A Featureless Approach to Improve Self-Consistency in Structured Light Bathymetry

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*—This paper describes a novel method for calibrating structured light setups using a featureless approach to quantify and improve self-consistency in bathymetric maps. The self-consistency and accuracy of seafloor reconstructions and information derived from them are limited by uncertainties in vehicle localisation, sensor models and their calibration. For high-resolution setups such as structured light, these uncertainties can be several orders of magnitude larger than the resolution of the maps generated. Although techniques such as simultaneous localisation and mapping and bundle adjustment can correct pose estimates and sensor calibrations to improve map self-consistency, both methods typically rely on finding and matching features in the data, which limits their application to structured light since a key advantage of this method is that it does not rely on seafloor features to be present in order to work. In this paper, we develop a fully featureless approach to improve self-consistency in structured light setups. Simulations are performed to validate the proposed method, and we analyse data that was collected using the Autosub6000 autonomous underwater vehicle equipped with the BioCam seafloor mapping instrument at a depth of approximately 1000m in the Darwin Mounds UK marine protected area. The results for independent parameter optimisation demonstrate that the fully featureless approach can converge towards optimal calibrations.**

I. Introduction

3D reconstructions of the seafloor have applications ranging from planning and inspection of seafloor infrastructure to monitoring of marine habitats. The accuracy of these maps is limited by uncertainty in both the estimated position and orientation (pose) of the vehicles used to gather data, and in the modelling and calibration of their sensors. For high-resolution methods such as structured light laser bathymetry [1], [2], these uncertainties can be several orders of magnitude greater than the resolution of the generated maps, limiting confidence in the reconstructions and any derived information.

Simultaneous localisation and mapping (SLAM) techniques can be used to reduce errors by comparing how consistent reobservations of areas in the map are with one another and adjusting vehicle pose estimates accordingly. SLAM approaches can be further augmented by bundle adjustments (BA) that optimise the calibration of sensor-models and finely adjust vehicle pose to improve the self-consistency of overlapping observations. Such approaches often rely on matching seafloor features between overlapping data. However, distinguishable features in seafloor terrain are not guaranteed to exist and are often sparsely distributed, such that only a small proportion of the data gathered can inform optimisation of the map.

Corresponding author: d.stanley@soton.ac.uk

Featureless methods can make use of every overlapping point in observations to inform optimisation. This kind of dense optimisation can potentially achieve greater global agreement across the generated map. Featureless approaches are also more robust as they are less dependent on the nature of the underlying terrain. Featureless approaches to optimise seafloor mapping observations were demonstrated by Roman and Singh in 2007 [3], who used a sub-map approach combined with an extended Kalman filter for acoustic multi-beam surveys. Barkby et al., [4] developed a particle-filter based approach that did not require sub-maps, and applied the method to multi-beam setups. Both of these approaches have since been applied to structured light by Massot et al., [5], [6]. In [3], [5], [6] map self-consistency was optimised using error metrics in the map reference frame. A limitation of this approach is that the transformation of observations from the sensor to the mapping frame also contains uncertainties, which makes it difficult to validate the accuracy of the final maps without a separate ground truth. In [4], previous observations were projected from the map to the sensor reference frame to allow for direct comparison with the sensor's measurements during reobservation. These direct measurements are known not to contain these errors since no transformations have been made, increasing confidence in solutions found, especially for featureless approaches where feature extraction and matching are not needed. For acoustic measurements, this was achieved using polar coordinates consisting of range and beam angle. However, this cannot be used for structured light measurements that make use of geometric relationships between a laser projection and a 2D camera detector.

