

.....

Received 3 June 2015; accepted 8 August 2016

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

light sectioning using line lasers leads to regularly resolved shape reconstructions, it does not map any color information of the seafloor. Inglis et al. presented a method (Inglis, Smart, Vaughn, & Roman, 2012) in which a bathymetry map generated from laser line scans is used to improve the accuracy of a stereo 3D reconstruction.

In the research presented in this paper, a system based on light sectioning is extended to directly add color to the bathymetry by back-projecting each 3D point into the illuminated area of a photo to retrieve the RGB (red, green, and blue) values for that point in space to produce dimensionally accurate, high-resolution 3D reconstructions of the seafloor and objects on it in colors as if it were on land. The proposed method uses a single camera for acquiring both the laser scan profiles and the color information of the seafloor by projecting a laser line so that it always appears in the upper part of the image, while the lower part is illuminated using LED panels.

1.1. Objectives and Contribution

The aim of this work is to enable the generation of 3D color reconstructions of the seafloor at consistent and high, subcentimeter-order resolution. This is achieved by bringing together aspects of underwater equipment development, algorithm development, data acquisition at sea using AUVs and ROVs, and data visualization and processing. While it is unrealistic to aim to completely overcome the influence of environmental factors such as turbidity, or to overcome limitations in vehicle performance, the method developed should be robust and provide results that are satisfactory under conditions that can be reasonably expected at sea. The algorithms should be scalable and adapt to different conditions of operation. The objectives are as follows:

- Develop a system to collect data that can be used on different underwater platforms.
- Develop a robust algorithm to transform the collected data into 3D models with accurate shape and color information while compensating for environmental characteristics of the water properties, which works for data collected over different types of terrain and in varying water conditions.
- Be able to handle large datasets, both for generation and visualization of data.
- Prepare the data in a way that lends itself to extraction of useful information automatically or semiautomatically.

The contributions of the work presented here are as follows:

- The development of hardware and software to record all data necessary to generate 3D reconstructions of underwater environments that can be used on both AUVs and ROVs.

- The development of robust software for transforming data collected underwater into highly resolved 3D reconstructions, while compensating for environmental characteristics of the water properties.
- Assessment of the accuracy of the system.
- Deployment of the system at sea for collection of data at more than 2,000 m depth.

1.2. Organization of the Paper

Section 2 discusses the challenges of underwater mapping and the approach chosen to overcome them. Section 3 introduces the hardware setup and the necessary software to record data at sea. Section 4 explains the algorithm developed to transform the raw data into 3D reconstructions, and it discusses specific difficulties of mapping in water and the steps taken to overcome these. Section 5 shows the results from a mapping survey at a seamount covering a 3.5-km-long transect up the slope of the mount between 2,125 and 1,400 m depth. Section 6 gives the conclusions of the presented research.

2. EXISTING UNDERWATER MAPPING METHODS AND APPROACH CHOSEN

Generating dimensionally accurate digital reconstructions of the seafloor in three dimensions has essentially two parts: obtaining shape information and obtaining the corresponding color information. Stereo vision and SFM are two passive methods that obtain both topography and color from the same set of images. Both methods are based on feature matching between photos, which makes the acquisition of data simple and efficient. However, these methods suffer from difficulties of matching features in turbid water (Garcia & Gracias, 2011) and over terrain with few visual features, and they lead to variable resolution across the map as a function of the abundance of visual clues.

Techniques for mapping bathymetry with consistently high resolution include different types of structured light by employing scanning point lasers (Kocak, Caimi, Das, & Karson, 1999; Moore, Jaffe, & Ochoa, 2000), line lasers (Inglis et al., 2012; Kondo et al., 2004; Tetlow & Spours, 1999), or light pattern projections (Bruno et al., 2011), which make use of the known relative positions of the camera and the projector. Measurements of the time of flight of a light impulse are used in LIDAR (Light Detection and Ranging) (Harsdorf et al., 1999) and serial imaging systems (Dalgleish et al., 2013). Scanning sonars (Singh, Whitcomb, Yoerger, & Pizarro, 2000) and multibeam sonars (Jakuba & Yoerger, 2003; Roman & Singh, 2007) use high-frequency sound waves for the same purpose, which are unaffected by the turbidity of the water and offer longer ranges, but they cannot match the resolution of optical systems at low altitude. These methods calculate shape with feed-forward algorithms that yield maps of relatively uniform resolution,




Figure 1. Concept of the 3D color mapping technique. The laser line projection reveals the shape of the seafloor, which is used to generate a high-resolution bathymetry map, while LED panels illuminate the lower half of the camera's field of view. When the vehicle has moved forward, the camera passes vertically above the area swept by the laser line, and the color information from the illuminated lower half of the images taken there is used to map color to the bathymetry. The shade mounted behind the LED panels prevents them from illuminating the upper part of the image, which is kept dark to enable robust detection of the laser line projection. Adapted from (Thornton, Asada, Bodenmann, Sangekar & Ura, 2013).

and some also reproduce the reflective intensity, but generally they do not give any information about the color of the seafloor. In Roman et al. (2013), the authors describe a method in which multibeam acoustics and laser line profiling is used to collect high-resolution bathymetry data while at the same time stereo imaging is performed. Kunz and Singh propose a method in Kunz & Singh (2013) in which images from a single camera are fused with multi-beam sonar data to create 3D maps of underwater terrains.

The most common type of structured light used underwater is light sectioning (or "laser scanning"), a simple yet robust and accurate way of measuring shape through triangulation using a line laser and a camera. A line is projected onto an object and observed from a camera mounted a certain distance away from the line laser. Based on the known position and orientation of the laser relative to the camera, the shape of the object can be calculated. The method was first proposed 1976 by Agin (Agin & Binford, 1976) to measure the shape of objects in a tabletop experimental setup, and it has also been used for underwater applications since then (Klepšvik, Torsen, & Thoresen, 1990; Kondo et al., 2004; Roman, Inglis, & Rutter, 2010). This technique was applied to underwater robots in the current research, and it was ex-

tended to map color by partly illuminating the image and projecting these color values onto the bathymetry map. This enables the simultaneous capture of both the structure and color of the seafloor. While Agin used a setup with a rotating mirror to deflect the laser sheet and scan the object, for the method presented here the camera and the line laser are mounted with a fix angle and distance to each other, and they scan the seafloor as the underwater vehicle moves relative to the seafloor. The laser is positioned so that its projection always appears in the upper part of the image, while at the same time the lower part is illuminated with white light in order to obtain color information. Figure 1 illustrates the two steps for generating colored 3D reconstructions. The upper part of the image, where the seafloor is scanned with the line laser projection, is used to generate a high-resolution bathymetry map, to which in the second step color is assigned from the illuminated area to build a texture map. To collect these images, the line laser is fixed on the underwater vehicle so that it points vertically downward and projects a line onto the seafloor that forms a 90° angle with the forward moving direction. The camera mounted a certain distance away from the laser records images of the seafloor where the laser line appears on the

ground. It is tilted so that its field of view (FOV) covers the area vertically below as well as the laser line projection optimally for mapping altitudes between 0.8 and 3 m. The texture map is corrected for the color-dependent attenuation of light in water to reproduce the colors as if the water was drained.

While the data presented in Section 5 of this paper were recorded at an average altitude of 2.7 m along a single transect, the mapping device SeaXerocks 1 described in the following section had previously been deployed from altitudes as low as 1 m to map the distribution of manganese crust deposits (Bodenmann, Thornton, & Ura, 2010a). The setup has also been deployed from an AUV for a grid survey of a hydrothermally active area in Kagoshima Bay, Japan (Bodenmann, Thornton, Nakatani, & Ura, 2011). The 3D data were then used to determine the height and calculate the volume of a mapped hydrothermal chimney.

