

Waterdrop Stereo

Paper ID: 0526




Figure 1: The pipeline of the proposed method. From left to right: Input image of water drops. Estimated 3D water drop shape. Estimated depth. Geometric Rectification.

Abstract

This paper introduces depth estimation from water drops. The key idea is that a single water drop adhered to window glass is totally transparent and convex, and thus optically acts like a fisheye lens. If we have more than one water drop in a single image, then through each of them we can see the environment with different view points, similar to stereo. To realize this idea, we need to rectify every water drop imagery to make radially distorted planar surfaces look flat. For this rectification, we consider two physical properties of water drops: (1) A static water drop has constant volume, and its geometric convex shape is determined by the balance between the tension force and gravity. This implies that the 3D geometric shape can be obtained by minimizing the overall potential energy, which is the sum of the tension energy and the gravitational potential energy. (2) The imagery inside a water-drop is determined by the water-drop 3D shape and total reflection at the boundary. This total reflection generates a dark band commonly observed in any adherent water drops. Hence, once the 3D shape of water drops are recovered, we can rectify the water drop images through backward raytracing. Subsequently, we can compute depth using stereo. In addition to depth estimation, we can also apply image refocusing. Experiments on real images and a quantitative evaluation show the effectiveness of our proposed method. To our best knowledge, never before have adherent water drops been used to estimate depth.

Keywords: Water drop imagery, single image surface construction, minimum energy surface, stereo

Concepts: •Computational Geometry and Object Modeling → Physically based modeling; •Enhancement → Geometric correction; •Scene Analysis → Stereo;

s

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). © 2016 Copyright held by the owner/author(s).

1 Introduction

Depth from real images is crucial information for many applications in computer graphics. A numerous methods have attempted to extract depth with various cues (e.g., shading [Ikeuchi and Horn 1981], multi-views [Hartley and Zisserman 2004], defocus [Favaro and Soatto 2005]). In contrast to all existing methods, in this paper, we explore a new possibility of using water drops adhered to window glass or a lens to estimate depth.

Water drops adhered to glass are transparent and convex, and thus each of them acts like a fisheye lens. As shown in Fig. 1.a, water drops' locations are normally scattered in various regions in an image. If we zoom in, each of the water drops displays the same environment from its own unique point of view. Due to the proximity to each other, some have similar visual content, but some can be relatively different, particularly when the water drops are apart in the image. Therefore, if we can rectify each of the water drops, we will have a set of images of the environment from relatively different perspectives, opening up the possibility of extracting the depth from the water drops, which is the goal of this paper.

To be able to achieve the goal, we need to rectify each water drop, so that planar surfaces look flat. Rectifying water drops, however, is problematic. In contrast to existing work in catadioptric imaging, which assumes the geometry of the sphere is known a priori, water drops shapes can vary in a considerable range and highly non-axial. To resolve this problem, we need to examine two physical properties of water drops. First, a static water drop has constant volume, and its geometric 3D shape is determined by the balance between the tension force and gravity. Because the water drop is in balance, it minimizes the overall potential energy, which is the sum of the tension energy and the gravitational potential energy. Based on this property, we introduce an iterative method to form the water-drop geometric shape. However, from a single 2D image, the volume cannot be directly obtained, since we do not know the thickness of the water drop. To solve this, we use the second physical property, i.e., water-drop appearance depends on their geometric shape and also the total reflection. The total reflection occurs near the water drop boundaries and triggers a dark band. We found that a water drop with a greater volume will have a wider dark band. Thus, we introduce a volume-varying-iteration framework that estimates the volume that best fit to the appearance. Having known the com-

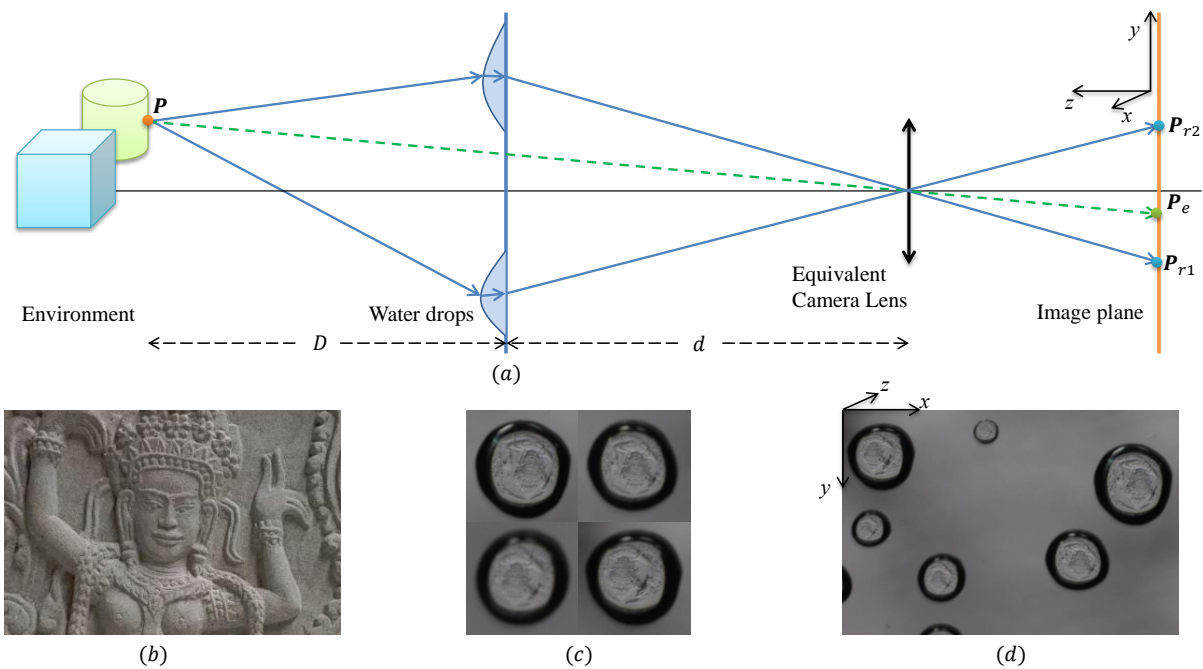


Figure 2: Image formation.

(a) The light path model assuming a pinhole camera. (b) Appearance of the environment when the camera focuses on the environment. (c) Appearance of water drops. (d) Image obtained by the camera.

plete 3D shape of water drops, we perform the multiple view stereo through backward raytracing and triangulation. Finally, we rectify the warped images. Having obtained the depth and rectified images, one of our applications is image refocusing. Figure 1 shows the general pipeline of our proposed method.