The method developed in this work performs featureless

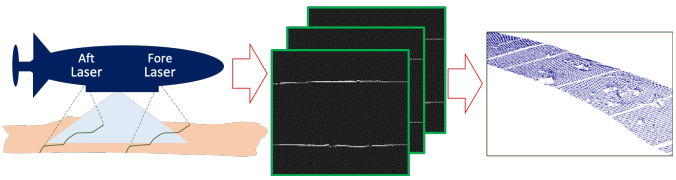

Fig. 1: The mapping process. A camera and laser equipped AUV takes images of laser projections on the seafloor. The camera's pose is estimated for each image and with this a point-cloud of the bathymetry is generated, then meshed.

optimisation using an error metric defined in the 2D detector frame. This is achieved by projecting previously observed areas of a map into the image frame for direct comparison with the new observations of the laser projection. Our setup employs two downwards pointing sheet lasers, one fore and one aft of a downwards pointing camera mounted on an autonomous underwater vehicle (AUV), as depicted in Fig. 1. As the AUV proceeds, the fore and aft lasers project lines onto the scene below that appear in the top and bottom halves, respectively, of images captured by the camera. This is modelled as two planes intersecting with the surface of the terrain and the intersection of these with the camera pixel vectors after passing through a lens. The laser lines projected onto the seafloor are identified in the images, and their 2D positions in the images are combined with estimates of the vehicle pose and sensor model parameters to generate a 3D point cloud, which is then meshed as described in [2]. However, errors in the calibration of the laser planes can lead to mapping errors [7] and inconsistencies in reobserved bathymetry.

II. Method

We propose a novel method to improve the reconstruction consistency of a dual laser stripe setup by comparing their seafloor observations. The fore laser is used to reconstruct the seafloor using a laser stripe reconstruction pipeline [2]. This terrain is then intersected with the aft laser plane, and the intersection projected onto the image frame. This reprojection and the detected laser lines are compared, and the laser orientation is optimised to minimise the distance between them. Although the method is developed for a dual laser setup, it is equally applicable to single laser setups when areas of the seafloor are reobserved during loop closures.

Differences between real and reprojected structured light observations, such as observations of laser lines projected using sheet lasers, arise where there are inconsistencies between overlapping data. This is caused by errors in: (1) the individual vehicle pose estimates, (2) the calibration of sensor models, the accuracy of those models, and (3) the extraction of laser line positions from images. For transects where the AUV proceeds along a largely constant heading, the dual laser setup considered here means that almost all points observed using the fore laser are reobserved with the aft laser after the AUV has travelled a few metres distance. It is therefore reasonable to assume that the relative pose error caused by drift in dead reckoning based navigation solutions is negligible, and inconsistencies between bathymetry observations will be largely due to sensor model and calibration inaccuracies. We validate our method with simulated data for varying conditions, before applying it to different transects of AUV survey data.

A. Reconstructing observations

Observations of the fore laser are used to reconstruct the terrain in the form of a meshed surface. For the given timestamp of an image captured by the camera, the aft laser is modelled as a plane intersecting the mesh, the parameters for which are derived from the calibration parameters of the laser and the estimated pose of the AUV for that timestamp. These points of intersection are then projected into the image frame of the camera using the camera's calibration parameters and the same estimated pose of the AUV. They are interpolated for each row-wise position of a pixel column to reconstruct the extracted positions of the aft laser line within the image.

B. Optimisation conditions

In the method used here, a sequential approach is taken to optimising parameters of the laser models. Each sheet laser is modelled as a plane that is roughly perpendicular to the surge axis of the AUV, crossing that axis at a given distance from the camera. The orientation of the laser planes is modelled by their normal vector's respective pitch and yaw-offsets from the surge axis. These two angles are optimised in turn by the proposed algorithm. To avoid large changes being made to the orientation of one laser and little to no change being made to the other, when the reverse could produce similar outcomes in reduction of error, changes made to the orientations of the lasers are coupled. Changes made to the pitch of one laser are copied exactly for the other, largely restricting the effects of such a change to raising the height of the observed terrain for one laser whilst lowering it for the other (due to the lasers being positioned on opposite sides of the camera's field of view). Changes made to the yaw of one laser are reversed in direction for the other but the magnitude maintained.

C. Bisection search

For the parameter of the laser planes that is being optimised, an iterative, bisection-type search technique is applied to find a local minimum of error. A transect's error is calculated as the root-mean-of-mean-square (RMMS) residual, where a residual is the columnar pixel distance between a real and a reconstructed position of the observed laser line (see Fig.2):

$$\epsilon_{transect} = \sqrt{\overline{\epsilon_{image}^2}} \quad \text{where} \quad \epsilon_{image} = \sqrt{\overline{r^2}} \tag{1}$$

Initial parameter values for the lasers are estimated using the method described in [7]. Bounds for possible deviation from these values are set and the transect's error is calculated for these three values (the lower bound, zero deviation, and the upper bound).

