Stereo Vision based Optic Nerve Head 3D Reconstruction using a Slit Lamp fitted with Cameras: Performance Trial with an Eye Phantom

Ian Coghill
Mario E Giardini

This is an author-generated copy of the accepted version.

Stereo Vision based Optic Nerve Head 3D Reconstruction using a Slit Lamp fitted with Cameras: Performance Trial with an Eye Phantom

Ian Coghill
Department of Biomedical Engineering
University of Strathclyde
Glasgow, Scotland, UK
i_coghill@outlook.com

Mario Ettore Giardini
Department of Biomedical Engineering
University of Strathclyde
Glasgow, Scotland, UK
mario.giardini@strath.ac.uk

Abstract—**Preventing vision loss from glaucoma, the second leading cause of blindness globally, relies on early detection and intervention. A key clinical indicator for early detection is structural change of the optic nerve head (ONH), which impacts the ONHs three-dimensional footprint on the retina. Clinically, such footprint is commonly assessed by subjective observation through a slit lamp, with significant inter- and intra-observer variability. Objective devices that rely on 3D imaging show promise in increasing diagnostic accuracy within clinics but are prohibitively expensive for many. In previous work, we demonstrated the initial feasibility of obtaining 3D surface images by applying stereo vision techniques to a slit lamp fitted with low-cost cameras. In this paper, we present the assessment of the precision and accuracy of our system with ONH targets in an eye phantom. Quantitative reconstructions, referenced to ground truth data, were obtained. Current results, albeit promising, suggest that more work is likely needed to improve accuracy and precision to a clinically significant level.**

Keywords—*3D Reconstruction, Optic Nerve Head, Stereo Vision, Slit Lamp, Eye Phantom*

I. INTRODUCTION

A. Problem

The second leading cause of blindness globally is glaucoma [1]. Its most prevalent form, primary open angle glaucoma (POAG) [2], is a slowly progressing optic neuropathy that irreversibly damages the neurons, more specifically retinal ganglion cells (RGCs), which connect the retina to the brain, leading to irreversible vision loss. Its irreversible nature means that early detection and treatment, which can halt or slow progression, is vital. Since visual degradation is often not noticed by the patient, or even detectable through visual testing [3], until the condition is advanced, early detection usually relies on detecting structural changes. Such changes impact the three-dimensional shape of the optic nerve head (ONH), the point at which the axons of the RGCs meet to exit the eye and form the optic nerve, directly visualizable using retinal imaging techniques, causing its dimple-like structure to appear more and more excavated as damage progresses, thus making it a key clinical indicator.

One of the major reasons that people are still becoming blind from POAG is that over half of those that suffer from it are unaware of it [4]. Part of this can be attributed to cases being missed by clinicians, though quantitative data is unclear [4]. Indeed, in clinics, the ONH is most commonly assessed in a subjective manner, by qualitative observation of a stereoscopic view of the structure [5], either directly using a slit lamp or in photographs taken with a fundus camera. Such assessment is not straightforward, leading to significant inter and intra-observer variability, even for specialists [6]. Quantitative imaging devices, such as optical coherence tomography (OCT) systems and the Heidelberg Retinal Tomograph (HRT, effectively a variant of a laser scanning ophthalmoscope), which are able to provide an objective assessment based on 3D surface reconstruction (or full depth 3D reconstruction in the case of OCT), have been shown to provide better diagnostic accuracy than general ophthalmologists who, in turn, have a higher level of training than the primary eye care health professionals who perform front-line referrals to hospital services. These systems are prohibitively expensive for many services. Much work, including [7]–[11], has been carried out on the idea of performing 3D surface reconstruction of the ONH by applying stereo vision techniques to stereo images acquired of it using a fundus camera, however these systems are still reliant on costly image acquisition and processing hardware.

B. Project Aim and Objective of this Work

Our work has focused on applying computer stereo vision techniques to retinal stereo image pairs obtained using a slit lamp biomicroscope with low-cost consumer-grade cameras fitted to its eye pieces. We previously reported on demonstrating the initial feasibility of this idea on ONH targets in a phantom eye [12]. With the present article, we report on the current accuracy and precision of the technique in reconstructing the surface of ONH targets in an eye phantom.