Contributions In this paper, we introduce a new way to recover depth using water drops from a single image. We also propose a novel method to reconstruct the 3D geometry of water drops by utilizing the minimum surface energy and total reflection. Aside from estimating depth, we also apply image refocusing through the information provided by water drops. Furthermore, the proposed non-parametric, non-axial algorithm can be generally applied to catadiptic imaging system.

The rest of the paper is organized as follows. Section 2 discusses related work in depth estimation, water modeling and shape from transparent objects. Section 3 explains the theory behind the water-drop physical properties. Section 4 introduces the methodology of the 3D shape estimation, stereo, as well as water-drop image rectification. Section 5 shows the three applications on stereo, image refocusing and image stitching. Section 6 shows the experimental results and evaluation. Section 7 concludes this paper.

2 Related Work

Three dimensional reconstruction of opaque objects from a single image have been explored for decades: Shape from shading [Ikeuchi and Horn 1981], shape from texture [Malik and Rosenholtz 1997], shape from defocus [Favaro and Soatto 2005] and piece-wise planarity [Horry et al. 1997]. A few approaches using silhouettes have been proposed to reconstruct a bounded smooth surface [Terzopoulos et al. 1988; Hassner and Basri 2006; Prasad and Fitzgibbon 2006; Joshi and Carr 2008; Oswald et al. 2012; Vicente and Agapito 2013]. The method of [Prasad and Fitzgibbon

2006] reconstructs a surface with minimum area, and [Oswald et al. 2012] proposes its speed-up version. However, none of these methods directly aim to model water or other transparent liquid from a single image.

Methods of [Garg and Nayar 2007; Roser and Geiger 2009; You et al. 2013; You et al. 2015] introduce airborne and adherent rain-drop modeling. Their goal is to detect and remove raindrops, and not to reconstruct 3D structure from raindrops. [Roser et al. 2010] exploits water drop surface fitting using B-splines and silhouettes using 1D splines. [Morris 2004; Tian and Narasimhan 2009; Or-eifej et al. 2011; Kanaev et al. 2012] exploit underwater imaging. They assume water surfaces are dynamic and dominated by transitions of waves, which do not suit to our specific problem.

Stereo and light field using perspective cameras with extra mirrors and lenses have also been explored. [Baker and Nayar 1999; Taguchi et al. 2010] propose algorithms using sphere mirrors. [Levoy et al. 2004] introduces arrays of planar mirrors. [Swaminathan et al. 2001; Ramalingam et al. 2006] address the use of axial cameras, and [Taguchi et al. 2010] extends the work to axial-camera arrays. Later, [Agrawal and Ramalingam 2013] proposes methods to automatically calibrate the system. All of these methods, however, assume radial or planar symmetry of the media (mirror/lens), which are not satisfied in the case of water drops, since water drops are highly non-axial.

3 Modeling

Theoretical background and modeling of water drops are discussed in this section. We first explain briefly the image formation, showing the correlations between the environment, water-drops and the camera. Subsequently, we model the raindrop 3D geometry, particularly the concept of minimum energy surface. Based on the image formation and the raindrop geometry, we study the total reflection inside water drops, which is necessary to determine the water-




Figure 3: Parameters of a water drop.

(a) In the global coordinates, the geometric shape is determined by the gravity. (b) The parameters in the camera coordinates. Point A is a two phase point (water-air), where the tensor force balances the pressure. Point B is a three phase point (water-air-material).

drop's volume. All these aim at water-drop image rectification.

3.1 Image Formation

Fig. 2 illustrates our image formation. Rays reflected from the environment pass through two water drops before hitting the image plane. Unlike in the conventional image formation, the passing rays are refracted by water drops, where each water drop acts like a fish-eye lens that warps the images. Assuming we have a few water drops that are apart to each other, the imageries of the water drops will be slightly different to each other, even though the environment is identical, as shown in Fig. 2.

From the illustration in Fig. 2, we can conclude that the image captured by the camera through water drops is determined by three interrelated factors: (1) the depth of the environment, (2) the 3D shape of water drops, which determine how light rays emitted from the environment are refracted and, (3) camera intrinsic parameters, which are assumed to be known. Therefore, to be able to recover the depth of the environment, we need to obtain the 3D shape of water drops.

3.2 Minimum Energy Surface

For the purpose of exploring the minimum energy surface to estimate the 3D shape of a water drop, we introduce a local coordinate system of the camera, which is illustrated in Fig. 2.a and d. In the coordinates, water drop 3D shape can be parameterized as:

$$\mathcal{S} = \{z(x, y), (x, y) \in \Omega_R\}, \quad (1)$$

where Ω_R indicates the raindrop area attached to glass. (x, y) is any point in the raindrop area and z is the height.

A static water drop has a constant volume, and its 3D shape $\tilde{\mathcal{S}}$ minimizes the overall potential energy E , which can be written as:

$$\begin{aligned} E(\tilde{\mathcal{S}}) &= \min E(\mathcal{S}) = \min(E_T(\mathcal{S}) + E_G(\mathcal{S})), \\ V(\mathcal{S}) &= \text{constant}, \end{aligned} \quad (2)$$

where E_T is the tension energy, and E_G is the gravitational potential energy, V is the volume. Therefore, to solve the geometry of a raindrop, we need to find the surface $\tilde{\mathcal{S}}$.

Figure 3 illustrates the 3D shape of a water drop. Point A is a two-phase (water-air) balanced point, where surface tension T balances pressure P . Point B is a three phase point (water-air-material), where the tension is from both water T_w and adhesion surface T_m . These two types of tension balance the gravity, G .

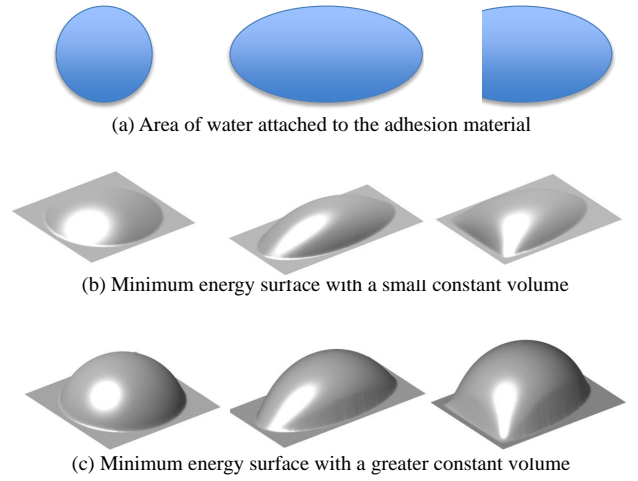


Figure 4: Minimum energy surfaces given the area and volume. (b) The minimum energy surface when the volume coefficient $\alpha = 0.10$, defined in Eq. (16). (c) $\alpha = 0.35$. Assuming the gravity is along z axis.

