Auto-calibration of line-laser structured-light seafloor mapping systems

David Stanley*, Adrian Bodenmann*, Miquel Massot-Campos*, Blair Thornton*†
*Centre for In Situ and Remote Intelligent Sensing, FEPS, University of Southampton, UK
†Institute of Industrial Science, The University of Tokyo, Japan

***Abstract*—The sub-centimetre resolution of line-laser generated maps exceeds acoustic bathymetry and provides detailed information needed for inspection and ecological research, but makes the maps more sensitive to errors stemming from (navigation and) mapping sensor calibration. The high resolution is achieved using a large relative baseline between the camera and laser apparatus. Their extrinsic parameters are often incidentally changed between missions when the vehicle they are attached to undergoes maintenance, making it desirable to have a means of calibrating these parameters in-situ and using natural scenes. Existing methods require the use of a second camera, which increases the complexity of the mapping hardware and often requires dedicated calibration manoeuvres to be performed. In this paper we present a novel in-situ calibration method using natural scenes that does not require the use of a second camera, but instead uses the ordinary data captured during a survey mission to perform a non-linear least squares optimisation that reduces inconsistencies between overlapping measurements. We take advantage of a dual line-laser setup to generate continuously overlapping measurements for calibration that have inherently small uncertainty in their relative pose estimates and demonstrate reduced inconsistency between overlapping data using our optimised calibration parameters over that of CAD-derived values.**

***Index Terms*—seafloor mapping, structured-light, featureless, calibration, error metric, underwater robotics**

I. Introduction

For the purposes of monitoring protected marine habitats; planning seafloor infrastructure; and surveying areas of archaeological interest, bathymetric maps are increasingly being made by 3D mapping systems attached to uncrewed underwater vehicles. While traditional acoustic systems, like multibeam sonar, can be used to measure bathymetry at a resolution of centimetres (typically tens of centimetres), optical mapping systems can be used to measure bathymetry at millimetre resolutions [1]. The accuracy of these measurements is typically limited by errors in positional estimates of the vehicle, but for especially high resolution systems; like those using a structured light approach; errors in the calibration of sensor models can also be significantly limiting [2].

Structured light mapping systems are formed of a camera and a means of projecting a known pattern of light onto the terrain. In captured images the pattern appears distorted by the terrain's topography and the relative perspective of the camera. Knowing the projection of the pattern, the positions of the terrain it is incident on can then be inferred. Line-laser systems are a type of structured light mapping system whose pattern is a line, projected onto the terrain by a laser that emits a sheet of light [1]. By orientating the line perpendicular to the direction of travel, and with the laser positioned along that axis of travel relative to the camera, the projected line shows the terrain in relief when imaged by the camera and measurements of the terrain across a swathe are generated, as illustrated in Fig. 1. A patch of terrain is then mapped in a push-broom fashion as the vehicle the system is attached to advances.

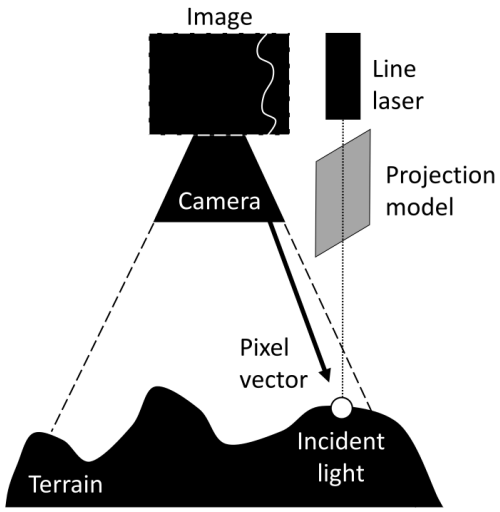

Fig. 1. The principle of line-laser mapping. Variations in the terrain that the laser projects a line across appear in relief within the images captured by the camera. By tracing back the vector of light rays entering a pixel within the camera's image sensor, its intersection with the laser's projection model in 3D space can be calculated. That intersection point is coincident with the surface of the terrain.

