

Multimodal Stereo Vision For Reconstruction In The Presence Of Reflection

Scott Sorensen
sorensen@udel.edu

Philip Saponaro
saponaro@udel.edu

Stephen Rhein
smrhein@udel.edu

Chandra Kambhamettu
<https://www.eecis.udel.edu/wiki/vims/>

Video/Image Modeling and Synthesis Lab
University of Delaware
Newark DE, USA

Abstract

Reflective and specular surfaces are problematic for traditional reconstruction techniques. Light projects non-linearly in scenes with these surfaces, and existing techniques to model this are poorly suited for real world applications. Accurately modeling the reflective surface is difficult without complete knowledge of the scene. To overcome this problem, we propose using different modalities of stereo vision to capture both the reflecting surface and the reflected scene. Using a four camera system consisting of a pair of visible wavelength cameras and a pair of long wave infrared cameras, we accurately reconstruct the reflective surface and ray trace reflected correspondences in the complementary modality. This approach allows for 3D reconstruction in the presence of a reflection, and does not require complete knowledge of the scene.

1 Introduction

Typical stereo reconstruction assumes a Lambertian scene, an assumption violated in scenes with reflective materials. Reconstructing objects in the presence of reflection is an ongoing area of research in computer vision, and some existing work is discussed in section 2. In essence, the problem is that without prior knowledge of the scene it is not possible to accurately model the reflecting surface and the reflected scene. Specular objects present a problem for many applications in vision as the images they form are a composite of their own surface and other components of the scene.

Reflectivity is wavelength dependent and different materials reflect and absorb different portions of the electromagnetic spectrum. Emissivity is the measure of the efficiency with which a surface emits thermal radiation, and metals typically have a very low coefficient values. This means that these surfaces emit very little thermal radiation. Long Wave Infrared (LWIR) Cameras detect light in the 8 – 14 μ m range, which is the range in which many materials emit radiation around room temperature. This allows the cameras to detect heat

given off by many surfaces. Surfaces with low emissivity do not effectively emit radiation in this band of the spectrum, which means light from these surfaces is reflected from other objects in the scene.

These reflective surfaces are not necessarily reflective to other modalities such as visible light. Brushed aluminum for example, is not very reflective in visible light but almost completely mirror-like in long wave infrared. Coatings like ink, paint, and anodization can have an effect on the emissivity, but in small amounts they do not affect appearance in LWIR, yet are apparent in visible wavelengths. This means that a textured surface in the visible band can appear highly mirrored in LWIR and vice versa. We propose using multiple modalities of imaging system to capture both the reflective surface as well as the reflected scene. We use a four camera system consisting of a visible band stereo pair and long wave infrared stereo pair. By using these different modalities we can simultaneously extract the reflecting surface, as well as capture the reflected scene. This allows for accurate reconstruction of the reflected scenes via ray tracing, and can be applied to a wide array of scenarios with reflecting surfaces. This sort of setup could be used in security applications where the camera system needs to stay out of sight, or in environments which could be hazardous to the cameras. We demonstrate that this approach works in both modalities, reconstructing a visible band scene as well as a LWIR scene using the other modality to extract the reflecting surface.

The rest of this paper is structured as follows. In section 2, we will elaborate on the problem and discuss other works which have dealt with it. In section 3, we discuss our approach. In section 4, we will describe our experimental setup. In section 5, we provide our results, and in section 6, we conclude with a brief summary.

2 Motivation and Related Work

Our environment contains many specular surfaces, and these surfaces offer challenges to most vision systems. Disambiguating a reflection from other objects in a scene can be challenging for human beings. Reconstruction in the presence of specular surfaces has been the focus of many works. Much of the work has been focused on refraction, such as [10, 11, 12, 13]. [10] used a 4D lightfield to analyze reflected scenes, leveraging the epipolar plane image space to separate reflected layers of the scene. [11] showed that Helmholtz reciprocity could be leveraged to exchange view and light sources in a scene using diffuse and specular reflections. Other works have built on this to recover 3D shape using time of flight cameras [14]. [15] used techniques from raytracing to reconstruct objects under refraction. [12] analyzed the relationship between the image of a calibrated scene and the geometry of a curved mirror surface on which the scene is reflected. [13] used a single camera viewpoint and the reflection of a planar target placed at two different positions to reconstruct a reflecting surface. Detecting specularities and identifying specular objects in scenes is an active area of research [8]. Additionally, LWIR and visible images have been used to classify materials [16], and this approach could be used to identify materials in the scene which are poorly emissive and reflective.