With the parameterized surface, we can write the surface tension energy as:

$$E_T(\mathcal{S}) = \int_{\Omega_R} \sigma dA = \int_{\Omega_R} \sigma \sqrt{1 + |\nabla z|^2} dx dy, \quad (3)$$

where σ is the surface tension index for water, dA denotes a unit surface area and ∇ is the gradient [Feynman et al. 2013]. As we can see, the tension energy is proportional to the area of the surface.

The gravitational potential energy can be expressed as:

$$E_G(\mathcal{S}) = \int_{\Omega_R} dx dy \int_0^z (x \cos \theta_x + y \cos \theta_y + z \cos \theta_z) g \rho dw, \quad (4)$$

where θ_x , θ_y and θ_z denote the angles between the x , y , z coordinates and the gravity correspondingly. g is the gravity and ρ is the density of water, which are generally known. Moreover, we can add a constraint that:

$$\int_{\Omega_R} z dx dy \equiv V. \quad (5)$$

Therefore, the parameterized surface \mathcal{S} is estimated by minimizing the overall potential energy determined by Eq. (2), (3) and (4) with the constraints of constant volume in Eq. (5). Figure 4 shows some examples of the surface estimated by using the technique. We will discuss the algorithm in detail in Section 4.

Note that, to uniquely determine the geometry of a water drop, we need to know both the 2D area where the water drop attached to glass, Ω_R , and the volume V . While the former can be directly inferred from the image, the latter is not straightforward to obtain. The subsequent section will discuss how we can possibly determine the volume.

3.3 Water-Drop Volume from Dark Band

As we can see in Fig. 2.c, the basic idea of our volume estimation is based on the dark band at the boundary of a water drop. We found that the wider the dark band the larger the volume of the water. This section discusses this idea further.




Figure 5: Refraction by a water drop.

(a) A ray coming from the environment is refracted twice before reaching the camera C . (b) Backward ray tracing, where a virtual ray emitted from the camera passes through the same path as in (a). For simplification, we remove the refraction between flat surfaces by moving the camera position to C' . (c) When θ_w is greater than the critical angle, light will not be transmitted but reflected inside. (d) Two polarized components, R_s and R_p , of the incidence ray.

Refraction model Fig. 5.a illustrates a ray coming from the environment is refracted twice before reaching the camera. Since, we are only interested in the rays that can reach the camera, we can use backward raytracing to know the paths of the rays. Moreover, we assume that the glass is so thin that we can ignore the refraction due to the glass. Because we are mostly interested in the refraction on the curved surface, to further simplify the model, we remove the refraction between the flat surface by moving the camera from position C to C' , as shown in Fig. 5.b. For approximation, when the incident angle is small, we can consider that the perpendicular distance from the camera to the refraction plane, denoted as C_z , is changed to $C'_z = \frac{n_w}{n_a} C_z$, where n_w and n_a are the refractive indices of water and air, respectively. Detailed derivation of the position of C' is discussed in Appendix A.

Dark Band and Total Reflection The dark band at the boundary of a water drop is caused by light coming from the environment reflected back inside the water, instead of being transmitted to the camera. This phenomenon is known as the total reflection, and applies to all light rays whose relative angles to the water's surface normal are larger than the critical angle, denoted as θ_w .

To analyze the correlation between the critical angle with the water-drop 3D shape $S : z(x, y)$, we refer to Snell's law, which indicates the critical angle:

$$\tilde{\theta}_w = \sin^{-1} \frac{n_a}{n_w}. \quad (6)$$

As indicated in Fig. 5.c, we denote the surface normal as \mathbf{N} , which can be derived from z as: $\mathbf{N} = (N_x, N_y, N_z)^\top = \frac{\mathbf{N}'}{\|\mathbf{N}'\|}$ where, $\mathbf{N}' = (\frac{\partial z}{\partial x}, \frac{\partial z}{\partial y}, 1)^\top$, and $\|\cdot\|$ denotes the ℓ_2 norm.

The angle between the surface normal and the z -axis denoted as θ_N is the sum of the incidence angle of water, θ_w , and the angle between the incidence ray and z -axis $\theta_{C'}$:

$$\theta_N = \theta_w + \theta_{C'}. \quad (7)$$

where $\theta_{C'}$ is determined by the position of the camera and the position of the refraction. Considering the z component of the normal N_z , also defined as: $N_z = \cos \theta_N$, we know that when $N_z \leq \cos(\theta_w + \theta_{C'})$, the corresponding water drop area is totally dark. For instance, when $\theta_{C'}$ is 0, and $\frac{n_a}{n_w}$ is approximately $\frac{3}{4}$, we




Figure 6: Dark bands.

(a) Water drop 3D geometry. (b) The z component of surface normal. (c) The dark bands. Second row: A greater water drop volume, wider dark band.

have

$$N_z \leq \tilde{N}_z \approx 0.661. \quad (8)$$

Where \tilde{N}_z is the denotation for the critical value. Figure 6 shows some examples of synthetically generated dark bands. As we can observe, a greater volume of the water drop indicates a wider dark band. Therefore, to infer the water drop volume from the dark band is possible.

Dark Band and Fresnel Equation While the dark band can be theoretically inferred from the water-drop geometry, detecting them from an image is nontrivial. Due to the sensor noise and the leak of light¹, dark bands are not totally dark. Moreover, there are textures in the environment that can be darker than dark bands. To

¹Back light from camera side which is reflected by the glass plate and goes into the camera; and the interreflection inside a water drop.

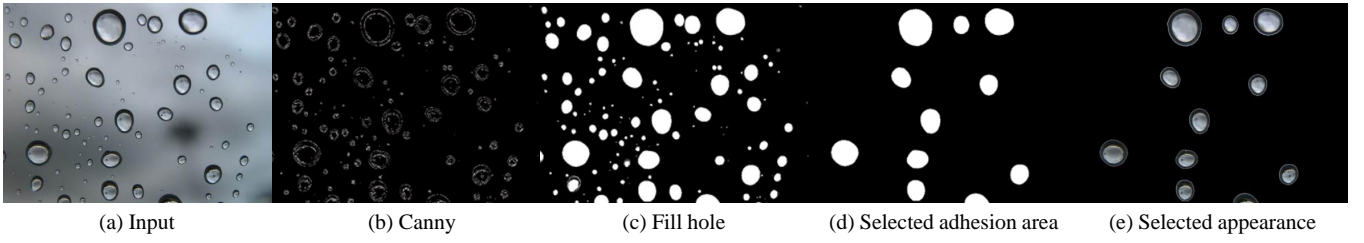


Figure 7: Selecting water drops from a single image.




Figure 8: Iteration of water drop 3D shape with a fixed volume.

resolve the problem, we employ the Fresnel equation and formulate the brightness values near the critical angle.

