relaxation variable θ is updated for every iteration as in [21]. The tensor $T^{\frac{1}{2}}$ is calculated as:

$$T^{\frac{1}{2}} = \exp(-\beta |I_0|^{\eta}) n n^T + n^{\perp} n^{\perp T} \quad (12)$$

where $n = \frac{\nabla I_0}{|\nabla I_0|}$ and n^{\perp} is the vector normal to ∇I_0 , while β and η are scalars controlling the magnitude and sharpness of the tensor. This tensor guides the propagation of the disparity information among neighboring pixels, while considering the natural image boundaries as encoded in n and n^{\perp} .

The resolvent operator [21] $(I + \tau_u \partial G)^{-1}(\hat{u})$ is evaluated using the thresholding scheme:

$$(I + \tau_u \partial G)^{-1}(\hat{u}) = \hat{u} + \begin{cases} \tau_u \lambda I_u & \text{if } \bar{\rho} < -\tau_u \lambda I_u^2 \\ -\tau_u \lambda I_u & \text{if } \bar{\rho} > \tau_u \lambda I_u^2 \\ \bar{\rho}/I_u & \text{if } |\bar{\rho}| \leq \tau_u \lambda I_u^2 \end{cases}$$

where $I_u = \frac{d}{du} I_1(\mathbf{x}, u)$. We summarize our approach in Algorithm 1.

The solved disparity is converted to depth by triangulating the unprojection rays using the unprojection function π^{-1} . This step is specific for the camera model used, hence we will not elaborate on methods to address this. Nevertheless, some camera models have closed-form unprojection function [9] [18] while others require non-linear optimizations [19].

IV. IMPLEMENTATION

In this section, we discuss our implementation choices to achieve accurate results and real-time processing, which includes image pre-processing, large displacement handling and our selected optimization parameters and hardware considerations.

A. Pre-rotation and calibration

We perform a calibration and pre-rotation of the image pairs before running the stereo estimation. We create a calibration field in the same manner as the trajectory field. The calibration field contains the rotation information as well as the difference in camera intrinsic properties (for binocular stereo case).

Again, we project a surface of arbitrary depth on the two cameras with projection function π_0 and π_1 while setting the translation vector $t = \mathbf{0}$. We then solve for the optical flow $\mathbf{w} = \mathbf{x}_1 - \mathbf{x}_0$. In this case the optical flow exactly represents the calibration field (see Figure 2). In case where $\pi_0 \neq \pi_1$, such as in binocular stereo, the calibration field will also contain the difference in intrinsic properties. For example, a difference in image center results in the diagonal warping in our binocular camera system as seen in Figure 2. Using the calibration field, we warp the second image I_1 once, resulting in a translation only transformation.

Algorithm 1 Algorithm for anisotropic TGV-L1 stereo for fisheye cameras.

Require: $I_0, I_1, \hat{\xi}_1, \pi$

Generate trajectory field (Sec. III-B)

$\omega = 0, \mathbf{w}_\omega = 0, u_\omega = 0$

while $\omega < N$ **do**

 Warp I_1 using \mathbf{w}_ω

while $k < nIters$ **do**

 Update primal-dual variables (11)

end while

 Clip δu_ω (Sec. III-C)

$u_{\omega+1} = u_\omega + \delta u_\omega$

$\mathbf{w}_{\omega+1} = \mathbf{w}_\omega + \delta u_\omega \gamma'(\mathbf{x})$

end while




Fig. 4. Trade-off between accuracy and processing time for choosing the warping iteration (better viewed in color)

B. Coarse-to-Fine Approach

Similar to most variational framework, we employ a coarse-to-fine (pyramid) technique to handle large displacement. Starting from a coarser level of the pyramid, we run N warping iterations and upscale both the current disparity and the warping vectors and carry the values on to the finer level.

One caveat of this approach on a fisheye image is the boundary condition especially for gradient and divergence calculations. To address this, we employ the Neumann and Dirichlet boundary conditions applied on a circular mask that rejects pixels greater than the desired FOV. The mask is scaled accordingly using nearest-neighbor interpolation for every level of the pyramid. Moreover, by applying a mask, we also avoid the problem of texture interpolation with a zero-value during upscaling when the sample falls along the boundary of the fisheye image.

C. Timing Considerations

We implemented our method with C++/CUDA on an i7-4770 CPU and NVIDIA GTX 1080Ti GPU. For TGV-L1 optimization and primal-dual algorithm, we use the parameter values: $\beta = 9.0$, $\eta = 0.85$, $\alpha_0 = 17.0$ and $\alpha_1 = 1.2$. Moreover, we fix the iteration values based on the desired timing and input image size. For an 800x800 image, we found that the primal-dual iteration of 10 is sufficient for our application, with pyramid size = 5 and scaling = 2.0 (minimum image width = 50).