3. THE MAPPING DEVICE: SEAXEROCKS 1

3.1. Hardware

The SeaXerocks 1 mapping system consists of a camera, a line laser, LED panels, and a PC in a pressure tight housing, which can be mounted on an AUV or ROV. If necessary, the hardware of SeaXerocks 1 allows a 1,200 kHz Teledyne Workhorse Navigator Doppler velocity log (DVL) and a Paroscientific 8CB4000-I depth sensor to be connected directly to the device for synchronized recording of navigation data. Alternatively, if available, the navigation data from the vehicle can be used, provided the SeaXerocks and vehicle clocks are synchronized. For ROV deployments, all devices are mounted on a payload skid, together with other sensors if required, as shown in Figure 2. A green (532 nm) 50 mW line laser with an opening angle of 64° in water is mounted at the front of the skid. A wide opening angle is desirable to maximize the scanned swath. Images are recorded with a Point Grey Firefly FMVU-13S2C color camera at a rate of 12 fps with a resolution of $1,328 \times 1,048$ pixels. It is fitted with a 2.8 mm Goyo Optical GMT12813SW lens, which yields a field of view of $62.7^\circ \times 52.3^\circ$ in water. The camera-lens assembly is enclosed in a 3,000 m depth rated flat-port underwater housing. A pair of 30 W LED panels are used for illumination because of their high energy efficiency and compact size. They are placed so that the overlapping volumes of their light cones and the field of view of the camera have minimum overlap. They are separated from the laser by a board that acts as a shade to prevent casting light on the laser line projection. All data are recorded by an Advantech PCM-9382 single board PC in a 3,000 m depth rated pressure tight housing. The system is normally operated from approximately 2 m altitude, which is a good compromise between area covered and image quality. At altitudes exceeding 3 m, the images start to get dark and low in contrast.




Figure 2. The devices of SeaXerocks 1 are fixed on a payload skid, shown in the lower part of the figure, which is mounted in the payload bay of the ROV Hyper-Dolphin, shown above.

While the line laser is mounted vertically, the camera makes an angle of 20° with the vertical, as this provides an optimal operation range for the system while keeping the area vertically below the camera in the field of view. The baseline between the camera and the line laser is chosen as a function of the target altitude of a mission and the available space on the underwater platform. Figure 3 shows that, for a given altitude, the longer the baseline is, the better is the vertical resolution. At the same time, the camera's field of view is limited, which limits the distance at which the laser can be mounted, so that its projection is still seen by the camera. The lower the altitude, the shorter the distance needs to be, and the optimum baseline is as long as possible (better resolution) but as short as necessary (out of frame if too short). The calculations behind Figure 3 show that this is the case for a baseline approximately equal to the altitude. As a too low altitude causes the laser line to be out of frame, the baseline needs to be adjusted for the lowest mapping altitude of a mission. While SeaXerocks 1 is normally operated at approximately 2 m over the seafloor, the constraints due to the size of the vehicle it was operated from and its ability to accommodate equipment have limited the distance between camera and laser to between 0.8 and 1.5 m depending on how the system was mounted.




Figure 3. The vertical resolution of the laser scanning setup depends on the altitude of the system above the seafloor and the length of the baseline between the camera and the line laser. The simulation shows that the longer the baseline is, the better (smaller values) is the resolution, but that if the altitude is too low, the laser line projection is not recorded anymore, which is apparent in the graph at the bottom left where there are no data because it is outside the camera's field of view.

3.2. Coverage and Resolution

The area that the system can map in a given time and the resolution of the generated 3D reconstructions depend on the altitude and speed of the underwater vehicle as well as the configuration of the camera and the line laser. The mapping system is usually deployed from an underwater vehicle maintaining a constant speed of 30 cm/s and hovering at an altitude of ~ 2 m, leading to a coverage of $\sim 2,500$ m²/h without overlap. At this altitude, the bathymetry resolution of the laser scanning system is 25 mm in forward (x), 2 mm in across track (y), and 5.4 mm in the vertical (z) direction, assuming a baseline of 0.8 m. The illuminated area directly below the camera is captured with a resolution of 2.1×1.7 mm per pixel. However, the mapping system can collect data from altitudes between 0.8 m and approximately 3 m, and Table I shows the corresponding resolutions and covered areas, as well as the environmental parameters of SeaXerocks 1. While the laser scanning bathymetry resolution along the track is limited by the image acquisition frequency, the resolution across the track and in the vertical direction changes with the mapping altitude. The accuracy of the system has been tested in a pool experiment (Bodenmann, Thornton, Ura, & Sangekar, 2010).

4. 3D RECONSTRUCTION ALGORITHM

The processing software is written in C++ and makes use of the open source libraries OpenCV (Open Source Computer Vision Library) (Bradski, 2000) and CGAL (Computational Geometry Algorithms Library) (The CGAL Project, 2015).

Table I. Specifications of the SeaXerocks 1 camera and lighting systems assuming a baseline of 0.8 m, a camera tilt angle of 20° , and a forward speed of 30 cm/s.

Mapping altitude	0.8 m (min)	2 m (typ)	3 m (max)
Swath below camera	0.9 m	2.3 m	3.4 m
Swath at laser projection	1.2 m	2.6 m	3.8 m
Area covered in 1 h	988 m ²	2,470 m ²	3,706 m ²
Bathymetry resolution (x/y/z)	25/0.9/1.2 mm	25/2/5.4 mm	25/2.8/11.2 mm
Size of pixel below camera (x/y)	1.2/0.7 mm	2.1/1.7 mm	2.9/2.6 mm
Camera resolution (h \times v)		1,328 \times 1,048	
Field of view in water (h \times v)		62.7° \times 52.3°	
Weight in air		12.5 kg	
Weight in water		6.9 kg	
Depth rating		3,000 m	
Power consumption		95 W	

The algorithm is structured into two major parts, both of which make use of the seafloor images and the navigation data. First the bathymetry is calculated from the laser line projection in the upper half of the images. In the second part, the texture map is built using the bathymetry map and the lower part of the images, which is illuminated by white light. The major steps of the algorithm are shown in the flowchart in Figure 4 and are explained in the remainder of this section.

4.1. Camera Model and Undistorting of Images

The difference in the refractive index between water and air causes the light rays to be refracted when entering the underwater housing, which leads to a distortion in the image observed by the camera, on top of the distortion caused by the lens. Treibitz, Schechner, Kunz, & Singh (2012) showed that when tracing rays corresponding to pixels through the viewport into the water and then extending the rays in water backward, due to the double refraction at the view port, they intersect in an apparent viewpoint that is generally at a different position from the actual camera. The location of the apparent viewpoint varies according to which ray is followed, hence a camera in an underwater housing generally does not satisfy the single viewpoint (SVP) model. However, in the present setup where the entrance pupil of the camera is 5 mm behind a 20-mm-thick view port, the position of the apparent viewpoint varies by less than 0.6 mm, hence the SVP model is considered a good approximation, and the error caused by this simplification is thought to be negligible.

The algorithm assumes the pinhole camera model (Hartley & Zisserman, 2004), and the images are undistorted using an algorithm based on Zhang (2000) and Bouguet (2004), implemented in OpenCV. There are other

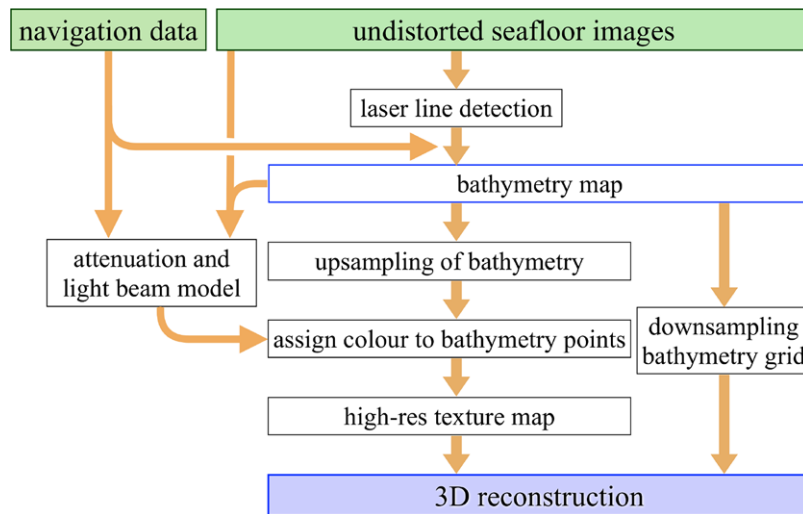


Figure 4. Flowchart illustrating the main steps of the 3D reconstruction algorithm.

approaches that use a more sophisticated lens model (Aggarwal & Ahuja, 2002) or correct for the distortion introduced by the double refraction at the view port by explicitly computing the refraction (Jordt-Sedlazeck & Koch, 2012), but for the current research the accuracy achieved with the pinhole camera model was found to be sufficient. We use a right-handed camera coordinate system where the x-axis is aligned with the camera's optical axis, the y-axis points to the right, and the z-axis points downward.