Unlike these works we aim to extract the reflecting surface by using different modalities of imaging. We not only reconstruct the reflecting surface, but the reflected scene using ray tracing. We use two calibrated stereo systems to reconstruct scene points on the reflecting surface in the visible band in which the surface is not reflective. Correspondences from the infrared cameras can then be ray traced and triangulated to find the 3D point in the scene. This technique works in both direction, as non reflective materials in the visible band can be

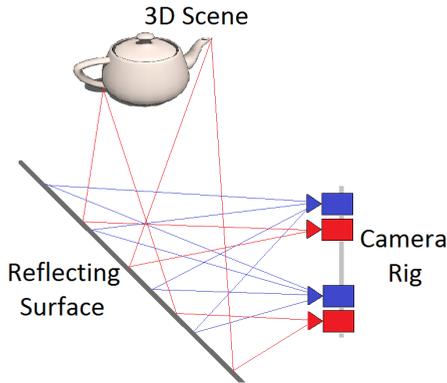


Figure 1: An illustration of the proposed approach. The visible band cameras in blue reconstruct the reflecting surface. The reflected 3D scene is captured by the LWIR cameras in red, and ray tracing is used to reconstruct the scene.

very reflective in LWIR and vice versa. We demonstrate this technique in both directions, reconstructing both visible and LWIR scenes.

3 Methods

In this section we will discuss our overall approach to reconstruction in the presence of reflection. We use a 4-camera rig consisting of two stereo pairs of different imaging modalities, namely long wave infrared and visible cameras. As illustrated in Fig. 1, we use different modalities to reconstruct points on the reflecting surface. Corresponding points from the reflected scene are then reconstructed via raytracing. Below we will coarsely divide this approach into a technique for calibrating, techniques for extracting the reflecting surface, and lastly a technique for stereo matching and reconstruction.

3.1 Calibration

Our system consists of four cameras, operating as two stereo pairs. These stereo pairs are of different imaging modalities and calibration is not a trivial matter. While these stereo pairs operate largely independently of each other, our approach requires a common coordinate system, necessitating calibration. First, each stereo system is calibrated independently. In the visible spectrum this is easily done with off the shelf calibration methods such as [9]. For long wave calibration the problem is more difficult. We use the method outlined in [14] which is briefly summarized below.

Essentially a ceramic backed paper calibration pattern is heated under a heat lamp. This causes the pattern to be visible in LWIR imagery due to increased heat and a change in emissivity from the printed surface. The pattern is however not uniformly heated. Artefacts of this process are mitigated by a pre-processing technique which involves masking out the calibration pattern using Otsu’s method [15]. The masked region undergoes iterative quadric fitting in the intensity space. This quadric is subtracted from the intensity image, and top hat filtering is applied. These steps are repeated and finally a sharpening filter is applied to the

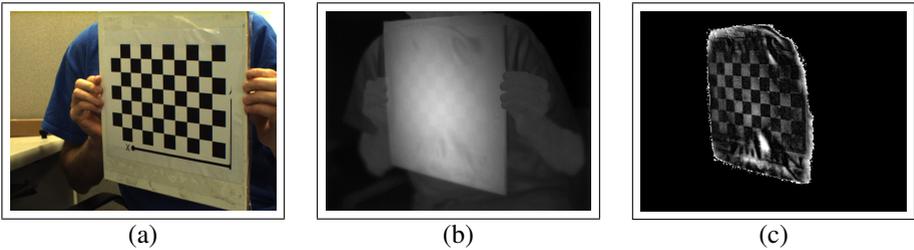


Figure 2: The calibration process: (a) The calibration board in the visible band; (b) The calibration board in LWIR. The low intensity variation makes it very difficult to calibrate; (c) The LWIR image after pre-processing. The image can be run through standard calibration methods now

images. This pre-processing technique allows standard calibration methods to be applied to the LWIR images.

Calibrating between modalities however requires a simple modification of this technique. This technique aims to use standard calibration tools for LWIR images, but in order to use it across modalities it is important to take emissivity into account. The printed pattern we use is a common checkerboard pattern used in many calibration techniques. In visible images this pattern consists of dark black printed squares and white spaces from the paper. In the LWIR imagery we similarly see dark and light squares but the cause is different. The surface of the calibration pattern is radiating heat, and the printed pattern changes the emissivity, in the case of our printed pattern the black toner is more emissive, radiating more energy and therefore higher intensity in the images. We therefore invert the intensity in the masked region of the pre-processed image, which allows for simple cross modality calibration using existing tools. The results of this process are shown in Fig. 2.