II. EXPERIMENTAL SETUP

A photograph, and CAD rendered 3D model, of the experimental setup (involving both the imaging system and eye phantom) are provided in Fig. 1. The photograph shows the eye phantom fixed in position, on a frame constructed from aluminum profile (KJN Automation Ltd, Leicester, UK), in front of the slit lamp lens of the imaging system.
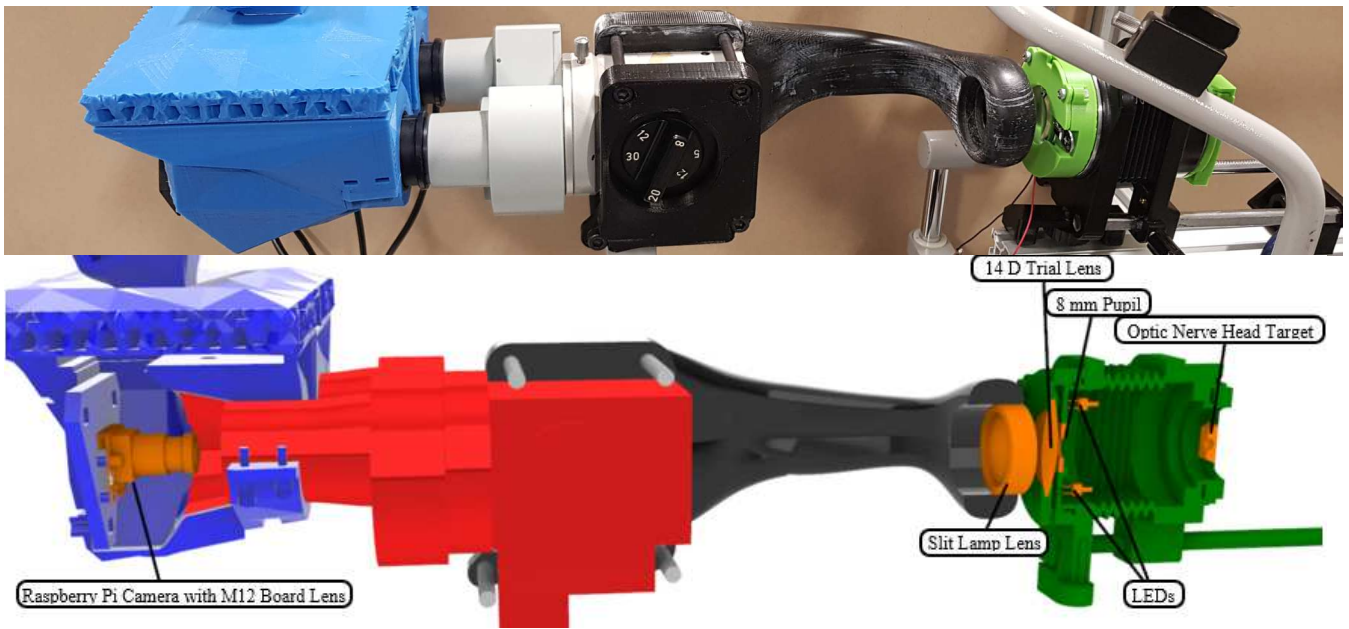
Fig. 1. Photograph (**TOP**) and corresponding color-coded and labelled CAD rendered 3D model with some sections cut to reveal internal components (**BOTTOM**) of the experimental setup. **Blue** - 3D printed camera mount; **Red** – Zeiss 30 SL-M slit lamp; **Black** – Slit lamp lens mount; **Green** – Eye Phantom; **Orange** – Imaging/optical/illumination components (excluding those of the slit lamp) and ONH target.

A. Imaging System

The assembled system featured a mount (shown in black in the CAD model in Fig. 1) which was clamped onto the slit lamp (shown in red) and used to hold a slit lamp lens, required for imaging the back of the eye with a slit lamp, in position in front of its objective lens. It also featured two cameras mounted, using a mount (shown in blue), onto its eye pieces. The slit lamp used in this work was a Zeiss 30 SL-M (Carl Zeiss AG, Oberkochen, Germany).

1) Indirect lens: The slit lamp lens used within this setup was a 90 D SuperField (Volk Optical, Mentor, OH, USA), a lens which is commonly available within clinics. This lens was press fit into a ring part, which was itself fit into the main mounting arm, allowing some adjustment of the lens' axial position for focusing. The lens was positioned approximately 140 mm in front of the slit lamp's objective lens. The parts used to mount the lens were all printed in black PLA+ (Sunlu, Los Angeles, California, US) on a Wanhao Duplicator i3 3D printer (Wanhao, Jinhua, China).