4.2. Subpixel Laser Line Detection

The position of the laser line within the seafloor images, together with the position and orientation data, allows calculation of the bathymetry of the seafloor as a list of 3D points. For this it is necessary to robustly and accurately detect the laser line. First, the transition between the illuminated and the nonilluminated area is detected, as shown in Figure 5(a), because the laser line should only be searched in the area not illuminated by the LED panels. A function scans every column of the image starting from the top, looking for a sudden rise in brightness. To avoid picking up the laser line instead of the shadow-to-light transition, the algorithm rejects points where the brightness drops again below a certain threshold.

The algorithm masks the area identified as illuminated and goes on to extract the laser line projection in the area above. Depending on the distance to the seafloor, which can vary strongly even within a single image due to the vertical profile of the seafloor, the brightness of the observed laser line can vary significantly, and it may be disrupted due to occlusions. As the line is not always sharp and there are other image features with even stronger

contrast, standard line detection algorithms, such as the Canny edge detector, do not perform well. If the laser was the only source of light, a relatively simple algorithm looking for the brightest pixels in the image could be used. However, stray light from the LED panels, and from the ROV lights if operated from an ROV, leads to other bright areas in the image. For these reasons, an algorithm that takes into account the color and a measure of continuity was developed.

In the first step, the color images are transformed into one-channel images, in which high values indicate that the corresponding pixel in the original image was likely to be illuminated by the line laser. It was found that rather than binarizing the image, an approach assigning different values depending on how strongly a pixel resembles the color of the laser line projection is more robust. For this conversion, a reward function is applied on every pixel that calculates a single value based on its RGB values. This "reward function" has the same concept as a cost function, but rather than minimizing a cost, the aim is to maximize a reward. It assigns high values to colors in which the laser line projection appears, and low values to the other colors. While the color of the laser has a very narrow band (green, at 532 nm), its brightness observed by the camera varies depending on the distance of the vehicle to the seafloor, and the reward function should be robust to that. Among different approaches tested, a function that separates the RGB color space into two parts, where one part is considered to represent the color of the laser line projection and the other part everything else, was found to be most robust. To train the reward function, a training set of pixels belonging to the laser lines from three to five photos and a set of pixels belonging to the background in the same photos are provided, which each form a cluster when plotted in the RGB space, as



(a) Photo of the seafloor illuminated by the LED panels in the lower half and with the laser line projection above. The blue line marks the identified transition between the illuminated and the non-illuminated area.

(b) Visualisation of the reward function assigning high values to pixels with a colour close to the colour of the laser. The area identified as illuminated by the LED panels is masked and shown here in blue.

(c) Image showing the identified laser line marked in red in the image of the reward function.

Figure 5. Steps for extracting the laser line projection in seafloor photos. In a later step, the lower, illuminated part of the image is used to map color to the 3D reconstruction.




Figure 6. Plot of color values of the pixels in the seafloor image used to train the SVM, with pixels pre-labeled as belonging to the laser line (green) and of pixels belonging to the background (black), each forming a cluster in the RGB color space. The separation plane (red) is determined using the SVM method.

shown in Figure 6. Support vector machines (SVMs) (Vapnik & Lerner, 1963) are used to calculate the most ideal plane that separates the two clusters. While SVMs can be used to construct hyperplanes, it was found that a flat plane was sufficient for the current application. The separation plane is then used to calculate the reward function, which indicates how strongly the color of a pixel resembles the color of the laser line. If the color values of a pixel input to the reward function are on the side of the laser line color, the distance from the plane is returned as the reward value, to reflect that the further a point is in that side, the more certain

it is that it actually shows part of the laser line; if it is on the opposite side, the reward is 0:

$$f_{\text{reward}}(u, v) = \max \left(\frac{a_1 R(u, v) + a_2 G(u, v) + a_3 B(u, v) - d}{\sqrt{a_1^2 + a_2^2 + a_3^2}}, 0 \right), \quad (1)$$

where a_1, a_2, a_3 , and d are the parameters of the plane determined with the SVM method described above, and $R(u, v)$, $G(u, v)$, and $B(u, v)$ are the red, green, and blue color values for the pixel at the coordinate (u, v) . This function is applied on the original image, and the reward function values can be represented as a gray scale image, as shown in Figure 5(b). This approach proved to be very robust, and one set of parameters can be applied for an entire dataset.

Finally, the exact position of the laser line is determined in the reward function image. As the line laser optics has approximately the same opening angle as the camera lens, the laser line spans the entire image except in areas where it is occluded. The line laser is aligned so that its projection is parallel with the vehicle's y-axis (axis pointing to the right-hand side), and so is the camera's horizontal axis. Therefore, the laser line projection is never vertical in the image, i.e., it appears only once in every image column, and the line extraction algorithm can proceed columnwise. Each column is convolved with a triangular kernel, and the maximum value within the column is identified. If the value is above a certain threshold, it is considered as the laser line in that column. To determine the position of the laser line with subpixel accuracy, the centroid is calculated in a window around where the maximum was found. This was proposed in Izquierdo, Sanchez, Ibañez, & Ullate (1999), based on work by Welch (1993). This is done for every column of the image. To eliminate false detections due to marine snow, a

measure of continuity is calculated, and secluded points are rejected as outliers. To do so, continuous patches of 10 pixels or more are identified and labeled as patches confirmed to show the laser line projection on the seafloor, based on the fact that marine snow usually appears as bright spots that are only two to three pixels wide. However, in this method, areas where the laser line detection is not continuous, for example because it appears faint in the image, are missed. Therefore, in a second round, single pixels identified as the laser line are also added to the confirmed laser line points if they lie between two patches of confirmed laser line points. Points not falling into either group are then rejected as outliers. Figure 5(c) shows the final result with the identified center of the laser line colored red.

4.3. Calculation of High-resolution Bathymetry Map

For each pixel in a seafloor image identified as part of the laser line, a vector pointing in the direction of its corresponding light ray is calculated using the rectilinear lens model. The intersection with the plane defined by the line laser leads to the position where the laser projection illuminated the seafloor, in the reference frame to the camera. Using the position and orientation of the vehicle, the coordinates are transformed to a global reference frame. The (u, v) coordinates of a pixel measured from the principal point identified as being part of the laser line lead to a direction vector \vec{q} from the camera,

$$\vec{q} \Big|_{\text{cam}} = \begin{pmatrix} f \\ u \\ v \end{pmatrix},$$

where f is the focal length of the camera measured as a multiple of the pixel size. This vector is in the reference frame with the same orientation as the camera, and in order to convert the camera reference frame to the vehicle reference frame, the following coordinate transformation is applied to correct for the camera's pitch β :

$$\vec{q} \Big|_{\text{camOrientVehicle}} = R_y(\beta) * \vec{q} \Big|_{\text{cam}},$$

where R_y is the rotation matrix for the y -axis (Arfken, 1985).

The vector in the vehicle's reference frame to point \vec{p} on the seafloor, corresponding to the pixel (u, v) in the image, can be written as

$$\vec{p} \Big|_{\text{vehicle}} = \vec{c} \Big|_{\text{vehicle}} + k \cdot \vec{q} \Big|_{\text{camOrientVehicle}}, \quad (2)$$

where \vec{c} is the position of the camera in the vehicle's reference frame, and k is an *a priori* unknown factor. Given the position of the line laser \vec{s} and the normal vector \vec{n} of the line

laser plane, k can be calculated according to the formula for line-plane intersections as

$$k = \frac{(\vec{s} \Big|_{\text{vehicle}} - \vec{c} \Big|_{\text{vehicle}}) \cdot \vec{n} \Big|_{\text{vehicle}}}{\vec{q} \Big|_{\text{camOrientVehicle}} \cdot \vec{n} \Big|_{\text{vehicle}}}. \quad (3)$$

By inserting the result from Eq. (3) into Eq. (2), the position of the point can be expressed with reference to the vehicle.