3.2 Extracting the Reflecting Surface

In this section we discuss our technique for extracting the reflecting surface in the scene. To accomplish this we fit a plane to correspondences reconstructed on the surface, however given dense correspondences far more complex surfaces could be used. We do not aim to detect specularity or identify which regions of the image constitute reflections. In this work we assume the position of the reflective surface is known in image space. In our experiments the reflective surface occupies nearly the entirety of the view in the images captured from all four cameras. The problem then becomes reconstructing the surface. This can be done using the other stereo vision modality.

In section 4.1, we show it is possible to add texture to a surface that is visible in one modality and not greatly affect the imagery from other modality. Adding texture enables the use of standard feature matching techniques, such as SURF matching [2]. To add texture to visible imagery that is invisible in LWIR we write on the material with a marker. Adding texture to LWIR imagery without affecting the visible imagery can be done by adding heat to an emissive surface. In section 4.1 we heat up emissive surfaces by placing a gloved hand on the surface for a few seconds before imaging. The resulting hand print is visible only in infrared.

While these methods of adding texture to both modalities of image are active and require physical access to the surface itself, it is easy to imagine using a pattern of structured light or

heat source that would be visible in one modality and not the other. Some materials, such as galvanized steel are already quite textured in the visible band, and depending on grain and polish could be effectively used without need for modification.

After adding texture we capture synchronized images with all four cameras. Within the non reflected modality with added texture, we match image features. In our work, SURF points are detected and matched. These matches are then triangulated using the method outlined in [2] to form a sparse point cloud. We place a threshold on Euclidian error and eliminate points from the sparse cloud with a high triangulation error, which helps ensure quality results.

To ray trace correspondence from the other stereo pair an implicit surface is needed. Since this is the first attempt using different modalities to extract the reflecting surface, we have modeled a simple reflecting surface, namely a plane, however, more complex surfaces could be modeled if dense correspondences can be found and a surface fit to the reconstructed points. A plane is defined by its normal and an origin. We perform a principal component analysis of the sparse point cloud, taking the third set of coefficients as the normal, and the centroid is taken as the origin. This plane is used to model the reflecting surface, and its implicit form

$$P \cdot n + d = 0 \quad (1)$$

can be used to intersect arbitrary rays for use in the reconstruction phase. Where P is the origin n is the plane normal and d is a constant.

3.3 Stereo Matching

Reconstructing the reflected scene requires correspondences. Stereo matching is an ongoing and active research area and LWIR stereo has been studied [9]. Stereo matching in this modality is challenging due to low variance in intensity. Additionally, LWIR cameras are typically lower resolution and have limited optics. These problems coupled with reflection make for challenging stereo matching and lead to noise in the reconstruction. Much of the research effort in stereo matching has been focused in the disparity domain, which requires rectified images. In a reflected scene rectification parameters from camera calibration will no longer accurately rectify the scene. In a scene with a more complex reflecting surface rectification may not be possible, and feature matching would be the best option.

To facilitate dense correspondence matching in the presence of reflection we calculate new rectification parameters using uncalibrated rectification [2]. Once the images are rectified, typical disparity matching techniques can be used. We employ Semi Global Block Matching (SGBM)[8] due to its record for good performance [8]. This facilitates dense correspondences, but these correspondences are of reflected objects, so reconstruction is not simply a matter of triangulation.

3.4 Ray Trace Reconstruction

To reconstruct the dense correspondences of the reflected scene we employ techniques from ray tracing. Ray tracing involves projecting rays through each pixel $p_i = [x_i, y_i]$ into a 3D scene. We refer the reader to [10] for an overview of ray tracing. The equation for a ray V_i is given by

$$V_i = C_0 + t \cdot \frac{\beta}{\text{norm}(\beta)} \quad (2)$$

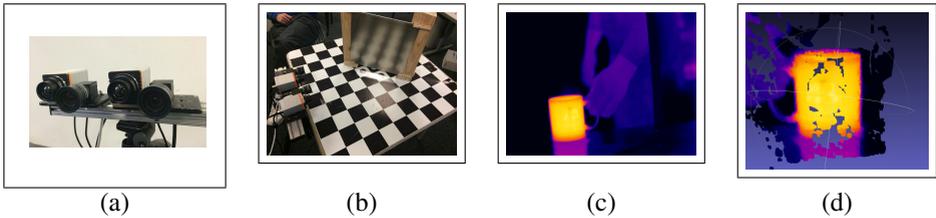


Figure 3: The experimental setup: (a) The 4 camera system with LWIR and visible cameras; (b) The experimental reflection rig; (c) A LWIR reflection of a mug with hot liquid; (d) Reconstructed results for the mug.