2) Cameras: Raspberry Pi NoIR v2.1 Cameras (Raspberry Pi Foundation, Cambridge, UK), which feature an 8-megapixel CMOS sensor, were used. Their stock lenses (f = 3.04 mm) were removed and replaced with M12 board lenses with a focal length of 8 mm (Hongkong Ansice Industrial Co, Ltd, Hong Kong, China) in order to obtain a narrower field of view. Replacement required that the stock lens be removed, the sensors housing be rigidly fixed to the PCB board backing it, using Loctite 4305 (Loctite, Düsseldorf, Germany), and screwing the board lens into a mount (3D printed in black PLA+ on a Wanhao Duplicator i3 3D printer) which iteslf was secured to the PCB board using Loctite 4305. The cameras were connected to a single StereoPi board (virt2real, St. Petersburg, Russia), containing a Compute Module 3 (Raspberry Pi Foundation, Cambridge, UK), which enabled viewing of a video stream from the cameras on a monitor, and image capture.

B. Eye Phantom

The eye phantom featured a single 14 D bi-convex trial lens (BiB Ophthalmic Instruments, Stevenage, Hertfordshire, UK), to provide refractive power; an air-filled interior; a 3D printed iris with an 8 mm pupil; interchangeable 3D printed ONH targets; and two 2 mm flat top white LEDs (Kingbright, Issum, Germany). These components, all shown in orange in the CAD model in Fig. 1, were held in place using 3D printed mounts fitted to the back and front of a camera macro-bellows (Fotga, Shenzhen, China). Such mounts and the macro-bellows are shown in green. All 3D printed components of the phantom were printed in green PLA (RS Components Ltd, Corby, UK) on an Ultimaker 2+ 3D printer (Ultimaker B.V., Utrecht, Netherlands), apart from the ONH targets, which were printed in ecoPLA (3Djake, Paldau, Austria) using the same printer. A calibration target, featuring a checkerboard pattern, required to calibrate the imaging system, was also created.

1) Optical Design: For simplicity at the current stage of research, the optical setup of the eye phantom was intended to mimic an emmetropic eye (i.e., one that is infinity focused). Albeit matching a human eye in terms of optical geometry and subtended (angular) sizes, the phantom's geometrical size was set to be 4.3 times larger than the typical human eye, for simplicity of handling. At this early stage of the research, the optics of the eye phantom were chosen to allow for refractive power, but not mimic the more complex optical aspects of the human eye, such as spherical aberration, coma, astigmatism and chromatic aberration.

2) Optic nerve head targets: Two different ONH targets, representing different ONH morphologies, were created – one resembling an ONH without pathology and the other resembling an ONH with significant glaucomatous damage. A photograph of these is provided in Fig. 2. The dimple structures are enlarged, according to the scale factor of the eye phantom,

Fig. 2. Photographs of ONH targets. **LEFT**: glaucomatous, **RIGHT:** normal.

so that they appear to be of the same size as a life-sized ONH in a life-sized eye when viewed through the optics of the phantom. The structures have a depth and diameter of around 3 mm and 8.6 mm, respectivey. The process taken to design and fabricate these is detailed in our previous work [12]. In this case, however, a larger prism has been used for embedding the dimple-like structures. To aid alignment of the reconstructions with ground truth data, 4 ink marks were made near the bottom of the dimple structures and 4 near the top surfaces, to act as optical anchor points for referencing the measurements to the ground truth data.

3) Calibration target: The calibration target contained a 10 x 7 checkerboard of squares, created photographically on Rollei RPX 25 black and white film (Rollei, Hamburg, Germany) by photographing a Philips BDM4350UC 43" monitor (Philips, The Netherlands) with 4K resolution displaying the pattern, using a Yashica FX3+ camera (Yashica, Japan) camera with stock 50 mm lens, and then having the film developed by a commercial developing lab. It was glued flat to a standard glass microscope slide using Loctite 4305. A thin white plastic film (from packaging) was taped to the other side of the slide to provide a white backdrop for the film which was transparent with the black squares of the checherboard. The target was retroilluminated by a 5 mm through-hole white LED (RS Components Ltd, Corby, Northamptonshire, UK). The size of the individual squares within the checkerboard were determined by photogrammetery against a reference scale, in this case the legs of a set of calipers (Mitutoyo Ltd., Andover, Hampshire, UK) set at 10 mm. In this, a length of squares was measured and the value divided by the number of squares measured. This was carried out with 4 different images, measuring different sides of the checkerboard, and the mean square size determined was 1.165 mm (sample standard deviation [SD] = 0.005 mm).

III. Performance Trial

The imaging system was used to capture images of the ONH targets and calibration target in the phantom, which were then processed using a stereo vision pipeline before quantitative validation against ground truth data was performed.