For each iteration of the technique, the error is calculated for the two midway points between the three values and the errors of the resultant five values compared. The deviation value that gives the lowest error is selected as well as the next nearest values, forming a new central value with bounds for the next iteration step. This is described in more detail by Alg. 1. The optimisation is stopped when a maximum number of iterations has been carried out or when a desired tolerance has been reached.

D. Simulating data

To validate the method we simulate data with ground-truth sensor parameters. We then apply it to the simulated data

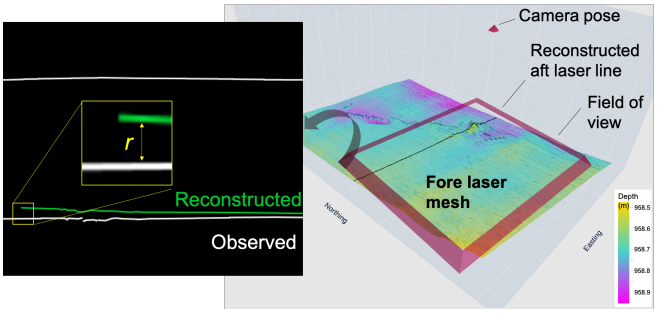

Fig. 2: The figure shows an area of seafloor that was reconstructed using the fore laser. The aft laser reobserves the same area seconds later. The image compares the simulated laser line (green) that is computed by reprojecting the intersection of the fore laser reconstruction using the camera and the aft laser plane model. The distance between this line and the line that was actually observed (white) forms the error.

Algorithm 1 Bisection search

procedure OPTIMISE(x_0, x_2, x_4)
  $y_0, y_2, y_4 = \epsilon(x_i)$, ▷ $i = 0, 2, 4$
  while tol(x_0, x_4) >min **and** count<max **do**
    $x_1, x_3 = (x_{i-1} + x_{i+1})/2$, ▷ $i = 1, 3$
    $y_1, y_3 = \epsilon(x_i)$, ▷ $i = 1, 3$
    $j \leftarrow min(y_j)$, ▷ $j = 0, 1, 2, 3, 4$
    if $j = 0$ **then**
      $x_0, x_2, x_4 \leftarrow x_0, x_1, x_2$
      $y_0, y_2, y_4 \leftarrow y_0, y_1, y_2$
    else if $j = 4$ **then**
      $x_0, x_2, x_4 \leftarrow x_2, x_3, x_4$
      $y_0, y_2, y_4 \leftarrow y_2, y_3, y_4$
    else
      $x_0, x_2, x_4 \leftarrow x_{j-1}, x_j, x_{j+1}$
      $y_0, y_2, y_4 \leftarrow y_{j-1}, y_j, y_{j+1}$
    end if
    $count \leftarrow count + 1$
  end while
  return x_j ▷ The optimal value
end procedure

with initial laser parameters offset from the known, ground-truth values. This allows us to validate that, under appropriate conditions, the method used here converges towards the true solution. It also allows us to investigate the effect of different initial conditions on convergence.

The data is simulated by taking a transect of real survey data and setting the modelled laser orientations to be exactly downwards-facing. We then generate a mesh using real observations of the fore laser and use it to reconstruct observations of the aft laser, as we would in our optimisation process. These reconstructed observations of the aft laser are then treated as real observations along with the genuinely real observations of the fore laser, and the set sensor parameters act as ground-truth parameters for this data. In this manner we produce simulated data that is entirely self-consistent and which closely reflects the nature of real terrain.

III. VALIDATION

In this section we detail the results of optimising simulated data initialised with sensor parameters offset from the ground truth in different ways. The data itself is simulated from a short, fairly plain transect of the same survey data used in IV. Table I details the four different initialisation conditions as well as the ground-truth. Bounds were set for the initial values as ± 3 deg, the maximum number of iterations was set to 7, and the desired tolerance was set to 10^{-5} deg.