The refraction coefficients, denoted as \mathcal{T}_s and \mathcal{T}_p , for two orthogonal polarized components for the light rays traveling from air to water are written as:

$$\mathcal{T}_s = 1 - \left(\frac{\sin(\theta_w - \theta_a)}{\sin(\theta_w + \theta_a)} \right)^2 \quad (9)$$

$$\mathcal{T}_p = 1 - \left(\frac{\tan(\theta_w - \theta_a)}{\tan(\theta_w + \theta_a)} \right)^2 \quad (10)$$

where θ_w and θ_a are depicted in Fig. 5.d. In our case, we assume the light from the environment is not polarized, and thus the overall refraction coefficient is

$$\mathcal{T} = \frac{1}{2}(\mathcal{T}_s + \mathcal{T}_p). \quad (11)$$

Concerning the dark bands, we are interested in two critical conditions. First, when the incidence angle θ_a is close to 0. In such a condition, $\sin \theta_a \approx \theta_a$, $\cos \theta_a \approx 1$, and consequently:

$$\mathcal{T} = \frac{4n_a n_w}{(n_w + n_a)^2}. \quad (12)$$

Substituting the value for water gives us $\frac{n_a}{n_w} = \frac{3}{4}$, and thus we have $\mathcal{T} \approx 0.980$.

Second, when incidence angle θ_a is close to $\frac{\pi}{2}$, (the locations near the dark band). Hence, $\sin \theta_a \approx 1$, $\cos \theta_a \approx \frac{\pi}{2} - \theta_a$, as a result:

$$\mathcal{T} = 2\sqrt{1 - \left(\frac{n_a}{n_w}\right)^2} \left(\frac{n_a}{n_w} + \frac{n_w}{n_a}\right) \left(\frac{\pi}{2} - \theta_a\right) \quad (13)$$

Similar to the first condition, substituting the value $\frac{n_a}{n_w} = \frac{3}{4}$, we obtain:

$$\mathcal{T} \approx 2.76\left(\frac{\pi}{2} - \theta_a\right). \quad (14)$$

Considering θ_a is connected with θ_w by Snell's law, and the connection between θ_N and θ_w (Eq. (7)), we establish the connection

between the z component of surface normal $N_z = \cos \theta_N$ and the refraction coefficient:

$$\mathcal{T} \approx 7.68(N_z - \tilde{N}_z). \quad (15)$$

This relation between the image brightness and surface normals give us a constraint of a contrast between the dark band and other parts inside the waterdrop image. Using this, the water-drop volume can be inferred using the brightness close to the dark band region. Details are provided in Sec. 4.1

4 Methodology

In this section, the detailed algorithm for rectifying images of water drops and estimating depth is introduced. As illustrated in Fig. 1, it has three main steps: (1) water drop 3D shape reconstruction by minimizing energy surface, (2) Multi-view stereo and (3) Image rectification.

Water Drop Detection Water-drops appearance is highly dependent on the environment, and thus detecting them is not trivial. Fortunately, in our case, we can assume water drops are in focus, and thus the environment image is rather blurred. Hence, we can utilize edge detection to locate water drops, as illustrated in Fig. 7. Having located water drops, we select those that are sufficiently large (e.g., the diameter is greater than 300 pixels). This is to ensure that rectified images are not too small.

4.1 Water Drop 3D Shape Reconstruction

Mesh representation and Initialization To reconstruct the 3D shape of water drops, we first represent the water surface using a parameterized mesh. Referring to Eq. (1), we can describe a surface as: $\mathcal{S} = \{z(i, j), (i, j) \in \Omega_R\}$, where (i, j) are the location of a pixel in the water drop area. Accordingly, the area of Ω_R is defined as: $B = \sum_{(i, j) \in \Omega_R} 1$, where 1 is the unit for a pixel's area.

At this initialization, we do not know the volume of the water drop, and make an initial guess based on:

$$V = \alpha B^{\frac{3}{2}}, \quad (16)$$

where α is the volume coefficient and set to 0.30 as default. Based on the equation, with α fixed, when the area B increases in square rate, the volume will increase in cubic rate. This means when performing scale change for the water drop surface, α remains the same value. Figure 4 gives some examples how α is related to the reconstructed surface.

We initialize the mesh as a cylinder by defining:

$$z(i, j) = \alpha B^{\frac{1}{2}}, (i, j) \in \Omega_R. \quad (17)$$

Figure 8 shows an example of the initial surface.

Iteration with fixed volume We solve the constrained minimum energy surface using the iterative gradient descent. For iteration t we update the mesh in three steps: tensor energy update, gravity update, and volume update. This strategy is an extension of the smooth surface reconstruction proposed by [Oswald et al. 2012].

- Step 1: Tension energy update. It attempts to construct the surface as smooth as possible:

$$z_{t+1} = z_t - \tau \sigma \cdot \frac{dE(\mathcal{S})}{dz_t}, \quad (18)$$

where τ controls the update speed with $\tau = 0.5$ as default, σ is the tension coefficient in physics. We define:

$$\frac{dE(\mathcal{S})}{dz_t} = -\text{div} \left(\frac{1}{\sqrt{1 + |\nabla z_t|^2}} \nabla z_t \right), \quad (19)$$

where div is the divergence.

In our settings, a water drop has size around $3mm$ or approximately 500 in pixel, and thus the size of a unit pixel is about $6\mu m$. Tension coefficient for water in room temperature is $73,000N/m$.

- Step 2: Gravity update. It intends to increase the height for the mesh points that lower the potential energy:

$$z_{t+1}(i, j) = z_t(i, j) - \tau \rho g ((y_g - i) \cos \theta_y + (x_g - j) \cos \theta_x), \quad (20)$$

where (x_g, y_g) is the geometry centroid of the water drop and defined as:


$$x_g = \frac{1}{B} \sum_{(i,j)} z(i, j) \cdot j, \quad y_g = \frac{1}{B} \sum_{(i,j)} z(i, j) \cdot i. \quad (21)$$

Substituting $\rho = 1Kg/L$ and $g = 9.8m/s^2$, we found when the adherent surface tilt is small, the waterdrop geometry is mainly dominated by the tension energy.


- Step 3: Volume update. Having updated the tension and gravity in the previous two steps, this step checks the current volume and compares it with the targeted volume V , and then re-adjusts the volume by adding the same value to all the mesh points:

$$z_{t+1} = z_t + \left(\frac{V - \sum_{(i,j) \in \Omega_R} z_t(i, j)}{B} \right) \quad (22)$$

After each iteration, we check the absolute change of volume: $\sum_{(i,j)} |z_{t+1}(i, j) - z_t(i, j)|$, and set the convergence threshold to $1e - 6V$ as default, where V is the targeted volume. We run the iterations up to 4000 times. Figure 8 shows the progress of the estimated volume.



