

Performance Evaluation of 3D Keypoint Detectors and Descriptors on Coloured Point Clouds in Subsea Environments

Kyungmin Jung, Thomas Hitchcox, James Richard Forbes

Abstract—The recent development of high-precision subsea optical scanners allows for 3D keypoint detectors and feature descriptors to be leveraged on point cloud scans from subsea environments. However, the literature lacks a comprehensive survey to identify the best combination of detectors and descriptors to be used in these challenging and novel environments. This paper aims to identify the best detector/descriptor pair using a challenging field dataset collected using a commercial underwater laser scanner. Furthermore, studies have shown that incorporating texture information to extend geometric features adds robustness to feature matching on synthetic datasets. This paper also proposes a novel method of fusing images with underwater laser scans to produce coloured point clouds, which are used to study the effectiveness of 6D point cloud descriptors.

I. INTRODUCTION

Conventional vision-based navigation has seen ample application in the field of subsea robotics. However, underwater environments pose unique challenges to vision-based navigation algorithms, such as light scattering, poor and inconsistent illumination, and feature-depleted scenes [1].

Newly developed high-precision subsea 3D optical scanners allow for more robust feature extraction in subsea environments [2]. For example, feature matching techniques have successfully been leveraged to initialize the alignment of bathymetric point cloud scans via a coarse alignment process [3]. This approach has been effective in bounding the navigation drift over time when used within an underwater simultaneous localization and mapping (SLAM) framework.

However, an evaluation of 3D keypoint detectors and descriptors for subsea environments appears to be absent in the literature. Moreover, a study on which detector/descriptor pairs work best in terms of robustness and reliability when used to initialize an iterative closest point (ICP) pipeline in subsea environments has yet to be thoroughly explored. This paper surveys the performance of 3D detector/descriptor pairs for point cloud alignment in a variety of underwater environments. The motivation for such a study stems from the need to incorporate loop closures within a bathymetric SLAM pipeline to correct odometry drift.

Furthermore, studies on synthetic datasets have shown that incorporating colour information to extend geometric

This work was supported in part by Voyis Imaging Inc. through the NSERC CRD Program and in part by the MEDA Program at McGill University. The authors are with the department of Mechanical Engineering, McGill University, Montreal, QC H3A 0C3, Canada. {kyungmin.jung@mail.mcgill.ca, thomas.hitchcox@mail.mcgill.ca, james.richard.forbes@mcgill.ca}.

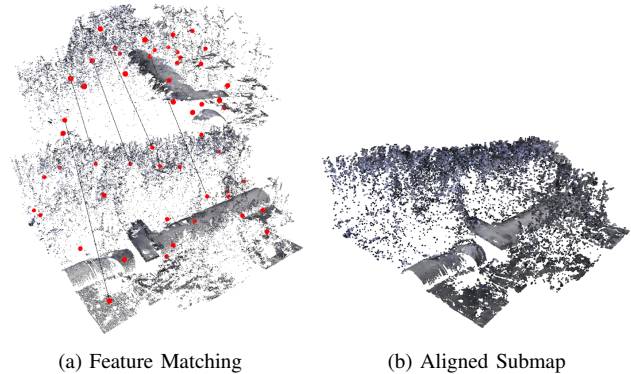


Fig. 1: Point cloud alignment via feature matching is shown here. Fig 1a shows two point cloud submaps generated from laser scans of a subsea pipeline. Note the submaps are separated in the vertical direction for easier visualization. The extracted keypoints from the two submaps are shown as red dots. The aligned submaps are shown in Fig 1b.

descriptors improves the robustness of feature matching [4]. As a secondary contribution, this paper proposes a novel method of colouring point clouds using an underwater laser scanner and an external colour camera. The resulting coloured point clouds are used to evaluate the performance of colour-enhanced feature descriptors.