where C_0 is the camera center, and

$$\beta = R' \cdot A^{-1} \cdot [x_i, y_i, 1] \quad (3)$$

where A is the camera matrix, and R is the camera rotation matrix.

These rays can then be intersected with surfaces in the scene. In our approach, we model the reflecting surface as a plane defined in equation 1. The intersection I_i of a given ray and a plane is obtained by substituting P from equation 1 with V_i from equation 2, and solving for t . We can then plug t back into equation 2 to find I_i the intersection point. Reflection is defined by,

$$V_{reflected} = V_i + (2 \cdot n \cdot C_i) \quad (4)$$

where

$$C_i = -(n \cdot V_i) \quad (5)$$

where, $V_{reflected}$ is the direction of the reflected ray, V_i is the camera ray defined in equation 2, and n is the surface normal. Thus the complete reflected ray is defined as $I_i + t_i * V_{reflected}$.

Corresponding sets of reflected rays are then triangulated by solving for the closest point of intersection using least squares similar to the triangulation used in section 3.2. Euclidian error thresholding is again applied to ensure a quality reconstruction.

4 Experiments

In this section we outline the experimental setup and present our tests to validate our method. For all experiments we utilize the same 4 camera setup shown in Fig. 3(a). This setup consists of 2 visible cameras, Point Grey Flea2G's capturing at 1280 x 960 resolution. The long wave infrared cameras are Xenics Gobi-640-GigE's capturing at 640 x 480 resolution and 50 mK thermal sensitivity. The whole setup is synchronized by software trigger to within a few milliseconds. The system was calibrated using the method described in section 3.1. We conduct a number of experiments with this system to test various aspects of our proposed approach.

4.1 Cross Modality Texture Experiment

To demonstrate that texture can be added in one modality while remaining invisible in the other, we capture images of a surface in both visible and long wave infrared, add texture to the surface and image it again. First a baseline image is taken, followed by a control image where no texture is added. We then add texture to the surface, and capture images

again. The control image is to find the natural variation from pixel drift and noise. We compare 7 materials adding texture in only one modality. The metals are highly reflective in the LWIR band but far less in the visible band. The plastic mirror is highly reflective to visible light, but not reflective in LWIR. For the metals and the whiteboard we add texture in the form of a marker which is apparent in the visible spectrum. For the plastic mirror and phenolic sheet, we placed a hand with a polyethylene glove on the surface to transfer heat to the surface without leaving a smudge that would be detectable in the visible spectrum. Additionally, we have conducted a short experiment to show how surface corrosion affects reflection and emissivity by comparing a corroded piece of galvanized steel with a polished one. We present results comparing the baseline to both control and textured images for both modalities in section 5.1.

4.2 Reflecting Surface Extraction

To validate our approach to extract the reflecting surface, we set up an experimental rig shown in Fig. 3(b) where different materials can be swapped in and out in a way that the surfaces are oriented and positioned the same each time. For a baseline measurement we placed a Lambertian textured surface, and reconstruct the surface as outlined in section 3.2. Subsequent materials are imaged and we compare surface orientation using cosine similarity. Results are reported in section 5.2.

4.3 Reconstruction Experiments

To evaluate our reconstruction technique we reconstruct objects reflected in each modality. Quantitative reconstruction results are obtained by comparing the reconstructed models to ground truth measurements made on the objects using a ruler and caliper. For visible reflection we reconstruct a textured cube. We measure the visible faces and compare the reconstructed result to ground truth measurements. For LWIR reflection we reconstruct the camera system itself as well as a mug filled with hot water, and compare to physical measurements made of the lenses and camera bodies, as well as the mug.

5 Results

In this section we present the results from the experiments described in section 4. We further analyze the results and briefly discuss the implications on our proposed methodology.