In [3] it's noted that having too large a baseline between the laser and the camera in a line-laser system will prevent terrain below a minimum altitude from being measured, but that it is otherwise desirable to maximise the baseline between laser and camera to maximise the resolution of any measurements made. Examples of setups in literature typically use baselines between 0.5 m and 1 m [1], [2], [4]. This long baseline means that the camera and laser elements are usually built as separate apparatuses, rather than a single unit, and attached to separate points on the host vehicle. If the extrinsic parameters

Funded by UKRI NERC BioCam NE/P020887/1 and EPSRC CDT-SIS

defining their relative positions and orientations are to be calibrated, therefore, this would have to be done while the system is attached to the vehicle. One way of achieving this is to calibrate the setup using data gathered in-situ during field deployments, with groups investigating ways to achieve this without the need to deploy and observe known calibration targets, instead using natural scenes.

In [4] a dual camera setup is used to identify features in observed laser lines and match them between pairs of simultaneously captured images. By exploiting stereo geometry, 3D positions of the incident laser light can be measured, and by doing this at different ranges enough points can be generated to determine the position and orientation of the projected laser sheet relative to the camera. As this method is feature-based, it is reliant on the terrain being of sufficient complexity. In [2] a dual camera setup is also used but with both cameras and the laser positioned along a single axis. This will typically allows points on the laser line in one image to be matched with respective points on the laser line in the other using the stereo geometry alone. In this way data correspondences can be formed between image pairs and 3D measurements generated without having to identify features in the laser line. The position and orientation of the projected sheet of light relative to the camera can then be calculated in the same fashion as the first method.

Using a second camera to capture simultaneous images of laser lines requires additional electrical power, additional computation, and additional data storage. In practice, therefore, data for calibration using these existing methods would be generated by performing a calibration manoeuvre or mission separate to an actual survey so that the time spent capturing simultaneous images with the second camera is minimised. Performing such a mission or manoeuvre may mean more ship-time is required, increasing monetary cost. Additionally, there is a risk that despite good data on an associated survey having been gathered, the calibration mission or manoeuvre is not successful and needs to be repeated. At best this would require even more ship-time to be spent on performing another calibration mission, but at worst this could simply leave the user with no good data for calibrating the system. A method that can be used in-situ with natural scenes and that doesn't require a second camera, instead using the ordinary survey data itself, is therefore desirable.

In this work we present such a method, based on the concept of bundle adjustment [5], for calibrating the extrinsic parameters of line-laser systems. We use a non-linear least squares approach to minimise inconsistency between overlapping transects of measurements by adjusting the values of extrinsic parameters. We use an error metric based in the images' frame of reference, for measuring inconsistency between transects, that is described in our previous work [6] and, via a dual laser setup, use overlapping transects captured mere seconds and metres apart whose error in positional estimates is expected to be negligible and whose identification is trivial. However, our method can still be used for a single laser setup whenever pose estimates are highly accurate (or have been corrected to be, for instance by SLAM) or else can be used alongside SLAM in an expanded bundle adjustment. We demonstrate that our method converges to broadly similar parameter values using a range of different overlapping pairs of transects collected from a single survey, improving on the results of our previous work [6]. We also demonstrate significantly reduced comparative depth error between transects using optimised parameters over that of CAD-derived values.

II. Method

Our mapping system uses two line-lasers positioned fore and aft of a monochrome camera, all pointed downwards, with each laser projecting a line of light normal to the direction of travel onto the terrain below. The laser lines projected onto the terrain appear in relief within images captured by the camera, with the fore and aft laser lines appearing in the top and bottom halves of the each image, respectively, as depicted in the top of Fig. 2. Their projection of light is modelled as a flat plane which intersects with both the laser emitter and the terrain. This is a commonly used laser model for line-laser mapping systems (for example in [1], [3], [4]), but we note that it does reduce the extrinsic parameter space by half, as the roll of the laser and its positional offset along axes other than that normal to the plane all fall within the plane itself and so are of no consequence. Non-planar models exist in literature, such as the elliptical cone model used in [7] for a flat viewport which would require all 6 extrinsic parameters to be estimated, but we expect our method to still work with these, although more data from a more diverse array of overlapping pairs of transects may be required for good calibration. The interaction of light with the camera apparatus, including view-port, is modelled as a pin-hole camera plus radial and tangential distortion. By identifying the position of the laser lines within images, 3D measurements of the terrain can then be calculated; as described in [3].