For the warping iteration, we plot the trade-off between accuracy and processing time in Figure 4 with fixed $\delta u^{max} = 0.2px$. From the plot, we can see that the timing linearly increases with the number of iterations, but the accuracy exponentially decreases. Choosing a proper value for N needs careful considerations according to the application.

V. RESULTS

We present our results in the following sections. First, we show the effect of limiting the magnitude of the incremental disparity solution per warping iteration to the accuracy of the estimation. Then, we compare our method with an existing rectified variational stereo method and a discrete stereo




Fig. 5. Limiting the magnitude of δu per iteration reduces the error around sharp image gradients and occlusion boundaries.

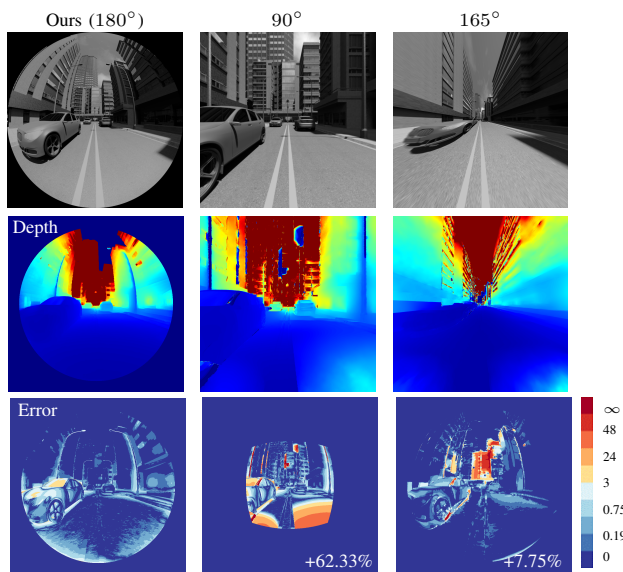


Fig. 7. Comparison between [2] with different field-of-view (90° and 165°) and our method. We compare the disparity error [22] as well as percentage of accuracy improvement by using our method.

matching method using both synthetic and real datasets with ground truth depth. Finally, we show some sample qualitative results on a commercial-off-the-shelf stereo camera fisheye system.

A. Limiting Incremental Disparity

To test the effect of limiting the incremental disparity, we measure the accuracy of our method on varying warping iteration and disparity limits. In Figure 5, we show the photometric error (absolute difference between I_0 and warped I_1) when $\delta u^{max} = 1.0px$ and $\delta u^{max} = 0.2px$. From the images, we can see that the photometric error is larger in areas with significant information (e.g. intensity edges and occlusion boundaries) when $\delta u^{max} = 1.0px$ compared to $\delta u^{max} = 0.2px$. This happens because it is faster for the optimization to converge in highly textured surfaces which results in overshooting from the tracked epipolar curve, as




Fig. 6. Accuracy of disparity (percentage of erroneous pixels, τ) with limiting the magnitude of δu for different warping iteration values N .

shown in Figure 3.

However, limiting the magnitude of δu has an obvious drawback. If the warping iteration is not sufficient, the estimated δu will not reach to its correct value which will result in higher error. We show this effect in Figure 6. Here, we plot various warping iterations N and show the accuracy of estimation with increasing δu^{max} using percentage of erroneous pixel measure ($\tau > 1$) [22]. Clearly, higher N and smaller δu^{max} results in a more accurate estimation.

B. Comparison with Rectified Method

We first compare our proposed approach with a rectified stereo method. To achieve a fair comparison, we use the same energy function and parameters in our implementation, except that we apply them in a rectified image. This rectified stereo approach is similar to the method presented in [2], except that we use intensity values instead of the census transform. We also explicitly applied a time-step preconditioning step and a relaxation after every iteration.

We compare our method with varying FOV for [2] on a synthetic dataset [23]. We use the same erroneous pixel measure from the previous section and summarize the result in Table I using an FOV of 120° for [2]. We also compare the disparity error [22] as well as the improvement additional accurate pixels (see Figure 7) using the full 180° for our method and a FOV of 90° and 165° for [2].

To better visualize the comparison, we transform the rectified error back to the original fisheye form. From the results, extreme compression around the center with ultra-wide angle (165°) rectification results in higher error especially for distant objects. With larger image area coverage, our approach do not suffer from this compression problem and maintains uniform accuracy throughout the image. Moreover, with the lower compression around the center (90°), the rectified method have increased error around the edges for closer objects (ground) due to increased displacement.