Optimal conditions are where the initial orientations of the lasers have been offset from the ground-truth in the same coupled fashion as the optimisation process attempts to converge back towards it. This means that it is mathematically possible for the process to arrive at the true pitch and the true yaw values. Pitch-optimal and yaw-optimal conditions are where the initial orientations of the lasers have only been offset from the ground-truth in same the coupled manner as the optimisation process in the stated axis of rotation — the orientation has still been offset from the ground-truth in the other axis of rotation, but not in a manner that mathematically allows for the optimisation process to arrive at the true values for both lasers. Non-optimal conditions are where the initial orientations of the lasers have been offset in both axes such that the optimisation process cannot mathematically arrive at true values for both lasers in either axis.

Fig. 3 shows the optimisation for the four initial conditions. Fig. 4 shows the effect on the consistency of the two meshes that are separately produced from observations of the fore laser and the aft laser, respectively. Table II details the optimal orientation values arrived at by the optimisation process. The results show a smooth error gradient in all instances. For optimal initial conditions, the optimisation process converges towards the true solution where it is mathematically possible. However, if either the initial pitch or yaw does not satisfy the

TABLE I: Initial Laser Orientations for Simulated Data

Simulation Name	Fore Laser		Aft Laser	
	Pitch (deg)	*Yaw (deg)*	*Pitch (deg)*	*Yaw (deg)*
Optimal	1.0	-1.0	1.0	1.0
Pitch-optimal	1.0	2.0	1.0	1.0
Yaw-optimal	1.0	-1.0	-2.0	1.0
Non-optimal	1.0	2.0	-2.0	1.0
(Ground-truth)	(0.0)	(0.0)	(0.0)	(0.0)

TABLE II: Optimised Laser Orientations for Simulated Data

Simulation Name	Fore Laser		Aft Laser	
	Pitch (deg)	*Yaw (deg)*	*Pitch (deg)*	*Yaw (deg)*
Optimal	-0.0078	-0.0625	-0.0078	0.0625
Pitch-optimal	-0.8750	5.0000*	-0.8750	-2.0000*
Yaw-optimal	-0.3359	1.7656	-3.3359	-1.7656
Non-optimal	0.6250	5.0000*	-2.3750	-2.0000*
(Ground-truth)	(0.0000)	(0.0000)	(0.0000)	(0.0000)

*Boundary value.

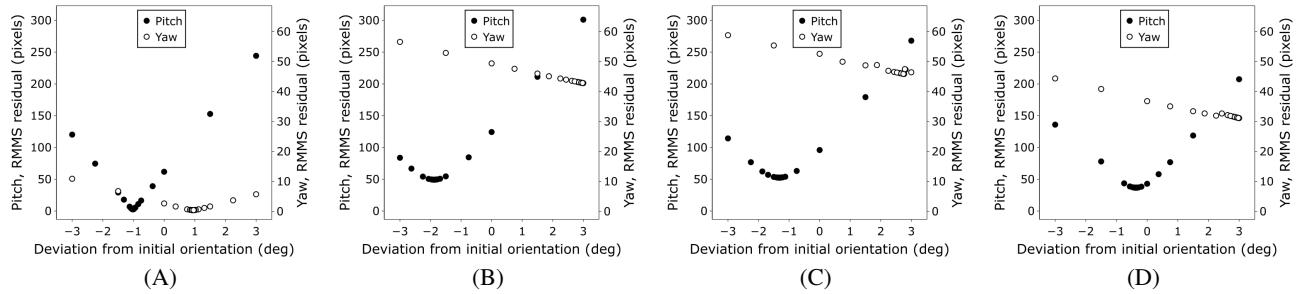

Fig. 3: Plots of the optimisation of simulated data. Showing optimisation of data with optimal (A), pitch-optimal (B), yaw-optimal (C), and non-optimal (D) offset calibration parameters. Deviation from initial orientation refers to that of the fore laser. Deviation of the aft laser can be inferred as the same as that of the fore laser in pitch and the negative in yaw.