A. Setup and Image Capture

The magnification setting on the slit lamp was set to 30 and fixed at this for the entirety of the trial. The cameras were focused to infinity and the eye pieces set to their zero-diopter setting. The axial position of the slit lamp lens was adjusted such that the imaging system was focused to infinity – allowing imaging of an emmetropic eye. One of the ONH targets was placed on the retina of the phantom, and on its optical axis. The LEDs that are inside the phantom that provide illumination to the targets from inside, were connected to a power supply and powered. Stereo images were acquired of the ONH target

already inside the phantom. The target was then rotated 90° around the phantom's optical axis and again imaged. This was repeated until 4 images were obtained, each with the target at a different rotation, and then the same was done for the other target after interchanging them, giving 4 repeated measurements of each target. The system was not re-aligned during the imaging session.

The LED of the calibration target was then connected to a power supply to power it. The target, held using a flexible helping-hand support (B07PYP4V5S, Amazon, Inc., Seattle, Washington, US), was then positioned in many different positions with different tilts, within the field of view of both cameras, and images taken at each position. In total, 100 calibration images were acquired.

B. Stereo Vision Pipeline

The calibration of the system, and interpolation of the sparse disparity map, were carried out using the MATLAB programming language and environment (The MathWorks, Inc, Natick, Massachusetts, US). All other tasks described below used custom code written in C++ using Visual Studio 2017 Community Edition (Microsoft Corp., Redmond, Washington, US).

1) Calibration: The distortion (3 radial and 2 tangential), intrinsic and extrinsic parameters of the stereo system (which include the eye phantom's lens) were estimated using the Camera Calibration Toolbox for MATLAB [13]. In order to gain some insight on whether a sufficient number of calibration images were taken, the parameters were estimated using 10 to 100 images, in increments of 10.

2) Rectification: Using the calibration parameters obtained from using 100 calibration images, the images of the ONH targets were corrected for distortion and rectified using the rectification functions of the OpenCV library [14]. This transformed the images such that corresponding points within the left and right images could now be found on the same horizontal scanline.

3) Stereo Matching: The stereo matching algorithm developed in this work followed a course-to-fine scheme, and included a texture threshold and left-right consistency check. Prior to performing matching, the images were converted to greyscale. The algorithm involved the following steps:

a) Initial course disparity estimation (Fig. 3, cell 1): A region of interest (shown in blue) is selected in the left image. The best matching patch of the same size (shown in red) is searched for in the right image. The best matching patch corresponds to that where the normalised cross-correlation (NCC), determined using the OpenCV library, is highest . The search range was the entire width of the right image (shown in green). The disparity value is equal to $R - L$ in units of pixels.

b) Creation of a grid of points (Fig. 3, cell 2): The region of interest is filled with a grid of points to be matched. In this work a grid of 40 by 40 points was used.

c) Setting of initial patch size (Fig 3, cell 2): The initial height and width of the patch to be matched (shown in blue) were set to be 6 times the vertial and horizontal spacing of the points in the grid, respectively.
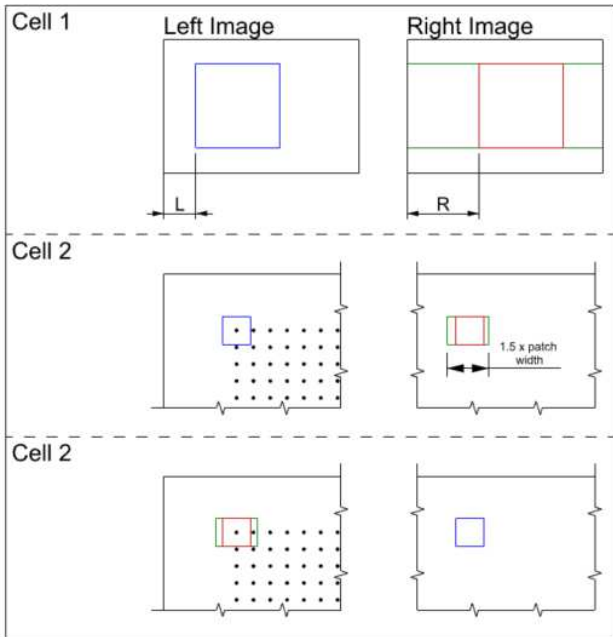
Fig. 3. Illustrations to complement the description of the stereo matching algorithm developed in this work.

d) Applying a texture threshold (Fig 3, cell 2): The population standard deviation of a normalised version of the patch to be matched is determined. If the value is below 0.2, the algorithm skips to stage i.