We aim to find the correct yaw, pitch, and distance-from-camera (DFC) values of our laser planes; shown inset in Fig. 2; to eliminate discrepancies between the measurements of the terrain that each laser independently generates. To do so we minimise a cost function that indirectly penalises such discrepancies, namely the error metric introduced in our previous work [6] and illustrated in Fig. 3. The observations of each laser form two sets of measurements that are each independently used in combination with their estimated parameters to reconstruct the terrain. Here we reconstruct the terrain as a triangulated irregular network (TIN) as this is a simple means of creating a continuous surface. We use Delaunay triangulation in the global lateral plane do so. Each laser's reconstruction is then used to predict the observations of the other respective laser by simulating the intersection of the laser sheet with the reconstructed terrain. The vertical offsets between the observed and predicted laser lines within images form residuals. The total error for a pair of overlapping transects is then the root mean square of image root mean square residuals, as detailed in (1) and (2) where r_i is an

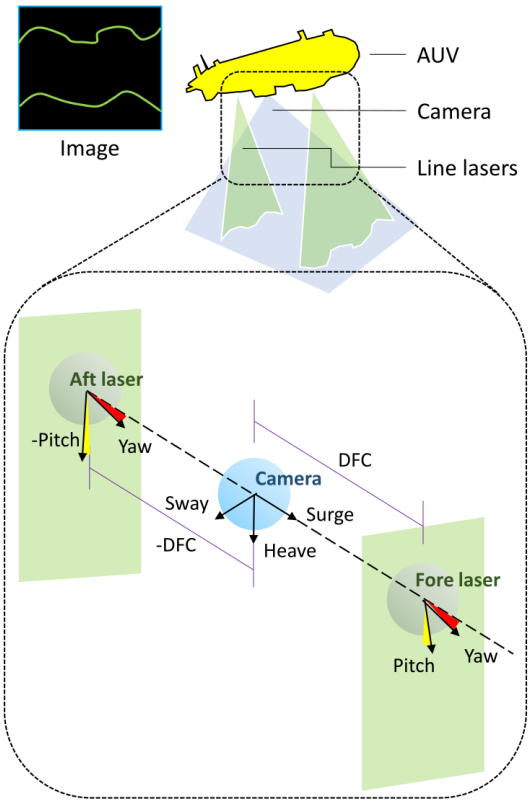

Fig. 2. A dual line-laser setup and its extrinsic laser parameters. Showing the setup (top) and extrinsic parameters (inset). Two line-lasers project sheets of light. These appear as lines of light that appear in the top and bottom half of images captured by a camera placed between the lasers. The sheets of light are modelled as planes. The directions of each laser's parameters are all with reference to the camera's frame of reference in the same way; the aft laser moving or its projection pitch towards the camera are therefore positive changes whereas they are negative for the fore laser. DFC is defined along the surge axis of the camera. Pitch is rotational about the sway axis. Yaw is rotational about the heave axis.

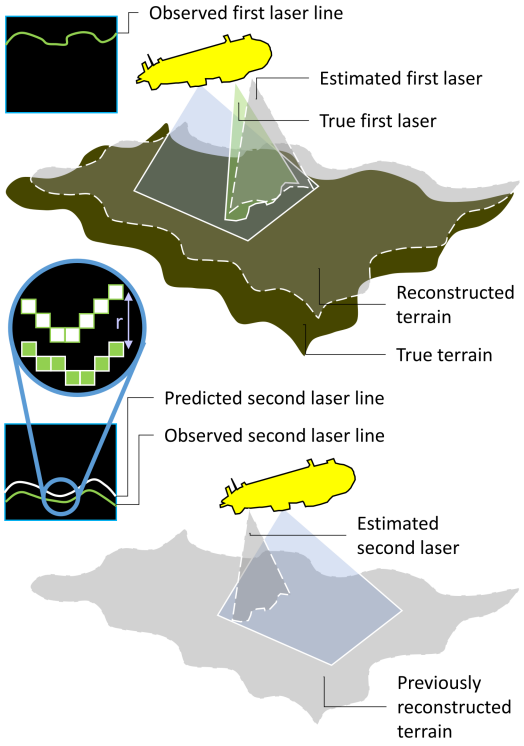

Fig. 3. Illustration of how residuals for use in the metric of inconsistency are calculated. The AUV carrying a line-laser mapping system moves along above terrain, generating bathymetric measurements of an area that are used to make a reconstruction (top). The AUV then passes over the same area again in a different direction or with a different laser. The new observations of the system are predicted and compared with the actual observations made (bottom), with the vertical offset between the laser lines in each pixel column forming residuals (r, inset). This is repeated with the second set of observations used to predict the first set of observations to generate all residuals.