In detail, this paper’s contributions include

- the generations of 3D coloured point clouds using an underwater laser line scanner and an external colour camera to add robustness to feature extraction,
- a robustness analysis of various 3D keypoints subject to different transformations and point cloud noise levels in different underwater environments, and
- a performance evaluation of 3D features in initializing an ICP fine alignment algorithm.

The remainder of this paper is organized as follows. Section II begins by outlining related work. Section III presents the novel method of colour point cloud generation. Section IV then explains the methods used to evaluate various detector/descriptor pairs with coloured underwater point clouds. Section V presents the evaluation results. Finally, Section VI summarizes the findings.

II. RELATED WORK

The effectiveness of different detector/descriptor pairs has been evaluated in the literature via terrestrial benchmark datasets. For example, a comprehensive evaluation of state-of-the-art 3D keypoints was conducted in [5], which mainly



Fig. 2: A laser line scanner, *Insight Pro* (top), and an imaging system, *Observer Micro* (bottom), provided by Voyis Imaging Inc. are used to collect point cloud and image data, respectively.

addressed their common traits in 3D object recognition tasks. The evaluation was carried out to analyze repeatability and computational efficiency using synthetic datasets from [6]. Keypoint repeatability was measured as a detector’s ability to find the same set of keypoints on different instances of a given point cloud that may be corrupted by noise. Overall, KPQ-SI, MeshDoG, and ISS yielded the best scores in terms of repeatability in simulated data. Further, [5] stated that the accuracy of feature matching could be improved by incorporating texture information into the descriptors.

The descriptiveness and robustness of ten 3D local descriptors were studied in [7]. The descriptiveness was evaluated using a precision-recall curve, and the robustness was evaluated with respect to varying support radius, noise values, and mesh resolutions. However, [7] analyzed the performance in terms of object recognition and shape retrieval and did not consider navigation applications.

More recently, [8] presented a comprehensive study of 3D point cloud descriptors, including more recently developed deep-learning-based descriptors. The paper also described the traits of local descriptors, which construct a histogram around keypoints within the input point cloud, and global descriptors, which generally estimate a single descriptor vector encoding the whole input cloud for scenarios like 3D object recognition and geometric categorization.

III. COLOURED POINT CLOUD GENERATION

This section discusses data collection methods, point cloud submap generation, and the novel method of point cloud colourization using a high-precision underwater laser line scanner and an external RGB camera.

A. Underwater Field Dataset

Field data was collected at Colpoy’s Bay, Ontario, Canada by Voyis Imaging Inc. using their surface vessel equipped with a Sonardyne SPRINT-Nav 500 Doppler velocity log-aided inertial navigation system (DVL-INS), a u-blox ZED-F9P high precision GNSS module equipped with a u-blox ANN-NM series high precision multi-band antenna, a Voyis *Insight Pro* underwater laser scanner, and an *Observer Micro* underwater imaging system. The laser scanner and the imaging system are shown in Fig 2. Laser ‘submaps’ are constructed by registering individual laser line profiles to the trajectory estimate produced by the DVL-INS. Although the DVL-INS provide very accurate navigation estimate, its




Fig. 3: The vehicle trajectory over a pipeline target is shown as a black line, while loop-closure crossings are shown as red circles. The colour of the underlying point cloud scan represents relative depth across the dataset.

TABLE I: A summary of the field data in number of loop closures.

Section	Structured	Semistructured	Unstructured	Total
Shipwreck A	9	6	6	21
Shipwreck B	13	16	54	83
Pipeline	6	2	0	8
Seabed	0	0	6	6
Total	28	24	66	118

drift becomes significant in long-term missions. A detailed explanation of point cloud generation and colourization is explained in Section III-B and Section III-C, respectively.

The collected field data is divided into four sections: two shipwreck sections, a pipeline section, and a seabed section. Each section of the data consists of multiple loop-closure events as the vehicle makes a loop around a previously visited region as shown in Fig 3. For each loop closure, a pair of submaps is generated around the two intersecting vehicle poses. Each submap pair is qualitatively assigned to one of three categories depending on the local environment. These categories are structured scenes such as shipwrecks and pipelines, semi-structured scenes such as debris fields, and unstructured scenes such as seabed and marine plant life. The number of classified submaps is shown in Table I.