e) Forward matching (Fig 3, cell 2): The best matching patch of the same size (shown in red) is searched for within a search range in the right image (shown in green) using the same matching method as in stage a. The search range is 1.5 times wider than the patch to be matched, and is centred according to the disparity determined in either stage a or, after the first courseness level, the disparity determined for the point at the previosus coarseness level. To obtain a sub-pixel estimate of disparity, a quadratic equation was fitted to the NCC values corresponding to the disparity values in a ± 3-pixel range centred around the point of maximum NCC and the disparity at which the quadratic reached its maximum was deemed to be the subpixel accurate disparity.

f) Backward matching (Fig 3, cell 3): The best matching patch in the left image (shown in red) for the patch in the right image (shown in blue) is searched for within a search range in the left image (shown in green) using the same matching method as in stage e, to obtain the sub-pixel disparity.

g) Apply a consistency threshold: If the disparity values determined in the forward and backward matching stages differ by more than 1.2 pixels then the algorithm skips to stage i. If the difference is less, the disparity value stored for this pass is the mean of the values determined in the forward and backward matching stages.

h) Reduce patch size and perform matching: The height and width of the patch are reduced by a factor of 1.2. Stages d to g are repeated with the new size. Once the algorithm has processed 5 different patch sizes, the algorithm skips to step i.

i) Patch size threshold: If matching was successful (it passed stages d and g) for 4 patch sizes, the disparity value

determined in the last successful pass is stored for the final sparse disparity map.

j) Perform matching on next point: Stages c to i are repeated for the next point in the grid. This is repeated until all points have been processed.

4) Interpolation: Smoothing splines are first fit to the rows of the 40 x 40 array of points in the disparity map and used to determine the disparity value for each pixel along those rows. Following this, the same procedure is repeated for each column of the disparity map output from the previous stage, resulting in a dense disparity map.

5) Reprojection: In the rectification step of this algorithm, the system was mapped to the non-convergent ideal stereo model. This model permits points from the rectified left image (x, y) to be reprojected back into 3D scene coordinates (X, Y, Z), if the disparity (d) is known, using the reprojection matrix (denoted Q). The reprojection matrix (which includes the focal length (f), principal point (c_x, c_y), and baseline distance (b) of the stereo camera) is output by the rectification functions of the OpenCV library, and takes the following form:

$$Q = \begin{bmatrix} 1 & 0 & 0 & -c_x \\ 0 & 1 & 0 & -c_y \\ 0 & 0 & 0 & f \\ 0 & 0 & \dfrac{-1}{b} & 1 \end{bmatrix} \qquad (1)$$

Reprojection is governed by the following relationship:

$$Q \begin{bmatrix} x \\ y \\ d \\ 1 \end{bmatrix} = W \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \qquad (2)$$

where the term W is a normalization factor that scales the 4th element of the right-hand term to one. This equation was used to reproject the dense disparity map data (x, y, d), along with RGB data, into the 3D scene space, yielding a textured dense point cloud.

C. Quantitative Validation

1) Ground Truth Data Acquisition: Ground truth scans of each of the targets were obtained using the commercial photogrammetry software 3DF Zephyr Free (3Dflow, Verona, Italy). In turn, 50 JPG images were aquired of the top surface of each target from different and random poses using the camera of a Samsung Galaxy S8 smartphone (Samsung, Seoul, South Korea), set to capture at a resolution of 4032 x 1960 pixels. The images acquired of each target were input into the software to yield proportionally accurate ground truth reconstructions. These were uniformly scaled in three dimensions to the correct size using the mesh processing software CloudCompare (CloudCompare v2.9.1, 2018). The scaling factor for each was determined through comparing measurements made on the surface of the physical targets (using the same measurement method as that used to measure the calibration checkerboard) to those made on the unscaled ground truth reconstructions using CloudCompare. Some speckle features of the targets were used as points where

measurements could be made between. Four measurements were made for each target, and the mean scale factors calculated were used for scaling.

2) Quantitative Comparison: CloudCompare was also used for this task. Two forms of comparison were made. In the first, two measurements were made of the dimple structure's diameter (measured using the anchor points present on the targets' top surfaces). In the second, a cloud-to-cloud comparison was performed. This involved first aligning the reconstructions with ground truths using the 8 achor points on each target. Then, for each target, the signed distance errors (closest point) between the ground truth and each reconstruction was made, for each point in the ground truth cloud. The mean and range of such errors for each point on the ground truth were determined, giving an idea of the accuracy and precision of the reconstructions, respectively.