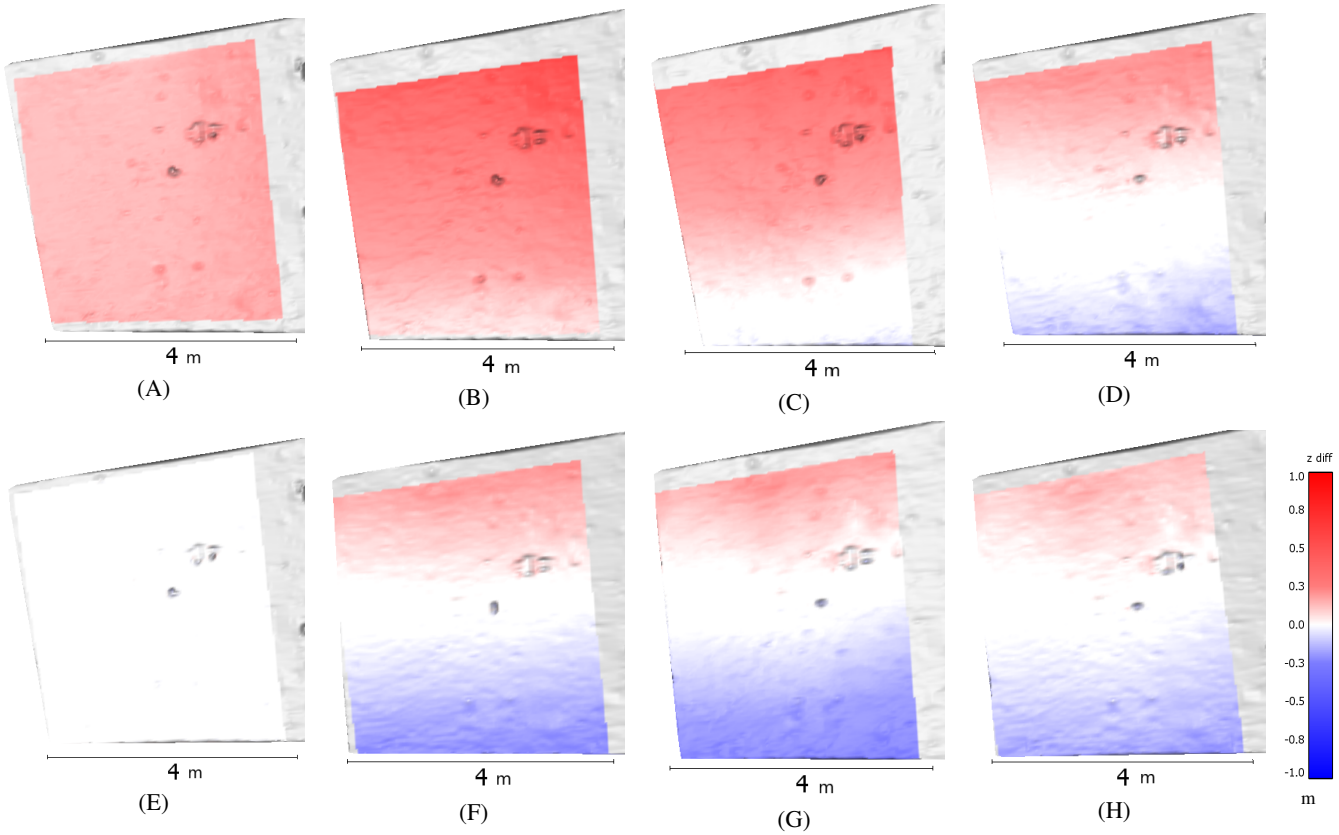

Fig. 4: Plots of the initial depth error (top row) and resultant depth error (bottom row) between the fore and aft laser meshes for the optimisation of the simulated data. Showing the optimal (A, E), pitch-optimal (B, F), yaw-optimal (C, G), and non-optimal (D, H) conditions. Depth error shown as the colouration of the fore laser mesh. Areas of no overlap between the two meshes are coloured grey.

optimal condition, it is not possible to convergence towards the true solution, even for the parameter in which it is mathematically possible to do so.