Additionally, we found no significant difference in processing time because the warping techniques are both run

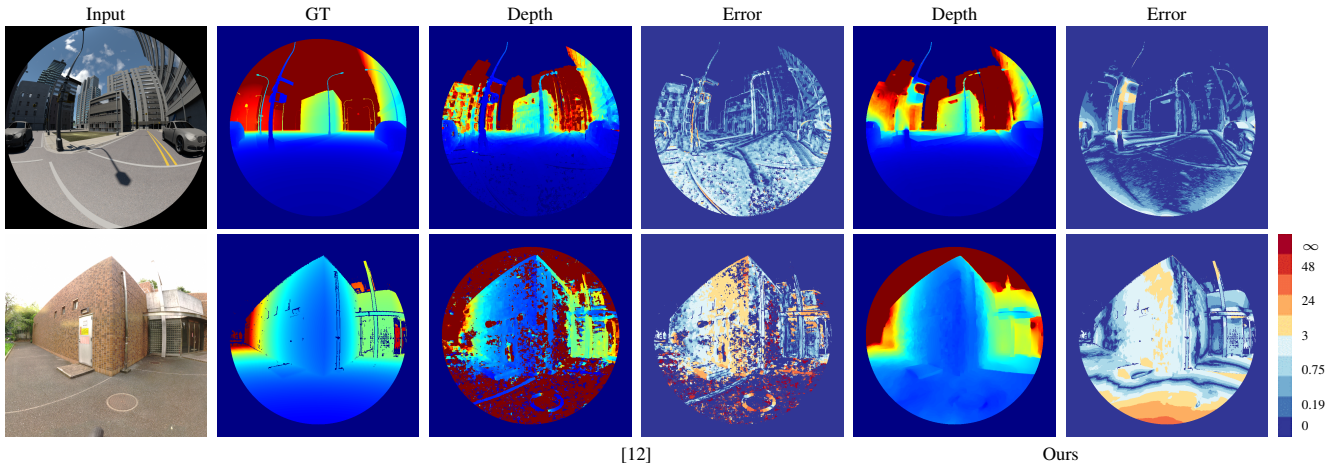


Fig. 8. Sample results on real and synthetic data with [12] and our method with disparity error [22].

in a single GPU kernel call and consumes the same texture memory access latency.

C. Comparison with Non-Rectified Method

In this section, we compare our method with planesweep implemented on a fisheye camera system [12] on real and synthetic scenes. The images were captured from two arbitrary camera location with non-zero translation. We show the sample results in Figure 8 and Table I.

One of the advantages of variational methods is the inherent density of the estimated depth when compared to discrete matching methods. In our experiments, we found that while our method is denser and significantly smoother than [12], it is more prone to miss very thin objects such as poles. Moreover, because our method is built upon a pyramid scheme, very large displacements are difficult to estimate which is visible in the results when the object is very close to the camera (nearest ground area).

Nevertheless, we show in Table I that our method is overall more accurate compared to [12] even after we removed the ambiguous pixels due to occlusion and left-right inconsistency. (In Table I, we use only the valid pixels in [12] for comparison).

D. Real-World Test

We tested our method on a laptop computer with NVIDIA GTX1060 GPU and an Intel RealSense T265 stereo camera, which has a $163 \pm 5^\circ$ FOV, global-shutter 848x800 grayscale image and a 30fps throughput. We show the sample results in Figure 1. We were able to achieve a 10fps with 5 warping iterations on a full image, and 30fps with 20 warping iterations on a half-size image. This system can be easily mounted on medium sized rover for SLAM applications.

VI. CONCLUSION AND FUTURE WORK

In this paper, we presented a warping technique for handling fisheye cameras specified for real-time variational stereo estimation methods without explicit image rectification. From our results, we showed that our approach can achieve higher and more uniform accuracy and larger FOV

TABLE I
DISPARITY ERROR COMPARISON WITH RECTIFIED, NON-RECTIFIED,
AND OUR METHOD.

	Frame	[2]		[12]		Ours	
		$\tau > 1$	$\tau > 3$	$\tau > 1$	$\tau > 3$	$\tau > 1$	$\tau > 3$
synthetic	04	14.72	8.06	28.46	3.40	6.60	0.78
	05	17.19	10.06	26.13	2.05	7.66	0.97
	06	11.76	8.33	27.76	1.55	5.51	1.51
	07	11.21	2.46	27.46	1.75	4.55	0.28
real	92	-	-	-	50.64	-	34.79
	100	-	-	-	43.85	-	20.15

compared to conventional methods without increasing the processing time.