IV. RESULTS AND DISCUSSION

A. Calibration

Graphs, showing the intrinsic and extrinsic parameters and overall mean reprojection error obtained using 10 to 100 image pairs, in increments of 10, are provided in Fig. 4 where the definition of each of the parameters can be found in [14]. The distortion parameters were not included due to the large number of them. All parameters, obtained with 100 image pairs, are tabulated in Table I. The overall mean reprojection error was 0.876 (SD = 0.589). It can be seen from the graphs that at 100 images, the parameters do tend to stabilize, indicating that more images may not have led to very different estimations.
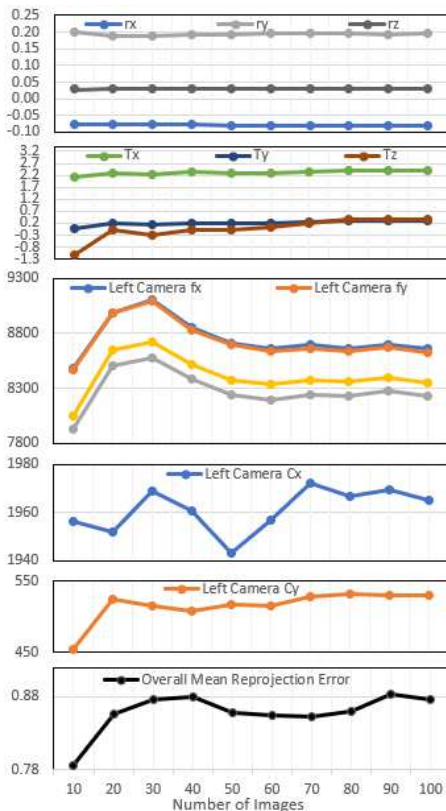

Fig. 4. Calibration parameters versus the number of image pairs used in calibration. The right camera's principal point is not shown since its value remained constant. See Table 1 for units.

Moreover, the error values seen in Table 1 represent a small fraction of the estimated values in the majority of cases. The reprojection error is large for this type of work, considering that the changes that occur to the ONH with POAG are subtle, and are likely to result in disparity changes on the order of single digit values. Potential sources of error include inaccurate checkerboard corner localization (possibly due to image quality at the edges), inaccurately fabricated checkerboard, and sub-optimal camera alignment. Estimates of the impact of camera resolution, that go beyond the scope of this paper, do not indicate image resolution as a significant source of error.

B. Reconstruction Performance

As an example, textured point clouds and color heightmaps of the ground truth and one reconstruction, for the glaucomatous ONH target, are shown in Fig. 5.

1) Measurements: The measurements made for each of the targets, on their reconstructions and ground truths, are tabulated in Table II. The values obtained indicate that the diameters measured on the reconstructions were close but slightly overestimated in both cases. There are a number of potential reasons for the discrepancy, including reconstruction error (as discussed below) and measurement error.

TABLE I.   SYSTEM CALIBRATION PARAMETERS.

Parameter	Intrinsic and Distortion Parameters	
	Left camera (value & ~3SD)	*Right camera (value & ~3SD)*
f_x	8655.2 pixels (47.0)	8229.2 pixels (45.6)
f_y	8624.3 pixels (47.0)	8352.9 pixels (46.6)
C_x	1965.4 pixels (10.4)	65.5 pixels (N/A)
C_y	530.3 pixels (12.8)	-41.7 pixels (N/A)
k_1	-1.819 (0.031)	-0.991 (0.018)
k_2	23.685 (1.013)	2.746 (0.124)
k_3	-132.572 (8.412)	-4.286 (0.340)
p_1	0.0123 (0.0014)	0.0123 (0.0005)
p_2	-0.00033 (0.00100)	0.04850 (0.00110)
	Extrinsic Parameters	
Parameter	Value & ~3SD	Parameter	Value & ~3SD
T_x	2.40660 mm (0.02080)	r_x	-0.08067 (0.00150)
T_y	0.29370 mm (0.01970)	r_y	0.19501 (0.00150)
T_z	0.35330 mm (0.10530)	r_z	0.02999 (0.00021)