(a) Estimated geometry with varying volume



(b) Estimated band of critical angle

Figure 9: Estimating the width of a dark band

Left: Underestimated volume where the darkband (shown in blue) is mostly covering the dark pixels. Middle: Correctly estimated volume, where the darkband cover about the same proportion of dark and bright pixels. Right: Overestimated volume, where the darkband is mostly covering bright pixels.

Iteration with varying volume Having estimated the surface with fixed volume, we can obtain the surface normals and evaluate the brightness values near the dark band and gradually adjust the volume.

In Section 3.3, we have built the relation between the surface normal and the luminance (Eq. (15)). According to Eq. (8), when the z component of surface normal is smaller than the critical value N_Z , the pixel brightness should be close to 0. Yet, when z is slightly greater than N_Z , the pixel brightness should follow Eq. (15). Thus, if we compute the average refraction coefficient, denoted as T_r , for pixels whose normals are within the range $N_z = N_Z \pm \tilde{N}_z$, we can do local linear expansion of Eq. (14) and obtain:

$$T_r = 3.84 \tilde{N}_z. \quad (23)$$

Specifically, we set $\tilde{N}_z = 0.02\pi$, and thus $T_r = 0.241$. Consequently, the average brightness of the band, I_r , is:

$$I_r = 0.241 I_b, \quad (24)$$

where I_b is the average brightness of the non water-drop areas.

In Fig. 9, we sample the brightness of the estimated band. As shown in Fig. 9.a, when the volume is underestimated, the dark band is wider than the real one, resulting in less bright pixels. On the contrary, when the volume is overestimated, the dark band is narrower than the real one, resulting in brighter pixels. With the above analysis, we update the volume every 400 iteration (as default) for the fixed volume algorithm introduced previously:

$$V_{t+1} = V_t + \tau_r \cdot V_t \cdot \left(1 - \frac{I_t}{I_r} \right), \quad (25)$$

where I_t is the sampled brightness, I_r is the targeted brightness value, and τ_r is a weighting coefficient which controls the updating speed and is set to 0.5 as default. We demonstrate the accuracy of the estimation by experiments in Sec. 5.1.

4.2 Waterdrop Stereo

Once the geometry of each water drop is obtained, we perform multi-view stereo to estimate depth. Unlike multiple view stereo




Figure 10: *Illuminance compensation of water drop images.*




Figure 11: *Ray-tracing based triangulation.*

For a set of corresponding points \mathbf{x}_j , the backward raytracing finds the set of rays \mathbf{R}_{oj} in the space. The position \mathbf{p} of the point in space is where the sum of Euclidean distance to all the rays is minimized.

based on perspective or radial catadioptric cameras, where the projection of each camera can be modeled using a few parameters, unfortunately the water drops are non-parametric and non-axial. To overcome this problem, we propose a raytracing based triangulation. (More detail about the implementation could be found in Appendix B.)

As illustrated in Fig. 11, a point $\bar{\mathbf{x}}_1$ has its corresponding points on other water drops, denoted as $\bar{\mathbf{x}}_j$, where j is the index of water drops. By knowing the geometry of each water drop, we can find the location where the refraction happens in each water drop, denoted as \mathbf{x}_j . Specifically, to obtain the location of \mathbf{x} , we need: (1) The position of the camera \mathbf{C} , which is known a priori, and thus can be converted to equivalent position \mathbf{C}' , as illustrated in Fig. 5.b. (2) The water drop geometry, which is already estimated.

At each refraction location, the incident angle is obtained using $\mathbf{R}_i = \frac{\mathbf{x} - \mathbf{C}'}{\|\mathbf{x} - \mathbf{C}'\|}$. The surface normal \mathbf{N} is known through geometry estimation. Through Snell's law, we can obtain the outbound

angle \mathbf{R}_o . The outbound ray is formulated as:

$$\mathbf{x} + \alpha \mathbf{R}_o, \alpha \in \mathbb{R}. \quad (26)$$

Hence, now we can perform the classical triangulation as illustrated in Fig. 11. Given a set of corresponding points on each waterdrop \mathbf{x}_j , we could obtain its outbound ray of fraction $\mathbf{x}_j + \alpha \mathbf{R}_{oj}, \alpha \in \mathbb{R}, j =$ the index of water drops. The triangulation aim to find the position of point \mathbf{p} , which minimizes the Euclidean distance to all the rays:

$$\mathbf{p} = \left[\sum_j (\mathbf{I} - \mathbf{R}_{oj} \mathbf{R}_{oj}^T) \right]^{-1} \left[\sum_j (\mathbf{I} - \mathbf{R}_{oj} \mathbf{R}_{oj}^T) \mathbf{x}_j \right]. \quad (27)$$

The detailed derivation refers to [Szeliski 2010]. The depth of each point \mathbf{p} is its z component. Figure 14 is an example of the depth map in water drop image.

4.3 Rectification of Water Drop Image

Having estimated the depth map on each water drop, we unwarped the distorted water drop image. Referring to the pinhole camera model (Fig. 2), for each water drop image, a space point P_r with projection at P_{r-i} is projected to P_e . Figure 12 and Figure 13 shows results of the rectified water drop images.

According to Eq. (9) and Eq. (10), with the water drop geometry obtained, we can compensate the brightness values according to the refractive coefficient \mathcal{T} . Figure 10 shows an example of the brightness compensation.

5 Experiments and Analysis

We conduct experiments using both synthetic data and real data to examine and analyze the performance of our method. In the experiment, we evaluate the estimated 3D shape of the water drops, and the depth estimation. Without loss of generality, our method can also be used to handle axial mirror/lens models [Taguchi et al. 2010].

5.1 3D Shape Reconstruction and Image Rectification

To evaluate the accuracy of the 3D shape of water drops, we utilize synthetic data. We cannot use real data, since automatic 3D acquisition systems, such as a laser range finder, cannot be used to estimate the 3D of water. We use real images for the evaluation of the image rectification. Some of the real images are taken by ourselves and some are downloaded from the Internet.

Figure 12 shows the generated synthetic water drops with a variety of boundaries. A quantitative evaluation is performed by comparing our estimation with the ground truth 3D shape. The error is




Figure 12: Quantitative evaluation of water surface reconstruction and rectification.

normalized to the percentage of the scale of the water drops. As one can observe, the reconstruction error is less than 3% even for the most irregular water drops.

Figure 13 shows a collection of the rectified water drop images from real data. The input image is cropped for better visualization, yet the camera center is not at the cropped image center.

Without loss of generality, our proposed method can also be used for axial models, which is considered as a specific case when the water drop is exactly radially symmetric. The last two rows of Fig. 13 show the results on spherical mirrors [Taguchi et al. 2010]. Because dark band estimation is not applicable on mirrors, we specify the volume parameter $\alpha = 0.40$ for the mirrors. And refraction is changed to reflection in raytracing.