individual pixel column's residual of image j within the transect and ϵ is the error.

$$\epsilon_{\text{img}} = \sqrt{\overline{r_i{}^2}} \tag{1}$$

$$\epsilon_{\text{total}} = \sqrt{\overline{\epsilon_{\text{img},j}{}^2}} \tag{2}$$

Improving on the sequential and coupled method used in our previous work which heavily restricted the parameter space and looked only at pitch and yaw [6], here we minimise the total error of an overlapping pair of transects using the Trust Region Reflective algorithm [8] with the Jacobian numerically estimated at each step, placing no restrictions on the parameter space and looking at pitch, yaw, and DFC. The yaw and pitch are defined in degrees, the DFC in metres, and the error in pixels.

With our dual laser setup, spatially overlapping transects are continuously generated throughout a survey as the vehicle moves forward and the aft laser sweeps across terrain that has just been swept by the fore laser. The CAD-derived value for the distance between laser observations is a mere 3.194 m. Due to the short time-frame involved between one laser measuring an area and the next, dead-reckoning estimates of the vehicle's relative position between overlapping measurements should have negligible error. Also, minimal effort is needed to identify overlapping transects; any period of time within the survey that is greater than that needed for the aft laser to catch up with the fore laser (which is dependent on the vehicle's speed) will yield an overlapping pair of transects.

III. Experiments and Results

For our experiments we used data gathered during the 2019 DY108-109 research expedition aboard the RRS Discovery to the Darwin Mounds Marine Protected Area in the North Atlantic, covering 29 hectares at millimetre-order resolution [9]. This is an area of mounds of cold-water coral situated at around 1,000 m depth and was surveyed using the University of Southampton's BioCam mapping system attached to the National Oceanography Centre's autonomous underwater vehicle, Autosub6000. Specifically, we used small subsets of this data; a total of 7 short transects of differing terrain variation and complexity, whose variance in ranging measurements is detailed in Table I. These vary in size between roughly 45 m^2

TABLE I
CHARACTER OF TRANSECTS

Transect	Standard deviation of ranging measurements (m)
A	0.493
B	0.284
C	0.422
D	0.362
E	0.119
F	0.387
G	0.045

TABLE II
INTENDED (CAD-DERIVED) VALUES FOR EXTRINSIC PARAMETERS

Parameter	Fore Laser Value	Aft Laser Value
Pitch (°)	0	0
Yaw (°)	0	0
DFC (m)	1.687	-1.507

and 80 m^2 and in the nature of their terrain; covering smooth sand, the beginning of a coral mound, and an embedded whale carcass. We optimised the pitch, yaw, and DFC of both lasers using each transect separately so as to test the robustness of our method using different data. The intended values for the extrinsic parameters during assembly (CAD-derived) are detailed in Table II and were used to initialise the optimisations.

The results of these optimisations are shown in Fig. 4. They show each transect converging to largely similar parameter values, improving on the results of our previous work where this was not the case [6]. However, for these optimisations of individual small transects, they have not converged to exactly the same results. There remains large variation between converged parameter values relative to the difference between the initial values and the mean of the converged values. This variation suggests that the nature of each transect (a combination of the complexity and shape of the terrain and the vehicle's path above it) is biasing the error space such that the optimisation does not converge to globally optimal values. Consequently, the robustness of this method to different types of transects and also to different initial conditions needs to be investigated in future work. It is also expected that using larger transects or combining transects would reduce the variance in optimised parameter values as small-scale biasing effects should then average out.