The position \vec{v} of the vehicle throughout time is determined using dead reckoning applied to the DVL speed measurements and using depth sensor measurements. Roll (ϕ), pitch (θ), and yaw (ψ) are also recorded by the DVL. While the navigation data from the DVL are saved at a rate of 4 Hz, images are recorded at 12 fps. The position and attitude measurements are interpolated for every image according to its time stamps. As there are no sudden accelerations due to the inertia of the vehicle's large mass and the surrounding water, the inaccuracy introduced by the interpolation is assumed to be negligible. Using the vehicle's position and orientation, the observed point on the seafloor can be represented in the global reference frame as

$$\vec{p} \Big|_{\text{world}} = \vec{v} \Big|_{\text{world}} + R_z(\psi) * R_y(\theta) * R_x(\phi) * \vec{p} \Big|_{\text{vehicle}}. \quad (4)$$

The above-described procedure is applied to every identified laser line point in every image. The result is a list of 3D points, which are then transformed into a mesh of triangles using Delaunay triangulation (Delaunay, 1934), representing the bathymetry. This mesh is saved as a polygon file format (ply) file. The algorithm outputs fixed point ascii ply files, which can be converted into binary ply files, for example using MeshLab,¹ to reduce the file size.

4.4. Assigning Color to Bathymetry Points to Generate a Texture Map

Color is matched to the bathymetry points from the lower, illuminated part of the seafloor photos. For each bathymetry point, the set of images in which the bathymetry point appears is determined. The bathymetry point is backprojected into each of these images, and if the backprojection is inside the LED illuminated area [e.g., bottom of Figure 5(a)], the color is retrieved and corrected for attenuation and the effects of nonuniform illumination; this is explained in more detail below. In case the bathymetry point appears in more than one photo, the colors from all the photos in which it appears are used and blended according to a function that assigns a higher relative weight to photos where the point

¹MeshLab is open source and available at <http://meshlab.sourceforge.net/>

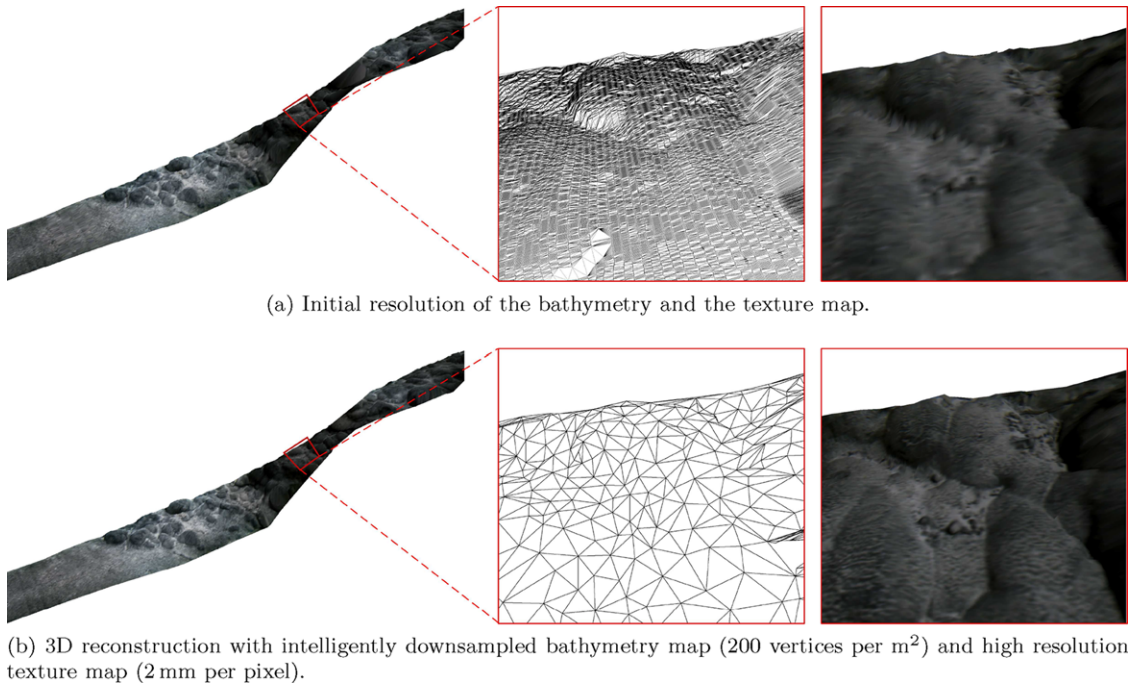


Figure 7. Texture map resolution improvement and mesh size optimization.

appears close to the center than those where it appears close to the border: the relative weight of the contribution from one image is linearly proportional to the distance of the corresponding pixel to the closest image border. This yields the color for every bathymetry point, which, after correcting the color as described in the following paragraph, is saved in a texture map as a jpg or png file and is linked by the ply file. Figure 7(a) shows an example of such a colored 3D reconstruction. Initially the texture map resolution is bound to the bathymetry resolution, which, assuming an altitude of 2 m, is 25 mm along the track and 2 mm across the track according to Table I. It is then upsampled to match the resolution of the original images, as shown in Figure 7(b).

4.5. Lighting and Attenuation Model Used for Color Correction

As light travels through water, it is absorbed and scattered at different rates depending on the wavelength (i.e., the color). Figure 8 shows the absorption coefficient $a(\lambda)$ as a function of the wavelength λ for water in the spectrum of visible light as published in Pope & Fry (1997). Red is absorbed most strongly in water, but the exact course of the absorption parameter varies depending on what other substances are dissolved, and therefore it changes between different sites at sea. Scattering, expressed by scattering coefficient $b(\lambda)$, also leads to a reduction of light reaching the camera. The sum




Figure 8. Absorption coefficient of pure water as a function of the wavelength of light.

of absorption and scattering is the attenuation, expressed by the attenuation coefficient

$$c(\lambda) = a(\lambda) + b(\lambda).$$

The total amount of light of a straight beam of light lost on a distance r is described by

$$E(r) = E(0)e^{-c(\lambda)r}, \quad (5)$$

where E is the irradiance. Additionally, scattering at small angles and refraction due to inhomogeneities of the refraction index lead to blurring in the image (Duntley, 1963).

While blurring does not change the position of the laser line projection within the image, it lets it appear less sharp, which also applies to the area illuminated by the LED panels. The effects of blurring cannot be easily overcome and pose a limit for the best resolution that can be achieved.

When retrieving the color from a seafloor image, the algorithm compensates for the color-dependent attenuation as well as for backscattering, based on the distance the light has traveled in water. It assumes Lambertian surfaces and constant turbidity and absorption properties of the water in the area where images are collected. Since the geometry of the mapping device and the relative position of the observed spot on the seafloor are known, the exact distance from the light sources to the seafloor, and from there to the camera, can be computed.

To estimate and correct the change in color, a simplified version of the Jaffe-McGlamery model (Jaffe, 1990; McGlamery, 1979) is used. The irradiance E_T reaching the camera is the sum of the direct (E_D), the forward-scattered (E_{fs}), and the backscattered (E_{bs}) irradiance:

$$E_T = E_D + E_{fs} + E_{bs}. \quad (6)$$

The direct component is the light that travels directly from the light source(s) to the spot on the seafloor and from there to the camera. The model is able to incorporate multiple light sources, which are assumed to be point light sources defined by their position, orientation, and their beam pattern, which can be expressed as a function of the angle φ from the optical axis. As light travels in water, it is attenuated according to Eq. (5), and the intensity is reduced additionally due to the point source spreading loss. When the beam is reflected off the seafloor, its color and intensity are changed according to the reflectance $M(P, \lambda)$ of the seafloor at that point P . The influence of the orientation of the reflecting surface is ignored in the model used in the current algorithm. The direct component of the light originating from light source L_i reaching the camera after being reflected off a point P on the seafloor can therefore be written as

$$E_{D,L_i}(P, \lambda) = E_{L_i}(\varphi) \frac{e^{-c(\lambda)R_i}}{R_i^2} M(P, \lambda), \quad (7)$$

where E_{L_i} is the irradiance of light source L_i , and R_i is the total distance from that light source to the point P on the seafloor and from there to the camera.

The backscatter component depends on the common volume between the illumination cone(s) and the camera's field of view. The volume of overlap increases with the distance to the seafloor, however the contribution from the volume close to the camera and the lights is much higher than the contribution from the overlapping volume at larger distances, due to the point spreading loss following the inverse-square law. Therefore, backscatter was approximated by a constant for distances above 0.8 m, at which mapping was performed.