TABLE II.   RECONSTRUCTION AND GROUND TRUTH MEASUREMENTS

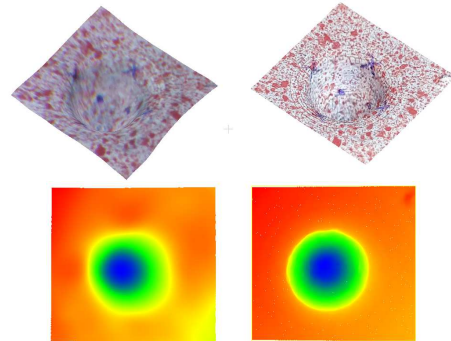| Target	Type	Measurement 1	Measurement 2
Normal	Reconstruction	Mean = 7.2 mm (SD = 0.1)	Mean = 7.1 mm (SD = 0.1)
	Ground Truth	7.0 mm	6.8 mm
Glaucomatous	Reconstruction	Mean = 6.0 mm (SD = 0.1)	Mean = 5.8 mm (SD = 0.1)
	Ground Truth	5.9 mm	5.7 mm

Fig. 5. Glaucomatous ONH target: point clouds (**TOP**) and color heightmaps (**BOTTOM**), of the ground truth (**RIGHT**) and one reconstruction (**LEFT**).
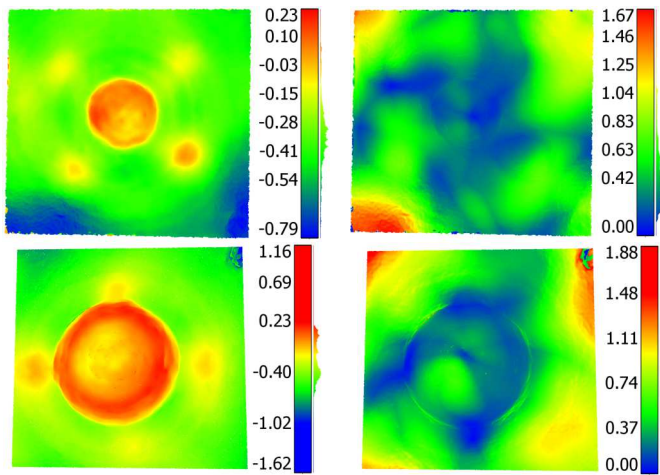
Fig. 6. Accuracy (mean distance errors; **LEFT COLUMN**) and precision (range of distance errors; **RIGHT COLUMN**) maps of the reconstructions obtained of the normal (**TOP ROW**) and glaucomatous (**BOTTOM ROW**) ONH targets. The colour scales, which are in units of millimetres, include histograms on their right-hand side to show the distribution of the errors.

2) Cloud-to-cloud comparison: Color error maps, showing the accuracy and precision of the reconstructions obtained for both targets, are shown in Fig. 6. It is apparent that areas far from the centre (i.e. the corners of the reconstructions) are the least accurate and precise. This provides some indication that radial distortion was not completely removed from the images. Moreover, a clear correlation appears to be present between the location of the anchor points and the pattern of the accuracy maps. This is likely an impact of the visual texture on the stereo matching algorithm, of a nature yet to be determined, since the only difference that these locations have compared to adjacent regions of the same depth are texture related. It also appears that, since the errors at the bottom of the dimple structrues are different to those of the peripheral regions, their depth is not accurate. This error likely results from either poor calibration or poor stereo matching performance. Since the reconstructions were performed at different angles of rotation, systematic errors (such as calibration errors) could have affected the precision maps obtained. Given the complex and subjective nature of early POAG detection, the required precision and accuracy for clinical usefulness is not clear. However, crude comparison against a device known to be clinically useful, and that is one of the poorer performing devices, can give an indication of the promise of the system used in this work. Indeed, when the errors in this work are scaled to what they would be in a life-sized eye, they are in general higher, but on the same order of magnitude as the depth measurement reproducibility error range seen in the HRT, 28.1 μm to 72.8 μm [15].

C. Limitations

The core limitations of this work relate to the ONH target appearance. Those used here differ from human ONHs which have large textureless regions, have semi-transparent layers, and are not usually rotationally symmetrical. Other limitations include the fact that illumination was provided from inside the eye, eliminating the challenges associated with corneal reflections. Moreover, the optics of the eye phantom used differ from that of a human eye.

V. CONCLUSION

This work assessed the accuracy and precision of performing 3D reconstruction of ONH targets in an eye phantom by exploiting computer stereo vision techniques to images obtained through a slit lamp fitted with cameras, an idea for which we demonstrated the initial feasibility of in previous work [12]. This study showed that while the reconstructions obtained resemble the ground truth data, the accuracy and precision of the system is likely not yet at the level needed to be clinically useful. Further work to improve calibration accuracy, image quality and stereo matching performance would be beneficial before moving to in-vivo work.