The error curve for the optimal conditions appears to be notably sharper than that of the other conditions. Any intricacies of the terrain used do not seem to have had a large effect on the smoothness of the error gradients. The terrain used to generate the simulated data is not particularly symmetric, however, and more symmetric terrain may reduce smoothness in the form of local, false minima. For the sub-optimal conditions, in all but the fully non-optimal case, the optimisation process led to at least slight reduction of depth error between the fore and aft laser meshes. Optimal conditions, for which the optimisation had resulted in values very close to the ground truth, had led to almost complete reduction of depth error. This small residual is due to the limit set on the number of iterations for optimisation used.

IV. Results and Discussion

In this section, we detail the results of applying the optimisation process to real data, gathered during the DY108-

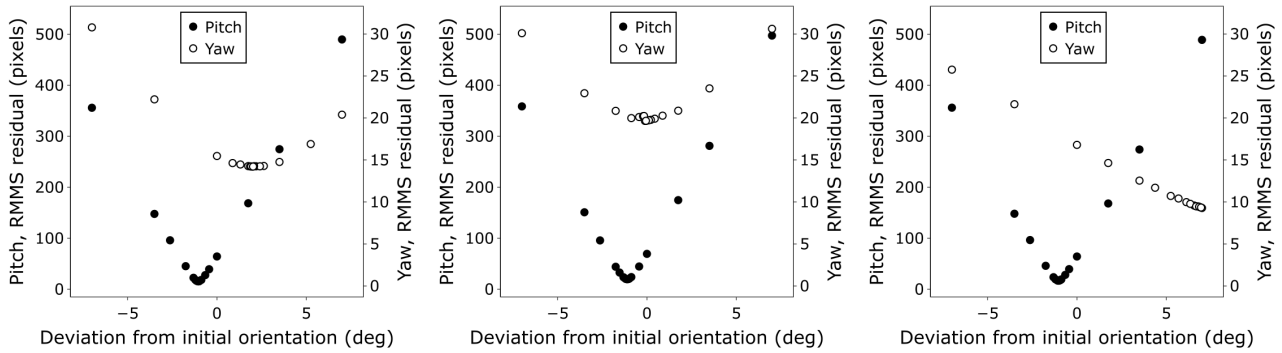

Fig. 5: Plots of the optimisation of real data. Showing optimisation of data for the whale (A), coral mound (B), and fairly plain (C) transects. Deviation from initial orientation refers to that of the fore laser. Deviation of the aft laser can be inferred as the same as that of the fore laser in pitch and the negative in yaw.