The optimal parameters for transect A, which had the greatest standard deviation in ranging measurements, have been applied to reconstruct each other transect and the depth errors for these transects, between the reconstruction of one laser and the other, have been calculated. The results are detailed in Table III and plotted in Fig. 5 and show a roughly 50% reduction in depth error from initial values.

IV. CONCLUSIONS

Although, as previously discussed, further work is needed to investigate the robustness of this method; based on the results shown so far, we can conclude that:

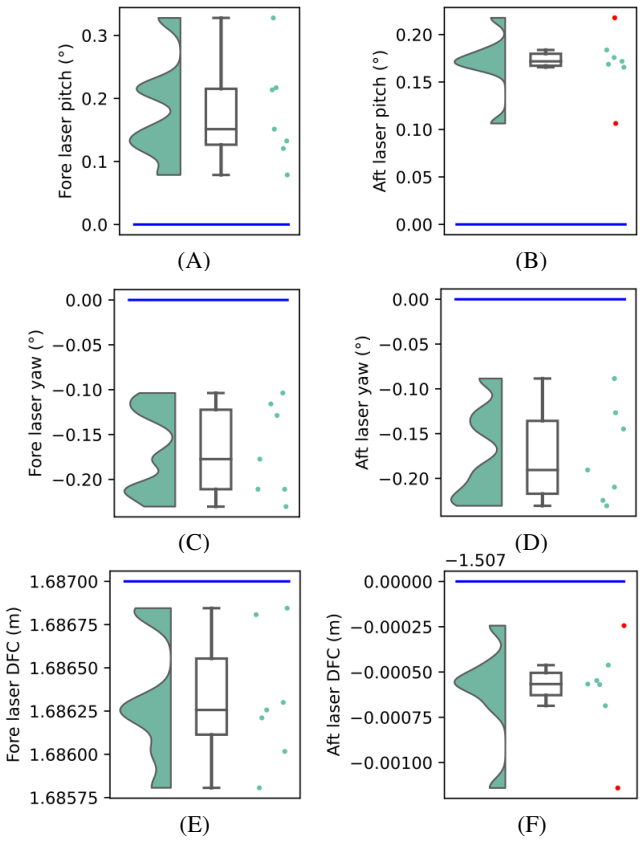

Fig. 4. Optimised extrinsic parameter values. Showing values arrived at using different transects (green), the initial parameter value (blue), and with boxplot outliers (red). Showing fore laser pitch (A), aft laser pitch (B), fore laser yaw (C), aft laser yaw (D), fore laser DFC (E), and aft laser DFC (F).

TABLE III
APPLICATION OF TRANSECT A CALIBRATION

Transect	Initial RMS depth error (m)	Resultant RMS depth error (m)
A	0.064	0.038
B	0.063	0.044
C	0.053	0.033
D	0.085	0.047
E	0.098	0.041
F	0.295	0.103
G	0.129	0.058

- A bundle adjustment-like approach using the featureless error metric from [6] can be used to auto-calibrate the pitch, yaw, and distance-from-camera of both lasers in a dual line-laser bathymetric mapping system. The proposed method can be used in-situ and, unlike [2], [4], does not require a second camera or a calibration manoeuvre to be performed. We demonstrated that the method converges to broadly the same solution when applied to a variety of transects with different types of terrain. It could also be used to calibrate single laser systems.
- A dual laser setup allows for auto-calibration, such as is demonstrated here, to be performed without having to first