5.1 Cross Modality Texture Results

As outlined in section 4.1, we compare our baseline image to a control as well as a textured image, and results are shown in table 1. Results are reported in absolute mean pixel intensity difference. Note that the LWIR images are captured as 16 bit intensity images, shown color mapped in our figures. The visible band images are 8 bit, which explains in part why there is such a large variance in the LWIR images. This variance can be seen even in the control images, however when we add LWIR texture (second set of materials in table 1) we see a dramatic difference from other tests. These results demonstrate that it is indeed possible to add texture in one modality without affecting the other.

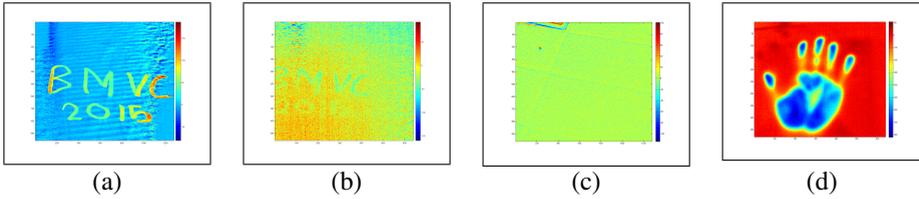


Figure 4: Difference images in the visible band and LWIR: (a) The difference image from writing on a surface in visible band; (b) The difference image from writing on a surface in LWIR band; (c) The difference image of a hand print in the visible band; (d) The difference image of a hand print in the LWIR band.

To further illustrate this we show a series of difference images in Fig. 4. Figures 4(a) and 4(c) are visible band difference images, and 4(b) and 4(d) are LWIR difference images. In (a) and (b) writing with a marker has been added to the surface and in (c) and (d) a hand has been placed on the image. The writing is clearly apparent in the visible image but not the LWIR image. Similarly the hand print is not apparent in the visible image, but obvious in LWIR. These images testify that texture can be added in one modality without affecting the other greatly. In Fig 4(b) and 4(c), the difference image mostly shows noise as well as some reflected parts of the room, which may have moved slightly relative to the camera or surface.

We compared corroded galvanized steel to polished steel by heating both with a gloved hand. Results are presented in table 2. The polished steel is much less emissive, and does not clearly show any signs of the added LWIR texture, however the corroded surface is more emissive, and therefore not only shows the added texture, but does not reflect. This shows that surface properties are critical, and even the same material can have drastically different reflection and emission based on corrosion.

	Control Visible	Textured Visible	Control LWIR	Textured LWIR
Polished Aluminum	1.24	7.00	80.96	2.20
Unpolished Aluminum	0.58	2.60	0.24	10.51
Galvanized Steel	0.09	12.97	20.04	65.71
Brushed Aluminum	0.22	4.94	4.87	15.72
Whiteboard	0.36	30.67	49.32	36.80
Plastic Mirror	0.62	1.10	5.89	150.55
Phenolic Sheet	0.34	0.56	7.31	194.07

Table 1: Results from experiment in section 4.1. Results are presented in absolute mean pixel intensity difference.

	Control Visible	Textured Visible	Control LWIR	Textured LWIR
Polished	0.20	1.7	2.43	17.01
Corroded	2.67	2.91	8.45	215.30

Table 2: Results from experiment in section 4.1 on Galvanized steel with and without corrosion. Results are presented in absolute mean pixel intensity difference.

5.2 Reflecting Surface Results

Results for experiment 4.2 can be found in table 3. These results show that by adding texture to a surface it is possible to extract these surfaces even though they are typically considered specular. The plastic mirror has been extracted using the LWIR modality, and has the highest error in part because of low resolution and texture in the images.

Material	Normal similarity
Unpolished Aluminum	0.97989
Polished Aluminum	0.84069
Plastic Mirror	0.81880

Table 3: Results for extracting the reflecting surface for reflective materials

5.3 Reconstruction Results

To demonstrate that the proposed reconstruction technique effectively handles reflection we have reconstructed a self-portrait of the camera system. We have placed our camera system in front of an aluminum plate which is highly reflective in the LWIR band, but much less reflective and textured in the visible band. Sample Visible and LWIR images are shown along with the resulting reconstructed model in Fig. 5. Note that the positions of cameras appear reversed between Figures 5(b) and 5(c), this is due to the fact that (b) shows a reflected image. Our approach captures the real geometry, and so the cameras are in the correct orientation. To obtain quantitative results we measure the camera system with a ruler and caliper, measuring the lenses and camera bodies where the reconstruction is not overly noisy. In total we took 6 measurements and report the RMS error in table 4. Additionally, we reconstructed a mug as shown in Fig. 3(c) and 3(d). We measured the height and radius of the mug in 5 places and compare to the reconstructed model.