B. Submap Generation

The Voyis *Insight Pro* laser line scanner measures a single array of points $\mathcal{P}_\ell = \{\mathbf{r}_\ell^{p_i s}\}_{i=1}^N$, where $\mathbf{r}_\ell^{p_i s}$ is the position vector of reflected point p_i relative to the sensor s resolved in the laser frame \mathcal{F}_ℓ , and N is the total number of points in a single scan line.

A submap at time t_k is composed of multiple nearby laser line measurements resolved in the vehicle body frame at time \mathcal{F}_{b_k} . Thus, laser points measured at some time t_k near a loop-closure time t_τ are transformed using the vehicle navigation estimate,

$$\mathbf{r}_{b_\tau}^{p_i z_\tau} = (\mathbf{T}_{ab_\tau}^{z_\tau w})^{-1} \mathbf{T}_{ab_k}^{z_k w} \mathbf{T}_{b_\ell}^{s z_\tau} \tilde{\mathbf{r}}_\ell^{p_i s}, \quad (1)$$

where $\tilde{\mathbf{r}}_\ell^{p_i s}$ is the homogeneous form of $\mathbf{r}_\ell^{p_i s}$, and

$$\mathbf{T}_{ab_k}^{z_k w} = \begin{bmatrix} \mathbf{C}_{ab_k} & \mathbf{r}_a^{z_k w} \\ \mathbf{0} & 1 \end{bmatrix} \in SE(3) \quad (2)$$

represents the pose of the vehicle as an element of matrix Lie group $SE(3)$. Here, $\mathbf{r}_a^{z_k w} \in \mathbb{R}^3$ denotes the position of the vehicle body, z_k , relative to some reference point w resolved in the world frame \mathcal{F}_a at some time t_k , and the direction cosine matrix $\mathbf{C}_{ab_k} \in SO(3)$ represents the attitude of the body frame \mathcal{F}_{b_k} at time t_k relative to the reference frame \mathcal{F}_a . Further, $\mathbf{T}_{b_\ell}^{s z_\tau}$ denotes the static extrinsics transformation from




Fig. 4: A shipwreck is reconstructed from laser scans (left). Occluded points from the camera’s viewpoint are removed using the method from [11] (right), so as to only add colour to points visible to the camera.

the laser sensor to the vehicle. Denote $\mathcal{P}_{b_\tau}^k = \{\mathbf{r}_{b_\tau}^{p_k, i, z_\tau}\}_{i=1}^{N_k}$ as the laser line measurement at time t_k resolved in the body frame at loop-closure time t_τ . Then a submap is constructed as

$$\mathcal{S}_{b_\tau} = \{\mathcal{P}_{b_\tau}^k\}_{k=1}^M, \quad (3)$$

where M is the number of laser line measurements in the submap. Following [9], the size of the submap is chosen as $5\text{ m} \times 5\text{ m}$ such that it is small enough to be assumed error-free and large enough to contain sufficient 3D information to be aligned to another submap. Further, all submaps are downsampled using the voxel grid filter from [10] with the grid size 0.05 m . The grid size is chosen to perform well with the dataset from Section III-A.

An accurate laser-to-body extrinsics estimate is necessary for accurate point cloud registration (1). The laser-to-body extrinsics are estimated a priori by minimizing the sum of squared reprojection errors.

C. Coloured Point Clouds

Incorporating texture information to extend geometric features has been shown to improve the feature matching accuracy [4]. The point clouds generated in Section III-B are textured using images from an underwater colour camera.