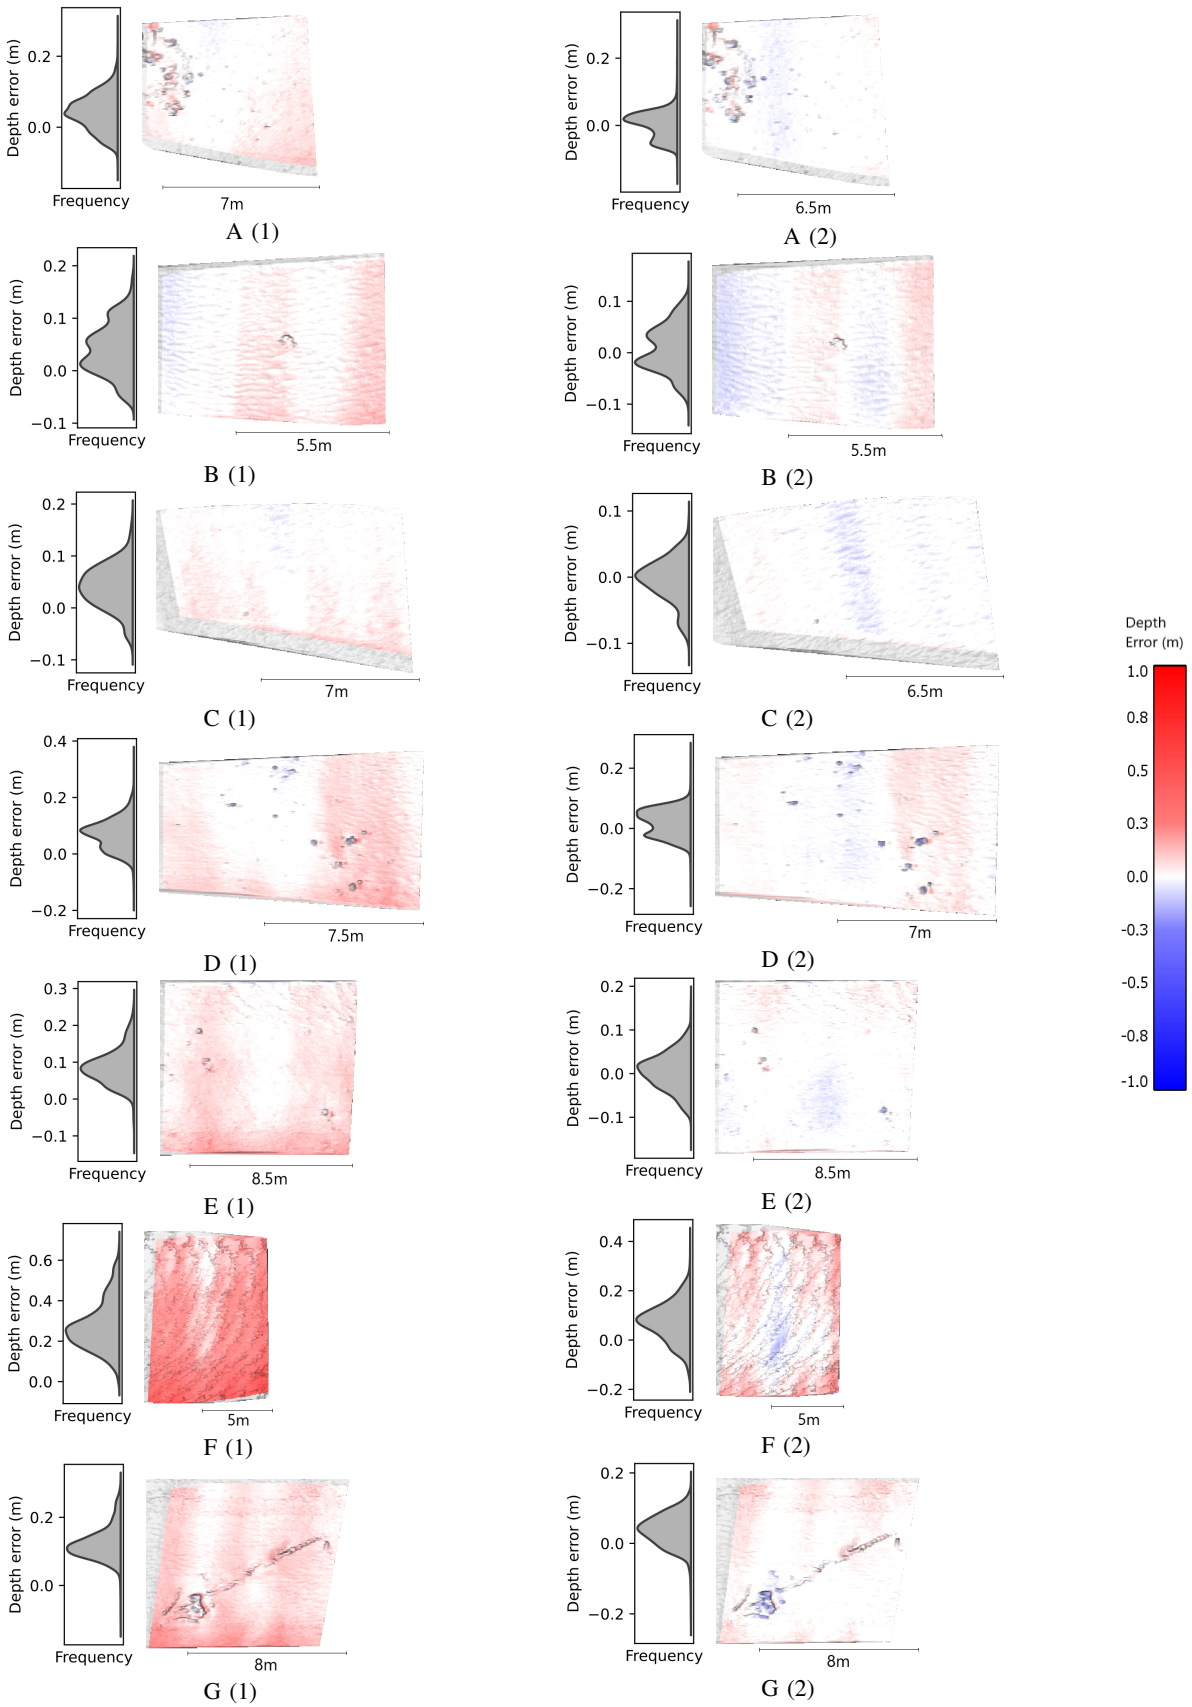

Fig. 5. The depth error of transects. (A-C). Showing initial depth error (1, left) and depth error when extrinsic parameters optimised using transect A are applied (2, right).

reduce errors in positional estimates through SLAM or other means, and with minimal effort required to identify transects of spatially overlapping data.

REFERENCES

[1] C. Roman, G. Inglis and J. Rutter, "Application of structured light imaging for high resolution mapping of underwater archaeological sites," OCEANS'10 IEEE SYDNEY, 2010, doi: 10.1109/OCEANSSYD.2010.5603672.

[2] M. Leat, A. Bodenmann, M. Massot-Campos, and B. Thornton, "Analysis of uncertainty in laser-scanned bathymetric maps," 2018 IEEE/OES Autonomous Underwater Vehicle Workshop (AUV), Porto, Portugal, 2018, pp. 1–7, doi: 10.1109/AUV.2018.8729747.

[3] A. Bodenmann, B. Thornton, and T. Ura, "Generation of high-resolution three-dimensional reconstructions of the seafloor in color using a single camera and structured light," J. Field Robotics, 2017, 34: 833–851. doi:10.1002/rob.21682.