For the visible band we reconstructed a textured box. The box has two faces visible in the reflected image, and we measure the seven edges and four hypotenuse of the reconstructed results. The RMS is reported in table 4. The visible results are less noisy and the reconstructed model looks better, but the sensor is higher resolution, and there is more contrast and less noise in the images. The LWIR reflected scene requires thermal variation in contrast, and most objects will come to thermal equilibrium with their environment over time. This makes the problem especially difficult in this modality, and is among the reasons these scenes were selected, as they contain objects which are hotter than the environment.

	RMS
Visible Cube Reconstruction	8.71 mm
LWIR Camera System Reconstruction	11.36 mm
LWIR Mug Reconstruction	6.34 mm

Table 4: Reconstruction results for the reflected scenes outlined in 4.3

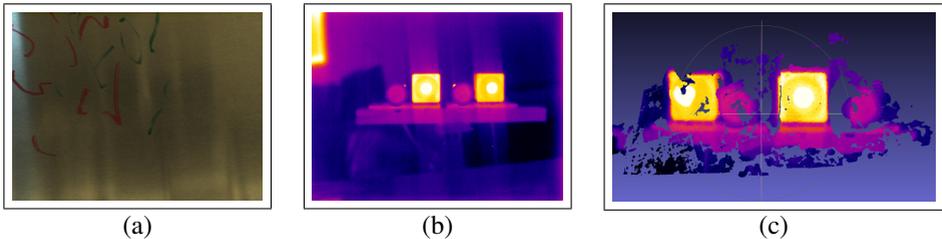


Figure 5: Results from reconstructing a self portrait of our camera system: (a) The reflecting surface with texture in the visible band; (b) The camera system reflected in LWIR; (c) The reconstructed model.

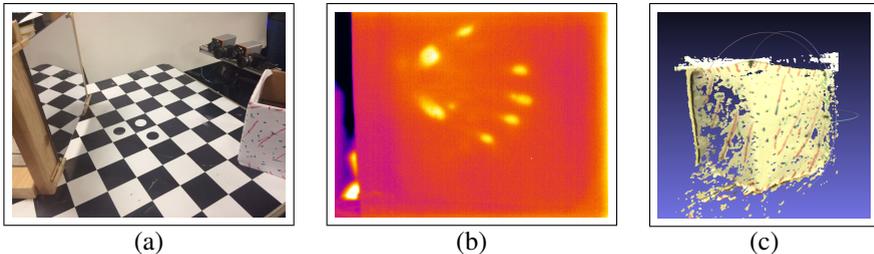


Figure 6: The visible band reconstruction results: (a) The reflection experimental setup; (b) The textured surface in LWIR; (c) The visible band reconstruction results.

6 Conclusion

6.1 Summary

In this work we have proposed using different modalities of imaging to reconstruct objects in a reflected scene. This scenario is difficult for traditional reconstruction approaches which assume a Lambertian scene. Many common materials are reflective, however different wavelengths are reflected differently and many objects are more reflective in certain parts of the spectrum. Using different modalities of stereo imaging, it is possible to simultaneously capture the reflecting surface as well as the reflected scene. We have demonstrated that textured surfaces in one modality can be reflective in another, and texture can be added to these surfaces in one modality without affecting the other. By accurately extracting the reflecting surface it is possible to ray trace correspondences from the reflected scene. We have shown that our technique works in both directions, reconstructing both visible and LWIR reflected scenes using the other modality to accurately model the reflecting surface.

6.2 Future Work

We have used visible band and long wave infrared cameras in this work, however this technique should be applicable to other modalities such as near infrared or ultraviolet imaging. Additionally, for simplicity we have used a reflecting plane in this work, however an implicit surface of greater complexity should work equally well given enough correspondences on the reflecting surface. We would also like to experiment with different materials and coatings like paint to quantify the affect these have on reflectivity and emissivity.