As the water at depths exceeding 1,000 m, where the system is deployed most of the time, is usually very clear, the blurring introduced at mapping distance is considered negligible at the resolution at which the images are captured. With these assumptions, and by inserting Eq. (7) into Eq. (6), we obtain

$$E_T(P, \lambda) = E_{bs}(\lambda) + \sum_i E_{L_i}(\varphi) \frac{e^{-c(\lambda)R_i}}{R_i^2} M(P, \lambda) \quad (8)$$

and solved for the reflectance:

$$M(P, \lambda) = \frac{E_T(P, \lambda) - E_{bs}(\lambda)}{\sum_i E_{L_i}(\varphi) \frac{e^{-c(\lambda)R_i}}{R_i^2}}. \quad (9)$$

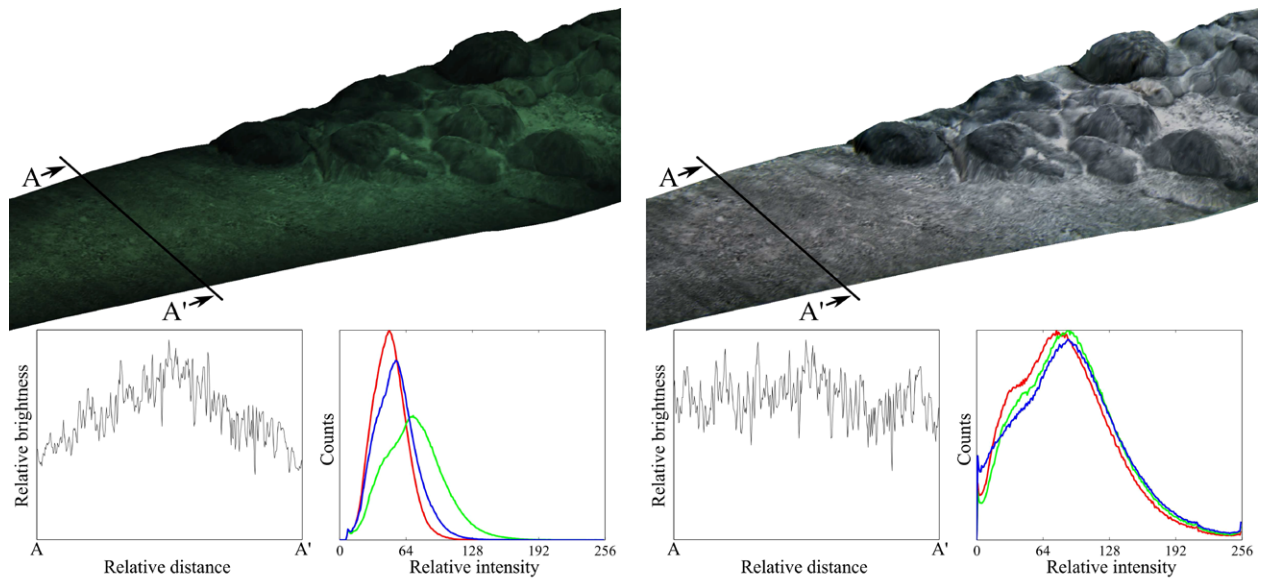
While these equations apply for a single wavelength, the camera bins all visible wavelengths into the three color channels red, green, and blue. To reconstruct the reflectance of a spot P on the seafloor appearing in the camera image with the colors (R_{raw} , G_{raw} , B_{raw}), Eq. (9) is applied. However, due to unmodeled effects, at large distances the measured intensity in the seafloor images deviates from what the term $\frac{e^{-c(\lambda)R_i}}{R_i^2}$ suggests. It is replaced by a parameter $F(R)$ (one per color band) that is estimated empirically using a number of seafloor images taken at different distances that show known objects, such as sand or rocks that were sampled, or if available manmade objects. The set of parameters (one per color channel) correcting for backscatter, which we call S , is estimated using images taken high enough above the seafloor, so that the ground is not visible and the image would be completely black in the absence of backscatter and other sources of noise. This allows us to compensate for backscatter and other sources of noise at the same time, such as CCD noise. The reflectance of a spot on the seafloor appearing in a pixel is then calculated using

$$\begin{cases} R_{corr} = \frac{R_{raw} - S_{red}}{\sum_i E_{L_i,red}(\varphi) F_{red}(R)}, \\ G_{corr} = \frac{G_{raw} - S_{green}}{\sum_i E_{L_i,green}(\varphi) F_{green}(R)}, \\ B_{corr} = \frac{B_{raw} - S_{blue}}{\sum_i E_{L_i,blue}(\varphi) F_{blue}(R)}. \end{cases} \quad (10)$$

These values are assumed to be the color as it would appear if there was no water. Figure 9(a) shows a 3D reconstruction where the color from the seafloor photos is applied directly without correcting it, and Figure 9(b) shows the same reconstruction where the color was corrected according to Eq. (10).

4.6. Reducing of Mesh Density

While a high resolution of bathymetry is desirable for applications such as mineral resource mapping and identification of benthic animals, rendering the entire 3D reconstruction with full bathymetry resolution is not practical because of the drop in performance due to the high polygon count. One



(a) 3D reconstruction where colour values were applied from the original photos without correcting them for attenuation or the non-uniform light beam. The brightness measured along one line across the transect is significantly higher in the centre than at the sides and the histogram is unbalanced, with red lagging behind blue and green.

(b) Colour balanced 3D reconstruction where the non-uniform light beam was compensated. The relative brightness is even across the transects, with small-scale undulations due to the presence of nodules on exposed sandy surfaces.

Figure 9. Adjusting the color balance and correcting for the nonuniform light beam. The graph at the bottom left of each subfigure shows the brightness measured along a line across the mapped transect, and the histogram shows the color distribution of the texture map, where the color of the lines is representative for each color channel. All units are arbitrary.




Figure 10. RMS error of the downsampled bathymetry map compared to the original.

way to overcome this would be to render the parts of the model further from the viewer with less resolution. However, rendering and level of detail processing were not the focus of this work, so a naive renderer was implemented, necessitating a reduction in the number of vertices. The bathymetry points are downsampled in such a way that the

shape is affected minimally, using an edge collapse algorithm developed by Lindstrom & Turk (1998) implemented in CGAL. The number of vertices is reduced by removing and repositioning vertices so as to alter the shape as little as possible. Figure 10 shows the root mean square (RMS) vertical error incurred by downsampling the number of vertices

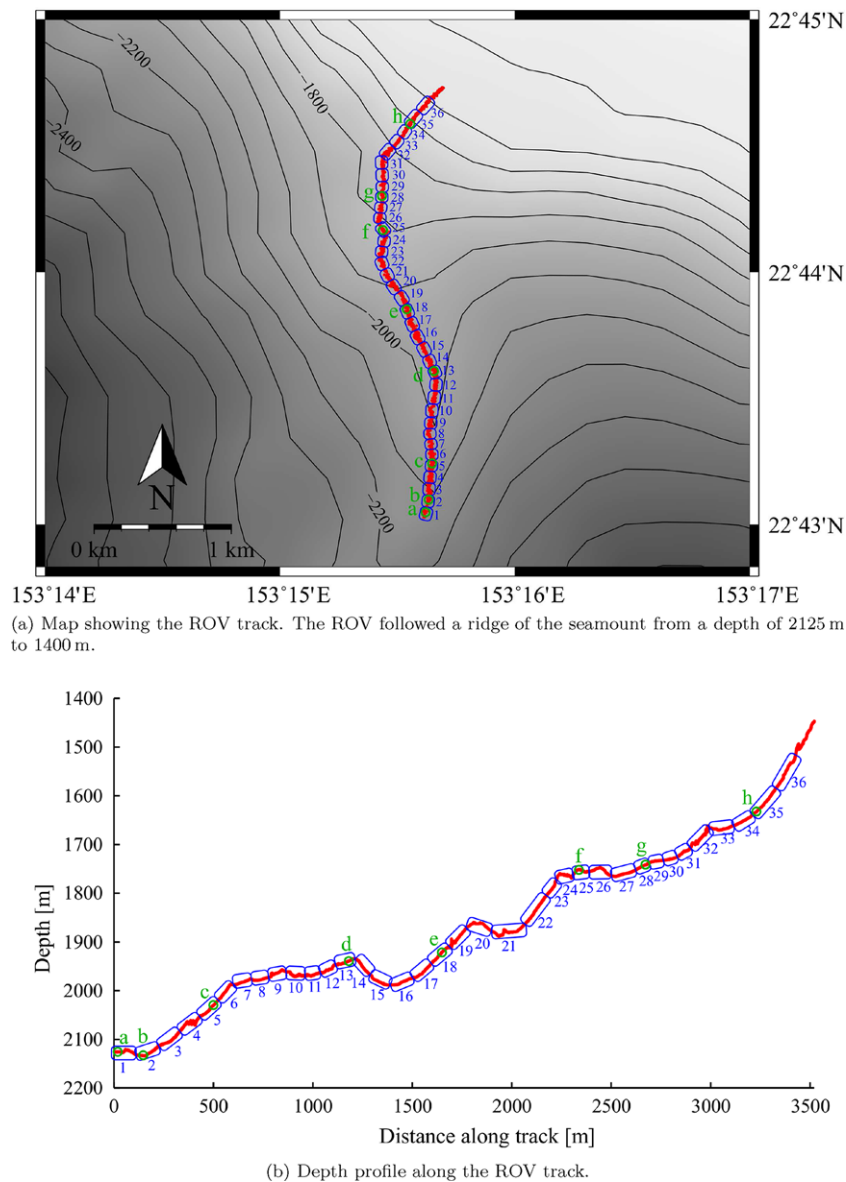


Figure 11. Mapping track during dive HPD1354 where a 3.5-km-long transect up a slope of the Takuyo #5 seamount was mapped in 6 h. The numbered boxes identify the sections of the map whose top views are shown in Figure 12, and the letters refer to closeup views shown in Figure 13.