[4] G. Inglis, C. Smart, I. Vaughn and C. Roman, "A pipeline for structured light bathymetric mapping," 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, Vilamoura-Algarve, Portugal, 2012, pp. 4425–4432, doi: 10.1109/IROS.2012.6386038.

[5] B. Triggs, P.F. McLauchlan, R.I. Hartley, and A.W. Fitzgibbon, "Bundle adjustment — a modern synthesis," IWVA 1999: Vision Algorithms: Theory and Practice, 2000, pp. 298–372.

[6] D. Stanley, A. Bodenmann, M. Massot-Campos and B. Thornton, "A featureless approach to improve self-consistency in structured light bathymetry," 2020 IEEE/OES Autonomous Underwater Vehicles Symposium (AUV), St. Johns, NL, Canada, 2020, pp. 1–6, doi: 10.1109/AUV50043.2020.9267891.

[7] A. Palomer, P. Ridao, J. Forest and D. Ribas, "Underwater laser scanner: Ray-based model and calibration," in IEEE/ASME Transactions on Mechatronics, vol. 24, no. 5, pp. 1986-1997, Oct. 2019, doi: 10.1109/TMECH.2019.2929652.

[8] M.A. Branch, T.F. Coleman, and Y. Li, "A subspace, interior, and conjugate gradient method for large-scale bound-constrained minimization problems," SIAM J. Scientific Computing, 1999, 21 (1): 1-23. doi:10.1137/S1064827595289108.

[9] V. Huvenne, B. Thornton, "Cruise report no. 66 RRS Discovery cruise DY108-109," 224p.