References

- [1] M. Alterman, Y.Y. Schechner, and Y. Swirski. Triangulation in random refractive distortions. In *Computational Photography (ICCP), 2013 IEEE International Conference on*, pages 1–10, April 2013. doi: 10.1109/ICCPPhot.2013.6528314.
- [2] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. *Computer Vision and Image Understanding (CVIU)*, 110:346–359, 2008.
- [3] Thomas Bonfort, Peter Sturm, and Pau Gargallo. General specular surface triangulation. In *Asian Conference on Computer Vision 2006*, 2006.
- [4] J.Y. Bouguet. Matlab camera calibration toolbox. 2000.
- [5] Andrey DelPozo and Silvio Savarese. Detecting specular surfaces on natural images. In *CVPR*. IEEE Computer Society, 2007. URL <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2007.html#DelPozoS07>.
- [6] Jannik Fritsch, Tobias Kuehnl, and Andreas Geiger. A new performance measure and evaluation benchmark for road detection algorithms. In *International Conference on Intelligent Transportation Systems (ITSC)*, 2013.
- [7] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, New York, NY, USA, 2 edition, 2003. ISBN 0521540518.
- [8] Heiko Hirschmuller. Accurate and efficient stereo processing by semi-global matching and mutual information. In *CVPR (2)*, pages 807–814. IEEE Computer Society, 2005. ISBN 0-7695-2372-2. URL <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2005-2.html#Hirschmuller05>.
- [9] S.J. Krotosky and M.M. Trivedi. A comparison of color and infrared stereo approaches to pedestrian detection. In *Intelligent Vehicles Symposium, 2007 IEEE*, pages 81–86, June 2007. doi: 10.1109/IVS.2007.4290095.
- [10] Nigel J. W. Morris and Kiriakos N. Kutulakos. Dynamic refraction stereo. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(8):1518–1531, 2011. URL <http://dblp.uni-trier.de/db/journals/pami/pami33.html#MorrisK11>.
- [11] Nigel Jed Wesley Morris. Image-based water surface reconstruction with refractive stereo. Technical report, 2004.
- [12] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *Systems, Man and Cybernetics, IEEE Transactions on*, 9(1):62–66, Jan 1979. ISSN 0018-9472. doi: 10.1109/TSMC.1979.4310076.
- [13] Philip Saponaro, Scott Sorensen, Abhishek Kolagunda, Stephen Rhein, and Chandra Kambhamettu. Material classification with thermal imagery. In *CVPR*, 2015.
- [14] Philip Saponaro, Scott Sorensen, Stephen Rhein, and Chandra Kambhamettu. Improving calibration of thermal stereo cameras using heated calibration board. In *ICIP*, 2015.

- [15] Silvio Savarese and Pietro Perona. Local analysis for 3d reconstruction of specular surfaces. In *CVPR (2)*, pages 738–745. IEEE Computer Society, 2001. ISBN 0-7695-1272-0. URL <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2001-2.html#SavareseP01>.
- [16] Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R. Marschner, Mark Horowitz, Marc Levoy, and Hendrik P. A. Lensch. Dual photography. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, pages 745–755, New York, NY, USA, 2005. ACM. doi: <http://doi.acm.org/10.1145/1186822.1073257>.
- [17] Peter Shirley and Steve Marschner. *Fundamentals of Computer Graphics*. A K Peters, Wellesley, MA, 3 edition, 2009. ISBN 978-1-56881-469-8.
- [18] Scott Sorensen, Abhishek Kolagunda, Philip Saponaro, and Chandra Kambhmettu. Refractive stereo ray tracing for reconstruction underwater structures. In *ICIP*, 2015.
- [19] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mounqi G. Bawendi, and Ramesh Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nat Commun*, 3:745+, March 2012. ISSN 2041-1723. doi: 10.1038/ncomms1747. URL <http://dx.doi.org/10.1038/ncomms1747>.
- [20] Sven Wanner and Bastian Goldluecke. Reconstructing reflective and transparent surfaces from epipolar plane images. In Joachim Weickert, Matthias Hein, and Bernt Schiele, editors, *GCPR*, volume 8142 of *Lecture Notes in Computer Science*, pages 1–10. Springer, 2013. ISBN 978-3-642-40601-0. URL <http://dblp.uni-trier.de/db/conf/dagm/gcpr2013.html#WannerG13>.
- [21] Jinwei Ye, Yu Ji, Feng Li, and Jingyi Yu. Angular domain reconstruction of dynamic 3d fluid surfaces. In *CVPR*, pages 310–317. IEEE, 2012. ISBN 978-1-4673-1226-4. URL <http://dblp.uni-trier.de/db/conf/cvpr/cvpr2012.html#YeJLY12>.