REFERENCES

[1] S. Kingman, 'Glaucoma is second leading cause of blindness globally.', *Bull. World Health Organ.*, vol. 82, no. 11, p. 887, Nov. 2004.
[2] Y.-C. Tham, X. Li, T. Y. Wong, H. A. Quigley, T. Aung, and C.-Y. Cheng, 'Global Prevalence of Glaucoma and Projections of Glaucoma Burden through 2040: A Systematic Review and Meta-Analysis', *Ophthalmology*, vol. 121, no. 11, pp. 2081–2090, Nov. 2014, doi: 10.1016/j.ophtha.2014.05.013.
[3] R. S. Harwerth, L. Carter-Dawson, F. Shen, E. L. Smith, and M. L. Crawford, 'Ganglion cell losses underlying visual field defects from experimental glaucoma', *Invest. Ophthalmol. Vis. Sci.*, vol. 40, no. 10, pp. 2242–2250, Sep. 1999.
[4] R. Susanna, C. G. De Moraes, G. A. Cioffi, and R. Ritch, 'Why Do People (Still) Go Blind from Glaucoma?', *Transl. Vis. Sci. Technol.*, vol. 4, no. 2, Mar. 2015, doi: 10.1167/tvst.4.2.1.
[5] J. A. Giaconi, S. K. Law, K. Nouri-Mahdavi, A. L. Coleman, and J. Caprioli, Eds., *Pearls of Glaucoma Management*, 2nd ed. Berlin Heidelberg: Springer-Verlag, 2016.
[6] N. J. Reus, M. de Graaf, and H. G. Lemij, 'Accuracy of GDx VCC, HRT I, and clinical assessment of stereoscopic optic nerve head photographs for diagnosing glaucoma', *Br. J. Ophthalmol.*, vol. 91, no. 3, pp. 313–318, Mar. 2007, doi: 10.1136/bjo.2006.096586.
[7] E. Corona, S. Mitra, M. Wilson, T. Krile, Y. H. Kwon, and P. Soliz, 'Digital stereo image analyzer for generating automated 3-D measures of optic disc deformation in glaucoma', *IEEE Trans. Med. Imaging*, vol. 21, no. 10, pp. 1244–1253, Oct. 2002, doi: 10.1109/TMI.2002.806293.
[8] T. Nakagawa *et al.*, 'Quantitative depth analysis of optic nerve head using stereo retinal fundus image pair', *J. Biomed. Opt.*, vol. 13, no. 6, p. 064026, Dec. 2008, doi: 10.1117/1.3041711.
[9] M. Bansal *et al.*, '3D optic disc reconstruction via a global fundus stereo algorithm', *Conf. Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. IEEE Eng. Med. Biol. Soc. Annu. Conf.*, vol. 2013, pp. 5877–5882, 2013, doi: 10.1109/EMBC.2013.6610889.
[10] J. Xu and O. Chutatape, 'Auto-adjusted 3-D optic disk viewing from low-resolution stereo fundus image', *Comput. Biol. Med.*, vol. 36, no. 9, pp. 921–940, Sep. 2006, doi: 10.1016/j.compbiomed.2005.05.001.
[11] L. Tang *et al.*, 'Robust Multiscale Stereo Matching from Fundus Images with Radiometric Differences', *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 11, pp. 2245–2258, Nov. 2011, doi: 10.1109/TPAMI.2011.69.
[12] I. Coghill, K. C. Jordan, R. A. Black, I. A. T. Livingstone, and M. E. Giardini, '3D Reconstruction of the Optic Nerve Head of a Phantom Eye from Images Obtained using a Slit Lamp Fitted with Low Cost Add-Ons', *Conf. Proc. Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. IEEE Eng. Med. Biol. Soc. Annu. Conf.*, vol. 2019, pp. 4717–4720, Jul. 2019, doi: 10.1109/EMBC.2019.8856654.
[13] J.-Y. Bouguet, 'Camera Calibration Toolbox for Matlab', 2015 http://www.vision.caltech.edu/bouguetj/calib_doc/index.html (Jun. 10, 2020).
[14] G. Bradski and A. Kaehler, *Learning OpenCV: Computer Vision with the OpenCV Library*. O'Reilly Media, Inc., 2008.
[15] A. W. Dreher, P. C. Tso, and R. N. Weinreb, 'Reproducibility of topographic measurements of the normal and glaucomatous optic nerve head with the laser tomographic scanner', *Am. J. Ophthalmol.*, vol. 111, no. 2, pp. 221–229, Feb. 1991, doi: 10.1016/s0002-9394(14)72263-9.