To texture the submap, the point cloud is first transformed to a camera frame to determine point visibility. In particular, each $\mathbf{r}_b^{p_i, z}$ is transformed via

$$\tilde{\mathbf{r}}_m^{p_i, c} = (\mathbf{T}_{bm}^{cz})^{-1} \mathbf{r}_b^{p_i, z}, \quad (4)$$

where $\mathbf{T}_{bm}^{cz} \in SE(3)$ is the camera-body extrinsic matrix estimated using stereo calibration. To determine the set of points visible to the camera, the approach presented in [11] is used, where a convex hull is constructed by transforming the point cloud to a new domain. The points that lie on the convex hull of the transformed set of points are considered visible to the camera. An example is shown in Fig 4.

After determining point visibility, the visible points are projected to the image plane using a pinhole camera model

$$\begin{bmatrix} u & v & w \end{bmatrix}^\top = \mathbf{K} \mathbf{r}_m^{p_i, c}, \quad (5)$$




Fig. 5: A 3D point projected to an image frame \mathcal{I} , shown as a red dot, is visible in multiple images at different pixel locations. Each colour candidate is weighted with the corresponding value of the 2D Gaussian. A unique colour is assigned by taking the sum of all weighted colour candidates (6).

TABLE II: 3D detectors / descriptors and their parameters.

Detector	Parameters	Descriptor	Size
ISS	Threshold, Radius Search	PFH	125
Harris3D	Threshold, Radius Search	PFHRGB	250
Lowe	Threshold, Radius Search	SHOT	352
Tomasi	Threshold, Radius Search	CSHOT	1344
Curvature	Threshold, Radius Search	3DSC	1980
Harris6D	Threshold, Radius Search	USC	1960
SIFT3D	Scales, Minimum Contrast		
SUSAN3D	Distance and Angular Threshold		

where \mathbf{K} is the camera intrinsic matrix. The colours from the corresponding pixels are assigned to the points.

In practice, the points in the submap may be seen in multiple camera frames at different observation angles resulting in multiple colour assignments. This issue is solved in [12] using a weighting scheme,

$$\bar{x} = \frac{\sum_i^{N_c} w_i x_i}{\sum_i^{N_c} w_i}, \quad (6)$$

where N_c is the total number of colour candidates, x_i is the i -th colour candidate, and w_i is the weight associated to each colour candidate. The weights are chosen from a 2D Gaussian function overlaid on top of the image with its peak at the optical centre as shown in Fig 5. The 2D Gaussian is used to weigh the points closer to the centre of the image higher. An example of the resulting coloured point cloud is shown in Fig 6.

IV. EXPERIMENTAL METHOD

A. Detector/Descriptor Pairs

The keypoints [13–16] and features [4, 17–20] in Table II are selected based on their performance evaluation found in the literature and their accessibility from the open source Point Cloud Library (PCL) [10]. Eight 3D keypoint detectors are evaluated for their robustness and computational complexity. For each chosen keypoint, 6 descriptors are evaluated in a point cloud alignment pipeline, for a total of 48 unique detector/descriptor pairs. The parameters for each detector and descriptor are set to default values provided in PCL.

B. Keypoint Evaluation

The most important characteristic of a keypoint detector is *repeatability*, defined as the detector’s ability to identify the