We implemented our method in Matlab and measured the computational time without parallelization. For water 3D shape estimation, the time varies depending on the water drop volume and the mesh resolution. Table 1 shows the computation time of varying volume and fixed mesh resolution. And Table 2 shows the time of varying mesh resolution. At typical case, the resolution of mesh is set to 200×200 and the reconstruction time is about 10s. Note that, because each of the water drop reconstruction are performed separately, we can simply parallelize each of the tasks. Thus, the overall computation time does not increase with the number of water drops.

Table 1: Computation time for water drop 3D reconstruction with varying volume.

Volume (α)	0.05	0.1	0.2	0.3
Iterations	1300	1600	2600	3200
Time (s)	5.5	7.1	11.7	15.1

The mesh resolution is fixed to 200×200 .

Table 2: Computation time for water drop 3D reconstruction with varying mesh resolution.

Mesh res.	50*50	100*100	200*200	400*400
Iterations	400	800	2600	5100
Time (s)	0.4	1.7	11.7	241

The volume is fixed to $\alpha = 0.2$.

5.2 Depth Estimation

We use both synthetic and real water drop data to demonstrate our stereo method. Furthermore, Our non-parametric, non-axial method could be applied to axial-mirror/lens model as well.

Figure 14 shows the generated synthetic data from the Middlebury




Figure 13: *Rectification of real water images.*

The first two row are water-drop images taken by ourselves. The first three columns in Row 3 and 4 are downloaded from the Internet. And the last three images are taken by sperical mirrors by [Taguchi et al. 2010]. There is slant between the background and the water drop in some data, however the rectified image is not necessary to be rectangles.

data set. As can be seen, the depth estimation result highly resembles the ground truth with only errors occur at object's boundaries.

The result on the real waterdrop images are shown in the first 4 rows of Fig. 15. For the first row, the data is taken using a micro-lens, and the resolution for each waterdrop is more than 600 pixels. The second and third data is taken by normal commercial lens, the resolution for each waterdrop is about 200-300 pixels. As shown, our proposed method can generally recover the depth structure. However, we find the bottleneck of our method is in finding the corresponding points between water drop images. Since, when the camera is zoomed-in to focus on the details of water drops, the sensor noise and the dust on the plate is no longer negligible, which adversely affect the accuracy and stability of both sparse and dense corresponding methods.

As mentioned previously, our non-parametric, non-axial method can be applied to axial-mirror/lens model. Fig. 16 shows our stereo estimation results. Since the image has sufficiently high resolution and less noise, the dense matching is significantly stable, and consequently it enables us to estimate the depth more accurately.

6 Discussion and Conclusion

In this paper, we had exploited the depth reconstruction from water drops. In our pipeline, there are three key steps: the water-drop 3D

shape reconstruction, depth estimation using stereo, and water-drop image rectification. All of these are done using a single image. We evaluated our method, and it shows that the method works effectively for both synthetic and real images. Nevertheless, there are still some limitations in it. One of the limitation is the common perspective camera and lens cannot obtain high resolution image of water drops. Which degraded the overall performance of the depth estimation, specifically, the sparse/dense correspondence quality is degraded because of the low-resolution images. For future works, we are considering improving the image quality. Furthermore, we will consider simultaneous estimation of the waterdrop geometry and depth.

References

- AGRAWAL, A., AND RAMALINGAM, S. 2013. Single image calibration of multi-axial imaging systems. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, IEEE, 1399–1406.
- BAKER, S., AND NAYAR, S. K. 1999. A theory of single-viewpoint catadioptric image formation. *International Journal of Computer Vision* 35, 2, 175–196.
- FAVARO, P., AND SOATTO, S. 2005. A geometric approach to shape from defocus. *Pattern Analysis and Machine Intelligence*,

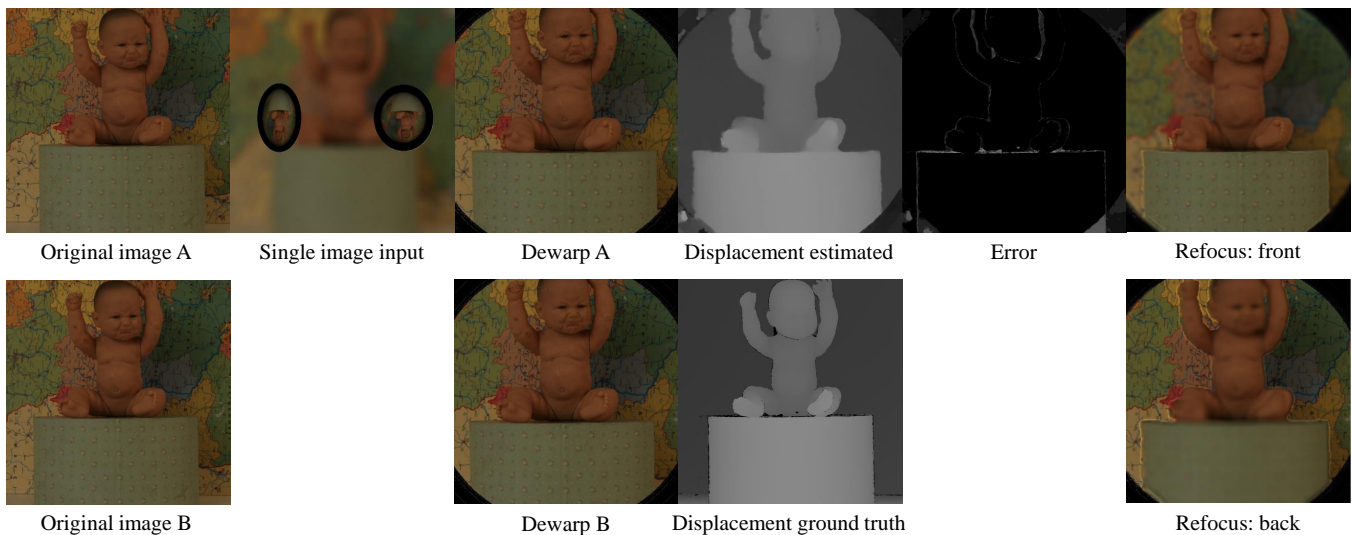


Figure 14: Stereo using two dewarped water drop images.




Figure 16: Stereo using axial mirror images.

Our non-parametric, non-axial method can be applied to axial data as well. The first column is the result using the original data provided by [Taguchi et al. 2010], the second and third column are based on low-resolution data, because the original data is not provided.

IEEE Transactions on 27, 3, 406–417.

FEYNMAN, R. P., LEIGHTON, R. B., AND SANDS, M. 2013. *The Feynman Lectures on Physics, Desktop Edition Volume I*, vol. 1. Basic Books.