Because of the wider FOV of fisheye cameras, the disadvantage of most variational methods, which is handling large displacement (wide baseline or near objects), is highlighted. However, this can be overcome by using large displacement techniques or initialization with discrete methods (such as planesweep).

REFERENCES

- [1] J. Stühmer, S. Gumhold, and D. Cremers, "Real-time dense geometry from a handheld camera," *Pattern Recognition. DAGM 2010. LNCS*, vol. 6376, 2010.
- [2] R. Ranftl, S. Gehrig, T. Pock, and H. Bischof, "Pushing the limits of stereo using variational stereo estimation," in *Proc. IEEE Intel. Vehic.*, June 2012.
- [3] J. Schneider, C. Stachniss, and W. Forstner, "On the accuracy of dense fisheye stereo," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, 2016.
- [4] Z. Arican and P. Frossard, "Dense disparity estimation from omnidirectional images," in *IEEE Conf. Adv. Vid. Sig. Based Surv.*, September 2007.
- [5] M. Schonbein and A. Geiger, "Omnidirectional 3d reconstruction in augmented manhattan worlds," in *Proc. IEEE Int. Work. Robots Sys.*, 2014.
- [6] L. Bagnato, P. Frossard, and P. Vanderghenst, "A variational framework for structure from motion in omnidirectional image sequences," *J. Math Imaging Vis.*, vol. 41, pp. 182–193, 2011.
- [7] W. Gao and S. Shen, "Dual-fisheye omnidirectional stereo," in *Proc. IEEE Int. Work. Robots Sys.*, 2017.

- [8] D. Caruso, J. Engel, and D. Cremers, "Large-scale direct slam for omnidirectional cameras," in *Proc. IEEE Int. Work. Robots Sys.*, 2015.
- [9] C. Geyer and K. Daniilidis, "A unifying theory for central panoramic systems and practical applications," in *Proc. IEEE Europ. Conf. Comput. Vis.*, July 2000, pp. 445–461.
- [10] R. Bunschoten and B. Krose, "Robust scene reconstruction from an omnidirectional vision system," *IEEE Trans. Robot. Automat.*, 2003.
- [11] B. Khomutenko, G. Garcia, and P. Martinet, "Direct fisheye stereo correspondence using enhanced unified camera model and semi-global matching algorithm," in *ICARCV*, 2016.
- [12] C. Hane, L. Heng, G. H. Lee, A. Sizov, and M. Pollefeys, "Real-time direct dense matching on fisheye images using plane-sweeping stereo," in *Proc. Int. Conf. 3D Vis.*, December 2014.
- [13] S. Im, H. Ha, F. Rameau, H.-G. Jeon, G. Choe, and I. S. Kweon, "All-around depth from small motion with a spherical panoramic camera," in *Proc. IEEE Europ. Conf. Comput. Vis.*, 2016, pp. 156–172.
- [14] B. Micusik and T. Pajdla, "Structure from motion with wide circular field-of-view cameras," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 28, no. 7, pp. 1135–1149, July 2006.
- [15] R. Zabih and J. Li, "Non-parametric local transforms for computing visual correspondence," in *Proc. IEEE Europ. Conf. Comput. Vis.*, 1994, pp. 151–158.
- [16] N. Papenberg, A. Bruhn, T. Brox, S. Didas, and J. Weickert, "Highly accurate optical flow computation with theoretically justified warping," *Int. J. Comput. Vis.*, vol. 67, pp. 141–158, 2006.
- [17] T. Svoboda, T. Pajdla, and V. Hlavac, "Epipolar geometry for panoramic cameras," in *Proc. IEEE Europ. Conf. Comput. Vis.*, 1998, pp. 218–231.
- [18] B. Khomutenko, G. Garcia, and P. Martinet, "An enhanced unified camera model," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 137–144, January 2016.
- [19] J. Kannala and S. S. Brandt, "A generic camera model and calibration method for conventional, wide-angle, fisheye lenses," *Proc. IEEE Int. Conf. Comput. Vis.*, vol. 28, no. 8, pp. 1335–1340, September 2006.
- [20] T. Pock and A. Chambolle, "Diagonal pre-conditioning for first order primal-dual algorithms in convex optimization," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2011.
- [21] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *Journal of Mathematical Imaging and Vision*, vol. 40, no. 1, pp. 120–145, May 2011.
- [22] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *Int. J. Robot. Res.*, 2013.
- [23] Z. Zhang, H. Rebecq, C. Forster, and D. Scaramuzza, "Benefit of large field-of-view cameras for visual odometry," in *Proc. IEEE Int. Conf. Robot Automat.*, 2016.