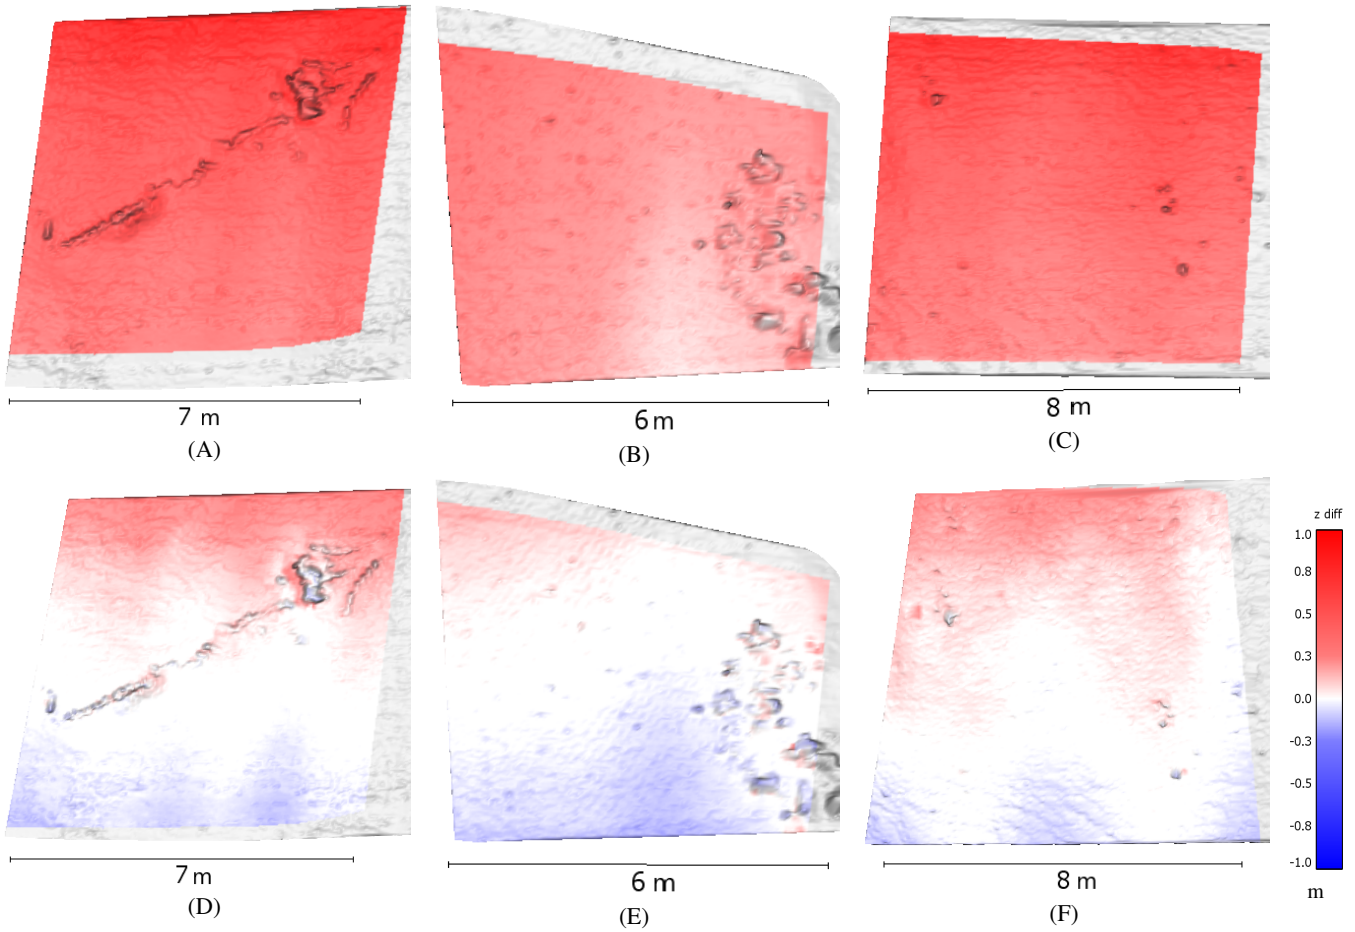

Fig. 6: Plots of the initial depth error (top row) and resultant depth error (bottom row) between fore and aft laser meshes for the optimisation of the real data. Showing the whale (A, D), coral mound (B, E), and fairly plain (C, F) transects. Depth error shown as the colouration of the fore laser mesh. Areas of no overlap between the two meshes are coloured grey.

109 research expedition of the RRS Discovery to the Darwin Mounds Marine Protected Area in the North Atlantic [9]. These are mounds of cold-water coral at around 1,000 m depth. The University of Southampton's BioCam mapping system was mounted on the National Oceanography Centre's Autosub6000 AUV and captured a dataset covering 29 hectares of the seafloor at millimetre-order resolution.

We optimised three short transects; featuring a whale car-

TABLE III: Optimal Deviation of Lasers from Initial Orientations for Real Data

Transect Name	Fore Laser		Aft Laser	
	Pitch (deg)	*Yaw (deg)*	*Pitch (deg)*	*Yaw (deg)*
Whale	-1.0391	2.0234	-1.0391	-2.0234
Mound	-1.0938	-0.1094	-1.0938	0.1094
Plain	-1.0391	7.0000*	-1.0391	-7.0000*

*Boundary value.

cass, the beginning of a coral mound, and some fairly plain terrain (a different transect to that used in III). Bounds were set for the initial, calibrated laser orientations as ± 3 deg, the maximum number of iterations was set to 7, and the desired tolerance was set to 10^{-5} deg. Fig. 5 plots the optimisation process for each transect. Fig. 6 shows the effect on the consistency of the two meshes that are separately produced from observations of the fore laser and the aft laser, respectively, for each transect. Table III details the optimal orientation values arrived at by the optimisation process.