on the example of a 58 m² large 3D reconstruction, part of which is shown in Figure 7. It compares naive subsampling, where vertices are evenly discarded in the grid until the desired vertex density remains, to Lindstrom-Turk's method. Decreasing the number of vertices reduces both the file size and the size of the 3D map in memory when loaded in the viewer software. While the original ply file in binary format storing coordinates as floating point values in meters is 90.5 MB, with 18,600 vertices per square meter,

downsampling to 10,000 vertices per square meter brings the file size to 48.5 MB (independent of the downsampling method) and to 0.49 MB with 100 vertices per square meter. While the error strictly keeps increasing with the reduced number of vertices in the case of Lindstrom-Turk's method, it fluctuates with naive subsampling because of the random change when eliminating vertices, but it has the same trend of increasing with a reduced number of vertices. The graph shows that the Lindstrom-Turk method always generates




Figure 12. Top view of the 3D reconstruction generated from the mapping data collected with SeaXerocks 1. The width of the mapping transect varies because of the fluctuation in the ROV's altitude while recording the data, and in places where the vehicle was flying too low or the laser line was occluded, gaps appear in the reconstruction. The position within the mapping track for each numbered section is indicated in Figure 11, and the closeup views of the features in the green boxes are shown in Figure 13.




Figure 12. Continued.

better results than the naive method. Down to 200 vertices per square meter, the RMS error only increases slightly, but it starts increasing rapidly below 120 vertices per square meter, which is why 200 vertices per square meter was determined to be an ideal compromise between file size and accuracy for this terrain. Figure 7(b) shows the downsampled 3D reconstruction with 200 vertices per square meter

with the texture map with 2 mm per pixel resolution. At this resolution, the file size of the dataset is 1.0 MB for the ply files and 2.1 MB for the texture map saved as a jpg file.

With this method, high-resolution 3D reconstructions of the seafloor of kilometer-long transects can be generated and displayed in a suitable viewer, such as MeshLab. The mapping data can also directly be analyzed with

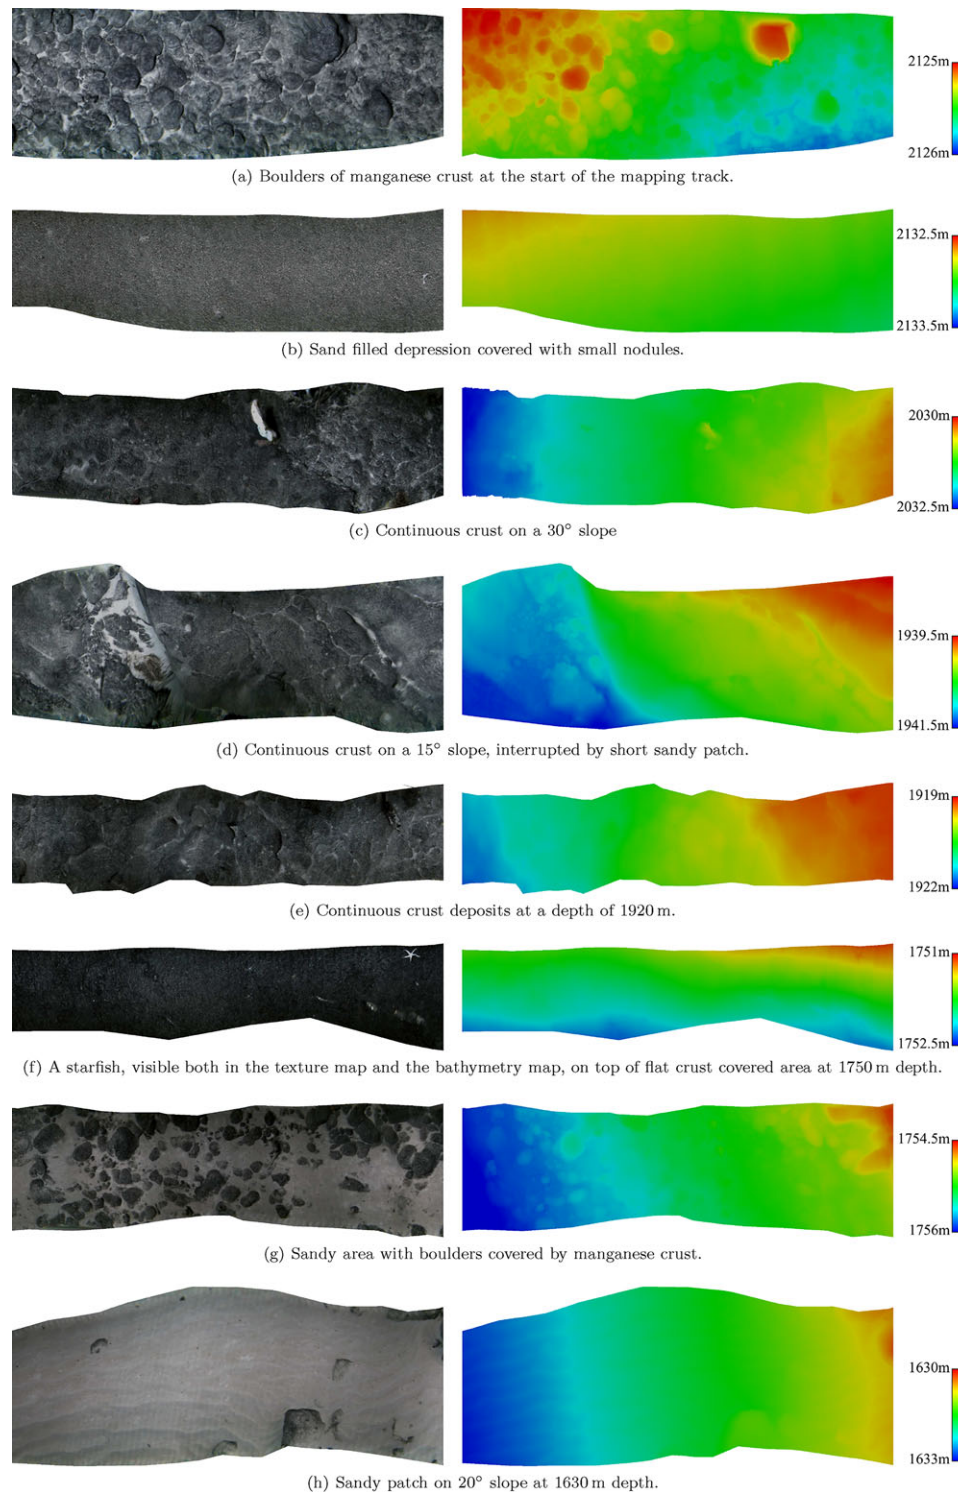


Figure 13. Closeup top views of a selection of areas of interest with a view of the colored 3D reconstructions on the left-hand side and the corresponding depth maps on the right-hand side. The color bars indicate the depth range in meters.

algorithms to extract scientifically useful information (Thornton, Asada, Bodenmann, Sangekar, & Ura, 2013). It is also possible to generate 3D reconstructions for sparse “lawnmower” grids and dense grids (Bodenmann et al., 2011) depending on which mode of observation is most suitable for the question being asked. In the present paper, a single transect was chosen since the difference in distribution of deposits at different depths on the slope of a seamount was of interest, and simple dead-reckoning was used for the navigation data.