GARG, K., AND NAYAR, S. 2007. Vision and rain. *International Journal of Computer Vision* 75, 1, 3–27.

HARTLEY, R. I., AND ZISSERMAN, A. 2004. *Multiple View Geometry in Computer Vision*, second ed. Cambridge University Press, ISBN: 0521540518.

HASSNER, T., AND BASRI, R. 2006. Example based 3d reconstruction from single 2d images. In *Computer Vision and Pattern Recognition Workshop, 2006. CVPRW'06. Conference on*, IEEE, 15–15.

HORRY, Y., ANJYO, K.-I., AND ARAI, K. 1997. Tour into the picture: using a spidery mesh interface to make animation from a single image. In *Proceedings of the 24th annual con-*

ference on Computer graphics and interactive techniques, ACM Press/Addison-Wesley Publishing Co., 225–232.

IKEUCHI, K., AND HORN, B. K. 1981. Numerical shape from shading and occluding boundaries. *Artificial intelligence* 17, 1, 141–184.

JOSHI, P., AND CARR, N. A. 2008. Repoussé: automatic inflation of 2d artwork. In *Proceedings of the Fifth Eurographics conference on Sketch-Based Interfaces and Modeling*, Eurographics Association, 49–55.

KANAEV, A. V., HOU, W., WOODS, S., AND SMITH, L. N. 2012. Restoration of turbulence degraded underwater images. *Optical Engineering* 51, 5, 057007–1.

LEVOY, M., CHEN, B., VAISH, V., HOROWITZ, M., MCDOWALL, I., AND BOLAS, M. 2004. Synthetic aperture confocal imaging. In *ACM Transactions on Graphics (TOG)*, vol. 23, ACM, 825–834.

MALIK, J., AND ROSENHOLTZ, R. 1997. Computing local surface orientation and shape from texture for curved surfaces. *International journal of computer vision* 23, 2, 149–168.

MORRIS, N. J. W. 2004. *Image-based water surface reconstruction with refractive stereo*. PhD thesis, University of Toronto.

OREIFEJ, O., SHU, G., PACE, T., AND SHAH, M. 2011. A two-stage reconstruction approach for seeing through water. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, IEEE, 1153–1160.

OSWALD, M. R., TOPPE, E., AND CREMERS, D. 2012. Fast and globally optimal single view reconstruction of curved objects. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, IEEE, 534–541.

PRASAD, M., AND FITZGIBBON, A. 2006. Single view reconstruction of curved surfaces. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, vol. 2, IEEE, 1345–1354.

RAMALINGAM, S., STURM, P., AND LODHA, S. K. 2006. Theory and calibration for axial cameras. In *Computer Vision-ACCV 2006*. Springer, 704–713.




Figure 15: Stereo using real waterdrop images.

The stereo uses 2-4 waterdrops, however only two of the dewarp results are shown.

ROSER, M., AND GEIGER, A. 2009. Video-based raindrop detection for improved image registration. *IEEE 12th International Conference on Computer Vision Workshops (ICCV Workshops)*.

ROSER, M., KURZ, J., AND GEIGER, A. 2010. Realistic modeling of water droplets for monocular adherent raindrop recognition using bezier curves. *Asian Conference on Computer Vision*.

SWAMINATHAN, R., GROSSBERG, M. D., AND NAYAR, S. K. 2001. Caustics of catadioptric cameras. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, vol. 2, IEEE, 2-9.

SZELISKI, R. 2010. *Computer vision: algorithms and applications*. Springer Science & Business Media.

TAGUCHI, Y., AGRAWAL, A., VEERARAGHAVAN, A., RAMALINGAM, S., AND RASKAR, R. 2010. Axial-cones: modeling spherical catadioptric cameras for wide-angle light field rendering. *ACM Transactions on Graphics-TOG* 29, 6, 172.

TERZOPOULOS, D., WITKIN, A., AND KASS, M. 1988. Symmetry-seeking models and 3d object reconstruction. *International Journal of Computer Vision* 1, 3, 211-221.

TIAN, Y., AND NARASIMHAN, S. G. 2009. Seeing through water: Image restoration using model-based tracking. In *Computer Vision, 2009 IEEE 12th International Conference on*, IEEE, 2303-2310.

VICENTE, S., AND AGAPITO, L. 2013. Balloon shapes: reconstructing and deforming objects with volume from images. In *3D Vision-3DV 2013, 2013 International Conference on*, IEEE, 223-230.

XU, L., JIA, J., AND MATSUSHITA, Y. 2012. Motion detail preserving optical flow estimation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 34, 9, 1744-1757.

YOU, S., TAN, R. T., KAWAKAMI, R., AND IKEUCHI, K. 2013. Adherent raindrop detection and removal in video. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*.

YOU, S., TAN, R., KAWAKAMI, R., MUKAIGAWA, Y., AND IKEUCHI, K. 2015. Adherent raindrop modeling, detection and removal in video. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*.

ZORICH, V., AND COOKE, R. 2004. *Mathematical analysis*. Springer.

Appendix A: Equivalent Camera Position for Flat Air-Water Refraction

Here we prove that we can further remove the refraction from water to air in Fig. 5.a by moving the camera to its approximated equivalent place at C' , as illustrated in Fig. 5.b.

We assume the camera position is $C = (0, 0, z_c)^\top$ and the flat plate is $z = 0$. Given a point on the plate $x = (x, y, 0)^\top$ where the refraction happens, the orientation of the incidence angle is:

$$\mathbf{R}_i = \frac{\mathbf{x} - \mathbf{C}}{\|\mathbf{x} - \mathbf{C}\|} = \frac{(x, y, -z_c)^\top}{\sqrt{x^2 + y^2 + z_c^2}} \quad (28)$$

The parallel and orthogonal components to the surface normal $N = (0, 0, 1)^\top$ are:

$$\begin{aligned} \mathbf{R}_{i\parallel} &= (\mathbf{R}_i^\top \mathbf{N}) \mathbf{N} = \frac{(0, 0, -z_c)^\top}{\sqrt{x^2 + y^2 + z_c^2}} \\ \mathbf{R}_{i\perp} &= \mathbf{R}_i - \mathbf{R}_{i\parallel} = \frac{(x, y, 0)^\top}{\sqrt{x^2 + y^2 + z_c^2}}. \end{aligned} \quad (29)$$

The orientation of angle of refraction, denoted as \mathbf{R}_o , can be obtained according to Snell's law:

$$\mathbf{R}_{o\perp} = \frac{n_a}{n_w} \mathbf{R}_{i\perp} = \frac{(x, y, 0)^\top}{\frac{n_w}{n_a} \sqrt{x^2 + y^2 + z_c^2}} \quad (30)$$