The results show the optimisation process converging for both pitch and yaw in the case of the whale and coral mound transects, but only the pitch has converged for the fairly plain transect with the yaw looking as though it might have converged if the bounds had not restricted it. However, although these transects are all from the same survey and so the true laser parameters should be constant, the optimisation process has arrived at different values of the yaw of the laser planes for different transects. The values arrived at for pitch are largely, although not entirely, constant. The reason for the difference in the yaw of the laser could be caused by differences in the terrain, and further experiments should be performed using larger transects of data.

Compared with the results of the simulations in III, these error gradients appear just as smooth. The error curves for pitch appear slightly less sharp than that of the optimal simulated data, but nearer the optimal case than the various sub-optimal cases. For all transects, the optimisation has greatly reduced the depth error between the fore and aft laser meshes. Future work should investigate how different laser models, e.g. cone instead of plane models, and multivariate techniques can be applied to further improve mapping self-consistency.

V. Conclusions

- The self-consistency of structured light bathymetry can be improved using a featureless approach that reconstructs overlapping mapping observations and uses an error metric in the sensor frame. The analysis of experimental data in this work showed that the featureless approach gives smooth error gradients for convergence.
- A dual laser setup increases the proportion of observations that have overlap, with the small time-delay involved increasing confidence in parameter optimisation due to the reduced effect of drift in position estimates.
- The sequential and coupled method for parameter optimisation used here only converges to a true solution where the initial parameter values are offset from their true values in the same coupled fashion. To generalise the optimisation, multivariate approaches that simultaneously alter each parameter of each laser independently should be investigated.
- While this sequential and coupled method showed significant improvements in the self-consistency of transects of data, some dependency of the optimised parameters on the analysed terrains was found. Future studies should optimise calibration parameters over a sufficiently large dataset to address this issue.
- To enable practical implementation on large datasets, data stratification methods should be investigated to limit computational costs.

Acknowledgements

This work was funded by the UK Research and Innovation Natural Environment Research Council under BioCam NE/P020887/1, part of the Oceanids Marine Sensor Capital program, and part-funded by the Engineering and Physical Science Research Council, via the Sustainable Infrastructure Systems Centre for Doctoral Training. The authors thank the Autosub team, the crew of the RRS Discovery, and their fellow scientific colleagues during DY108-109.

References

[1] 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, 2012, pp. 4425-4432, doi: 10.1109/IROS.2012.6386038

[2] 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

[3] Roman, C. and Singh, H. (2007), "A Self-Consistent Bathymetric Mapping Algorithm," J. Field Robotics, 24: 23-50. doi:10.1002/rob.20164

[4] S. Barkby, S.B. Williams, O. Pizarro, and M.V. Jakuba, "A Featureless Approach to Efficient Bathymetric SLAM using Distributed Particle Mapping," *J. Field Robotics*, vol. 28, no. 1, pp 19-39, 2011, doi: 10.1002/rob.20382

[5] M. Massot-Campos, G. Oliver-Codina, A. Bodenmann and B. Thornton, "Submap Bathymetric SLAM using Structured Light in Underwater Environments," *AUV2016 – Tokyo*, Tokyo, Japan, pp. 1-8, 2016, 10.1109/AUV.2016.7778669

[6] M. Massot-Campos, G. Oliver-Codina, and B. Thornton, "Laser Stripe Bathymetry using Particle Filter SLAM," *OCEANS 2019 — Marseille*, Marsielle, France, pp. 1-7, 2019, doi: 10.1109/OCEANSE.2019.8867106

[7] 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.

[8] K. Istenic, V. Ila, L. Polok, N. Gracias and R. García, "Mission-time 3D reconstruction with quality estimation," OCEANS 2017 - Aberdeen, Aberdeen, 2017, pp. 1-9, doi: 10.1109/OCEANSE.2017.8084708.

[9] V. Huvenne, B. Thornton, "Cruise Report No. 66 RRS Discovery Cruise DY108-109," 224p.