5. RESULTS FROM DEPLOYMENT AT TAKUYO #5 SEAMOUNT IN THE NORTHWEST PACIFIC

At the time of this writing, the SeaXerocks 1 mapping system has been deployed on 11 cruises and consistently collected data from various locations on the seafloor around Japan. During the NT12-05 cruise of JAMSTEC, a 3.5-km-long transect up the slope of the Takuyo #5 seamount was mapped in a single dive, as shown in Figures 11 and 12. There are manganese-rich crust deposits on the slopes of the mount, which precipitated out of the water over several million years (Usui et al., 2007), making them an interesting target for geological studies, and they are of interest as a potential resource for mining. Visual and acoustic mapping was performed in order to assess the volumetric distribution of the deposited crust. While the acoustic data supplied thickness data of crust-covered areas, the visual maps were used to assess the horizontal distribution of manganese crust and to verify whether the acoustic reflections were from crust or from other surfaces, such as sediments (Thornton et al., 2013).

The mapping device was mounted onto the ROV Hyper-Dolphin shown in Figure 2. It followed a ridge up the seamount at an average altitude of 2.7 m, leading to a mapped swath averaging 3 m. Figure 12 shows the top view of the entire 3,566-m-long mapped transect that was mapped in 6 h, covering an area of approximately 10,000 m², and Figure 13 shows eight closeup views of areas of particular interest with their corresponding depth maps. At a depth of 2,126 m, the mapped track starts with big boulders of manganese crust, as shown in Figure 13(a). At 100 m further up the slope at a depth of 2,133 m there is a sand-filled depression that led to a flat 50-m-long sandy patch covered with small pebbles, as can be seen in Figure 13(b). After the sand-filled depression, the terrain continues similarly as before for another 300 m, with boulders of manganese crust and intermittent sandy patches. At a depth of 2,030 m at approximately 500 m traveled distance along the track, the crust cover becomes more continuous on a 30° steep slope and slabs of continuous crust appear, as the detail in Figure 13(c) shows. After a change in steepness at around 600 m distance traveled, the terrain becomes flatter and the crust cover becomes less continuous, with boulders of crust interrupted by sandy patches. At 1,200 m distance

traveled, large slabs of crust begin to appear again, as Figure 13(d) shows. On the following down-hill slope, the surface changes again into large boulders of manganese crust deposits interrupted by patches of exposed sand, with the ratio of crust to sand shifting more and more toward sand as the inclination becomes flatter. This trend is reversed after the lowest point is passed, with the ratio of exposed crust boulders increasing with the angle of the slope on the way up the slope. At 1,600 m distance traveled, large continuous deposits of crust appear, as shown in Figure 13(e), before the crust cover becomes again more patchy. This kind of terrain continues up to 2,300 m of the distance traveled along the track at a depth of 1,750 m, where the terrain changes abruptly to solid continuous and flat crust cover, as shown in Figure 13(f), stretching over approximately 120 m. In the same figure, a starfish residing on the manganese crust can be seen. It also shows up in the bathymetry map, which demonstrates its high vertical resolution. On the following slope, the crust cover becomes less dense, and boulders of crust on an otherwise sandy surface, as shown in Figure 13(g), dominate most of the next 700 m. At a depth of 1,630 m, after 3,300 m of mapped terrain on a 20° slope there is a large sandy area, with no crust deposits, as shown in Figure 13(h).

The 3D reconstruction makes it possible to objectively assess the crust cover on a section of the seamount. It shows that the changes in coverage are sudden, and while there is some relationship to the slope's inclination, there is no obvious pattern, stressing the need to map areas of interest at high resolution.

6. CONCLUSIONS

This paper describes a method for mapping both the shape and color of the seafloor or objects on it at high resolution. It extends laser line scanning, which builds a millimeter-order resolution bathymetry map, to a real-color 3D reconstruction by mapping color information from illuminated areas of an image to the bathymetry. The method has been implemented using a single camera with a partially illuminated field of view. This setup has been deployed multiple times at sea in order to generate data products that cover ~10,000 m² in a single dive with bathymetry data at millimeter-order resolution, which is higher than that typically achieved using acoustic methods and is comparable to methods that use visual features. The 3D reconstruction of a 3.5-km-long transect up a slope of a seamount demonstrates its application to manganese-rich crust surveys. The reconstruction allows us to clearly distinguish between exposed crust, sand, and transitional surfaces, allowing for an objective estimate of crust coverage along the mapped trajectory. The system can be applied to map almost any type of seafloor at depths of up to 3,000 m.

In the context of other methods available for high-resolution studies of the seafloor, the major advantage of the

method described lies in the consistent high resolution that it can achieve due to the relatively long baseline between laser and camera. Since the laser line is a strong visual feature and additionally is less susceptible to backscatter than lights or strobes that illuminate the entire image, the method remains relatively robust to a certain level of turbidity in the water. Furthermore, since the method does not rely on visual features of the seafloor but actively projects a laser line, it also works on featureless terrain and with images with little contrast or with high image noise. This means that the system can potentially be used in areas or under conditions where it would be difficult to map using other visual techniques.

Another advantage of the proposed method lies in the simplicity of the setup and the computations necessary to generate data products. Many underwater platforms use cameras to view the seafloor and have some form of navigational setup. Here we demonstrate that high-resolution data products can be generated using a low-end camera (1.4 megapixel resolution at 12 fps) and a standard navigational setup (DVL, compass and pressure sensor) to generate visually consistent data products. While navigational drift is problematic for any form of mapping survey, the results show that even in the presence of such drift, since the time is short between when bathymetric information and color information is recorded using our setup (on the order of seconds), the drift does not have any noticeable effect on the assignment of color information, and the transitions in shape and color remain consistent. The implication is that the algorithm can be implemented on a wide range of AUVs or ROVs by the simple addition of a line laser to produce outputs comparable to what has been shown in this work. The algorithm described is also computationally efficient, since all calculations are feed-forward, which allows the algorithm to handle large data sets. To improve the usability of the final data product, a method to decrease the number of bathymetry points is used to enable interactively visualization of large data sets. The resulting georeferenced imagery is suitable for further analysis, and in our future work we aim to extend the algorithm to incorporate measurements from other seafloor observation instruments such as subsurface acoustic measurement, acoustic reflection intensity, or the results of interpreted visual data to generate multimodal data products that can be used to aid the understanding of seafloor environments.

ACKNOWLEDGMENTS

The authors thank the Hyper-Dolphin team and the R/V Natsushima crew for their help and advice during the NT12-05 cruise. The project is funded by the Japanese Ministry of Education, Culture, Sports, Science and Technology under the Program for the development of fundamental tools for the utilization of marine resources.

REFERENCES

- Aggarwal, M. & Ahuja, N. (2002). A pupil-centric model of image formation. *International Journal of Computer Vision*, 48(3), 195–214.
- Agin, G. & Binford, T. (1976). Computer description of curved objects. *IEEE Transactions on Computers*, C-25(4), 439–449.
- Arfken, G. (1985). *Mathematical Methods for Physicists*. Academic Press, Orlando, FL, 3rd edition.
- Bodenmann, A., Thornton, B., Nakatani, T., & Ura, T. (2011). 3D colour reconstruction of a hydrothermally active area using an underwater robot. In *Proc. of MTS/IEEE Oceans*, Kona, HI.
- Bodenmann, A., Thornton, B., & Ura, T. (2010a). 3D mapping of the seafloor in color using a single camera. *Sea Technology*, 51(12), 51–53.
- Bodenmann, A., Thornton, B., Ura, T., & Sangekar, M. (2010b). Pixel mapping for generation of 3D coloured seafloor bathymetry using a single camera. In *Proceedings of AUV 2010*, Monterey, CA.
- Bouguet, J.-Y. (2004). Camera calibration toolbox for matlab.
- Bradski, G. (2000). The OpenCV Library. *Dr. Dobb's Journal of Software Tools*.
- Bruno, F., Bianco, G., Muzzupappa, M., Barone, S., & Rationale, A. (2011). Experimentation of structured light and stereo vision for underwater 3D reconstruction. *ISPRS Journal of Photogrammetry and Remote Sensing*, 66(4), 508–518.
- Dalgleish, F., Ouyang, B., & Vuorenkoski, A. (2013). A unified framework for image performance enhancement of extended range laser seabed survey sensors. In *Proceedings of the International Symposium on Underwater Technology*, Tokyo.
- Delaunay, B. (1934). Sur la sphere vide. *Izvestia Akademii Nauk SSSR, Otdelenie Matematicheskikh i Estestvennykh Nauk*, 7(6), 793–800.
- Duntley, S. (1963). Light in the sea. *Journal of the Optical Society of America*, 53(2), 214–233.
- Garcia, R., Campos, R., & Escartín, J. (2011). High-resolution 3D reconstruction of the seafloor for environmental monitoring and modelling. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 58–65), San Francisco.
- Garcia, R., & Gracias, N. (2011). Detection of interest points in turbid underwater images. In *Proceedings of MTS/IEEE Oceans*, Santander, Spain.
- Harsdorf, S., Janssen, M., Reuter, R., Toeneboen, S., Wachowicz, B., & Willkomm, R. (1999). Submarine lidar for seafloor inspection. *Measurement Science and Technology*, 10(12), 1178–1184.
- Hartley, R., & Zisserman, A. (2004). *Multiple view geometry in computer vision*, 2nd ed. Cambridge University Press.
- Inglis, G., Smart, C., Vaughn, I., & Roman, C. (2012). A pipeline for structured light bathymetric mapping. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 4425–4432), Algarve, Portugal.