Considering \mathbf{R}_o is normalized: $\|\mathbf{R}_o\| = \|\mathbf{R}_{o\perp} + \mathbf{R}_{o\parallel}\| = 1$, and

thus:

$$\mathbf{R}_o = \frac{\mathbf{R}'_o}{\|\mathbf{R}'_o\|} = \frac{(x, y, \frac{n_w}{n_a} z_c \sqrt{1 + \frac{n_w^2 - n_a^2}{n_w^2} \frac{x^2 + y^2}{z_c^2}})^\top}{\|\mathbf{R}'_o\|}. \quad (31)$$

Hence, the equivalent camera position is: $\mathbf{C}' = (0, 0, \frac{n_w}{n_a} z_c \sqrt{1 + \frac{n_w^2 - n_a^2}{n_w^2} \frac{x^2 + y^2}{z_c^2}})^\top$ when the incidence angle is close to the optic axis of the camera, *i.e.*, $x, y \ll z_c$, we can have the approximation that the equivalent camera position $\mathbf{C}' = (0, 0, \frac{n_w}{n_a} z_c)^\top$.

Appendix B: Detailed Implementation of Water-drop Multiple View Stereo

Once the geometry of each water drop is obtained, we perform multiple view stereo to estimate depth. We propose a raytracing based triangulation. As illustrated in Fig. 11, for two corresponding points \mathbf{x}_1 and \mathbf{x}_2 and their rays of refractions, the triangulation aims to find the position of point \mathbf{p} which minimizes the Euclidean distance to all the rays. This idea can be directly extended to more than two water drops.

There are 3 main steps in our multiple view steres: (1) Inverse raytracing, (2) Corresponding points for different water drops, and (3) Triangulation.

Inverse Raytracing The goal of the inverse raytracing is to find the orientation of the ray of refraction. We call it inverse raytracing because we assume the ray is originated from the camera, refracted by the water drops and arrives at the objects.

As illustrated in Fig. 11, we show the inverse raytracing on the left water drop. Without loss of generality, we assume the camera position is $\mathbf{C} = (0, 0, z_c)^\top$ and the flat plate is $z = 0$, and the image-plane has corresponding pixels with the flat plane using rotation and scaling. According to Appendix A, the equivalent camera position is $\mathbf{C}' = (0, 0, z'_c)^\top = (0, 0, \frac{n_w}{n_a} z_c \sqrt{1 + \frac{n_w^2 - n_a^2}{n_w^2} \frac{x^2 + y^2}{z_c^2}})^\top$.

For a pixel on the image plane, with a corresponding point on the flat plate $\bar{\mathbf{x}} = (\bar{x}, \bar{y}, \bar{z})^\top$, we can find the refraction location, denoted as $\mathbf{x} = (x, y, z)^\top$, by using the constraints:

$$\frac{x}{\bar{x}} = \frac{y}{\bar{y}} = \frac{z - z'_c}{\bar{z} - z'_c}. \quad (32)$$

In practice, because finding the intersection between a flat plane and a line is easier than finding the intersection between a line and a curved surface, we specify \mathbf{x} and find the corresponding pixel $\bar{\mathbf{x}}$.

At point \mathbf{x} , the angle of incidence is:

$$\mathbf{R}_i = \frac{\mathbf{x} - \mathbf{C}'}{\|\mathbf{x} - \mathbf{C}'\|} = \frac{(x, y, z - z'_c)^\top}{\|\mathbf{x} - \mathbf{C}'\|}. \quad (33)$$

The surface normal is obtained according to the water drop geometry:

$$\mathbf{N} = \frac{\mathbf{N}'}{\|\mathbf{N}'\|} \quad (34)$$

$$\mathbf{N}' = (\frac{\partial z}{\partial x}, \frac{\partial z}{\partial y}, 1)^\top.$$

Then, the orientation of ray of refraction \mathbf{R}_o is obtained according to Snell's law:

$$\begin{aligned} \mathbf{R}_{i\parallel} &= (\mathbf{R}_i^\top \mathbf{N}) \mathbf{N} \\ \mathbf{R}_{i\perp} &= \mathbf{R}_i - \mathbf{R}_{i\parallel} \\ \mathbf{R}_{o\perp} &= \frac{n_w}{n_a} \mathbf{R}_{i\perp} \\ \mathbf{R}_o &= \mathbf{R}_{o\perp} + \mathbf{R}_{o\parallel} \\ \|\mathbf{R}_o\| &= 1. \end{aligned} \quad (35)$$

Finding Correspondence Between Water Drops Finding the corresponding pixels between different warped water drop images is a challenging task. Compared to the normal cameras, the distortion between water drop images is significantly worse. Moreover, unlike spherical mirrors/lenses where all the mirror/lenses share the same distortion, each water drop has its own distortion.

To solve this problem, we try to find the corresponding pixels on the angular-dewarped images. Note that, since the depth of image is not yet obtained, we cannot accurately dewarp the image. Thus, we dewarp the images solely according to the angle of refraction. Nevertheless, we find the angular-dewarping can significantly recover the images.

As introduced in the inverse ray-tracing, for a pixel $\bar{\mathbf{x}}$, with ray of refraction $\mathbf{R}_o = (r_{ox}, r_{oy}, r_{oz})^\top$, we project the pixel to

$$\begin{pmatrix} r_{ox} & r_{oy} \\ r_{oz} & r_{oz} \end{pmatrix}^\top. \quad (36)$$

Because the water drop surface is smooth and convex, implying the Jacobian on the surface is always positive, it means the angular mapping is one-to-one [Zorich and Cooke 2004]. Thus, we can map back the dewarped corresponding pixels to the warped image.

The third and fourth columns of Figure 14 show examples of the angular dewarping results and the dense correspondence on the dewarped images. Specifically, we use [Xu et al. 2012] for the dense correspondence estimation.

Triangulation Now we can perform the classical triangulation as illustrated in Fig. 11. Given a set of corresponding pixels on each water drop \mathbf{x}_j , we can obtain its outbound ray of refraction:

$$\mathbf{x}_j + \alpha \mathbf{R}_{oj}, \alpha \in \mathbb{R}, j = 1, 2, \dots \quad (37)$$

The triangulation's goal is to find the position of point \mathbf{p} , which minimizes the Euclidean distance to all the rays [Hartley and Zisserman 2004]:

$$\mathbf{p} = \left[\sum_j (\mathbf{I} - \mathbf{R}_{oj} \mathbf{R}_{oj}^\top) \right]^{-1} \left[\sum_j (\mathbf{I} - \mathbf{R}_{oj} \mathbf{R}_{oj}^\top) \mathbf{x}_j \right]. \quad (38)$$

The depth is the z component of \mathbf{p} .