- Izquierdo, M., Sanchez, M., Ibañez, A., & Ullate, L. (1999). Sub-pixel measurement of 3D surfaces by laser scanning. *Sensors and Actuators A: Physical*, 76(1-3), 1–8.
- Jaffe, J. S. (1990). Computer modeling and the design of optimal underwater imaging systems. *IEEE Journal of Oceanic Engineering*, 15(2), 101–111.
- Jakuba, M., & Yoerger, D. (2003). High-resolution multi-beam sonar mapping with the autonomous benthic explorer (ABE). In *Proceedings of the International Symposium on Unmanned Untethered Submersible Technology (UUST03)*, Durham, NH.
- Johnson-Roberson, M., Bryson, M., Friedman, A., Pizarro, O., Troni, G., Ozog, P., & Henderson, J. C. (2016). High resolution underwater robotic vision based mapping and three dimensional reconstruction for archaeology. *Journal of Field Robotics*.
- Johnson-Roberson, M., Pizarro, O., Williams, S., & Mahon, I. (2010). Generation and visualization of large-scale three-dimensional reconstructions from underwater robotic surveys. *Journal of Field Robotics*, 27(1), 21–51.
- Jordt-Sedlazeck, A., & Koch, R. (2012). Refractive calibration of underwater cameras. In *Proceedings of the 12th European Conference on Computer Vision*, vol. 7576 of *Lecture Notes in Computer Science* (pp. 846–859), Florence, Italy. Berlin-Heidelberg: Springer.
- Klepšvik, J., Torsen, H., & Thoresen, K. (1990). Laser imaging technology for subsea inspection: Principles and applications. In *Proceedings of Intervention Sous-Marine* (pp. 317–332), Toulon, France.
- Kocak, D., Caimi, F., Das, P., & Karson, J. (1999). A 3-D laser line scanner for outcrop scale studies of seafloor features. In *Proceedings of MTS/IEEE Oceans* (vol. 3, pp. 1105–1114), Seattle.
- Kondo, H., Maki, T., Ura, T., Nose, Y., Sakamaki, T., & Inaishi, M. (2004). Relative navigation of an autonomous underwater vehicle using a light-section profiling system. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (vol. 2, pp. 1103–1108), Sendai, Japan.
- Kunz, C., & Singh, H. (2013). Map building fusing acoustic and visual information using autonomous underwater vehicles. *Journal of Field Robotics*, 30(5), 763–783.
- Lindstrom, P., & Turk, G. (1998). Fast and memory efficient polygonal simplification. In *Proceedings of Visualization '98* (pp. 279–286), Research Triangle Park, NC.
- Marks, R., Rock, S., & Lee, M. (1995). Real-time video mosaicking of the ocean floor. *IEEE Journal of Oceanic Engineering*, 20(3), 229–241.
- McGlamery, B. L. (1979). A computer model for underwater camera systems. In *Proceedings of SPIE 0208, Ocean Optics VI* (pp. 221–231), Monterey, CA. International Society for Optics and Photonics.
- Moore, K. D., Jaffe, J. S., & Ochoa, B. L. (2000). Development of a new underwater bathymetric laser imaging system: L-bath. *Journal of Atmospheric and Oceanic Technology*, 17(8), 1106–1117.
- Mora, C., Tittensor, D., Adl, S., Simpson, A., & Worm, B. (2011). How many species are there on earth and in the ocean? *PLoS Biology*, 9(8), e1001127.
- Murez, Z., Treibitz, T., Ramamoorthi, R., & Kriegman, D. (2015). Photometric stereo in a scattering medium. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 3415–3423), Santiago, Chile.
- Narasimhan, S. G., Nayar, S. K., Bo, S., & Koppal, S. J. (2005). Structured light in scattering media. In *Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV'05)* (vol. 1, pp. 420–427), Beijing, China.
- Negahdaripour, S., & Madjidi, H. (2003). Stereovision imaging on submersible platforms for 3-D mapping of benthic habitats and sea-floor structures. *IEEE Journal of Oceanic Engineering*, 28(4), 625–650.
- Nicosevici, T., Gracias, N., Negahdaripour, S., & Garcia, R. (2009). Efficient three-dimensional scene modeling and mosaicing. *Journal of Field Robotics*, 26(10), 759–788.
- Palterson, R. (1975). Future developments in deep-sea imaging. *Oceanus*, 18(3), 17–23.
- Pizarro, O., Eustice, R., & Singh, H. (2009). Large area 3-d reconstructions from underwater optical surveys. *IEEE Journal of Oceanic Engineering*, 34(2), 150–169.
- Pizarro, O., & Singh, H. (2003). Toward large-area mosaicing for underwater scientific applications. *IEEE Journal of Oceanic Engineering*, 28(4), 651–672.
- Pope, R., & Fry, E. (1997). Absorption spectrum (380–700 nm) of pure water. II. Integrating cavity measurements. *Applied Optics*, 36(33), 8710–8723.
- Roman, C., Inglis, G., & Rutter, J. (2010). Application of structured light imaging for high resolution mapping of underwater archaeological sites. In *Proceedings of MTS/IEEE Oceans*, Sydney, Australia.
- Roman, C., Inglis, G., Vaughn, J., Smart, C., Dansereau, D., Bongiorno, D., Johnson-Roberson, M., & Bryson, M. (2013). New tools and methods for precision seafloor mapping. *Oceanography*, 26(1), 10–15.
- Roman, C., & Singh, H. (2007). A self-consistent bathymetric mapping algorithm. *Journal of Field Robotics*, 24(1-2), 23–50.
- Singh, H., Whitcomb, L., Yoerger, D., & Pizarro, O. (2000). Microbathymetric mapping from underwater vehicles in the deep ocean. *Computer Vision and Image Understanding*, 79(1), 143–161.
- Tetlow, S., & Spours, J. (1999). Three-dimensional measurement of underwater work sites using structured laser light. *Measurement Science and Technology*, 10(12), 1162–1167.
- The CGAL Project (2015). *CGAL User and Reference Manual*. CGAL Editorial Board, 4.6 ed.
- Thornton, B., Asada, A., Bodenmann, A., Sangekar, M., & Ura, T. (2013). Instruments and methods for acoustic and visual survey of manganese crusts. *IEEE Journal of Oceanic Engineering*, 38(1), 186–203.

- Treibitz, T., Schechner, Y. Y., Kunz, C., & Singh, H. (2012). Flat refractive geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(1), 51–65.
- Usui, A., Graham, I. J., Ditchburn, R. G., Zondervan, A., Shibasaki, H., & Hishida, H. (2007). Growth history and formation environments of ferromanganese deposits on the Philippine Sea Plate, northwest Pacific Ocean. *Island Arc*, 16(3), 420–430.
- Vapnik, V., & Lerner, A. (1963). Pattern recognition using generalized portrait method. *Automation and Remote Control*, 24(6), 774–780.
- Welch, S. (1993). Effects of window size and shape on accuracy of subpixel centroid estimation of target images. Nasa Technical Paper 3331.
- Zhang, H., & Negahdaripour, S. (2005). Fast and robust progressive stereo reconstruction by symmetry guided fusion. In *Proceedings of MTS/IEEE Oceans* (pp. 551–556), Brest, France.
- Zhang, Z. (2000). A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(11), 1330–1334.