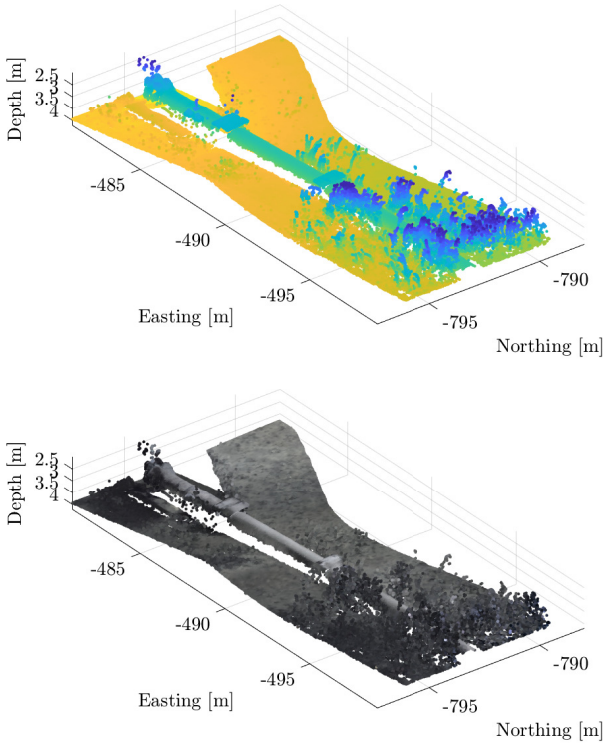


Fig. 6: A submap generated from the pipeline data displayed in elevation (top) and in colour (bottom).

same set of points on different instances of a given model [5]. The different instances discussed for keypoint evaluation are viewpoint change and noise corruption.

Optical sensors in underwater applications tend to point downward as they scan structures on the seafloor. Therefore, keypoint repeatability is evaluated on point clouds rotated about the vertical basis vector of the world frame, a_3 , which is often perpendicular to the seafloor. A stationary submap, S_{b_1} , resolved in the body frame \mathcal{F}_{b_1} is rotated with angles from 0° to 180° with 10° step to create rotated point clouds.

The robustness of the keypoints is tested on noisy point clouds as well. Gaussian white noise is injected to S_{b_1} with 1σ increments of 0.005 from 0 to 0.05. The transformed or the noisy submap is defined as the *modified* submap, S_{b_2} .

For each pair of S_{b_1} and S_{b_2} , two sets of keypoints, $\mathcal{K}_{b_1} = \{\mathbf{r}_{b_1}^{f_i z_1}\}_{i=1}^{N_1}$ and $\mathcal{K}_{b_2} = \{\mathbf{r}_{b_2}^{f_j z_2}\}_{j=1}^{N_2}$, are extracted from each submap, respectively. Similar to [5], the keypoints are said to be repeatable if they are approximately coincident,

$$\left\| \mathbf{C}_{b_1 b_2} \mathbf{r}_{b_2}^{f_j z_2} + \mathbf{r}_{b_1}^{z_2 z_1} - \mathbf{r}_{b_1}^{f_i z_1} \right\| < \epsilon, \quad (7)$$

where $\mathbf{C}_{b_1 b_2} \in SO(3)$ and $\mathbf{r}_{b_1}^{z_2 z_1}$ define the perturbing transformation. A threshold value $\epsilon = 1e-2$ is chosen to capture the proximity of corresponding keypoints even in noisy clouds. Given a set of repeatable keypoints \mathcal{RK} , the relative repeatability r is given by

$$r = \frac{|\mathcal{RK}|}{|\mathcal{K}_{b_1}|}. \quad (8)$$

The processing time for each keypoint is also evaluated.




Fig. 7: Block diagram of point cloud alignment pipeline.

C. Feature Evaluation

Each loop closure produces two submaps at two different vehicle poses from which keypoints are detected and features are extracted. Extracted features are matched using the k -nearest-neighbour algorithm. The features from source clouds are matched to the target clouds and vice-versa to produce unique matches.

Matched keypoints are used to perform a coarse alignment using [21], which attempts to solve an intractable optimization problem in the presence of a large number of outlier correspondences. Thus, the coarse alignment algorithm is used as an outlier rejection method. Feature matching is considered to have failed if the number of inliers is insufficient to allow the coarse alignment algorithm to accurately align the two submaps.

Once the coarse alignment has been performed, a modified recall equation from [22] is used to evaluate the inlier ratio,

$$\text{recall} = \frac{\#\text{inliers}}{\#\text{features}}. \quad (9)$$

Further, coarse alignment estimates the relative transform between the two submaps $\hat{\mathbf{T}}_{b_2 b_1}^{z_2 z_1} \in SE(3)$. Since ground truth information is not available, the alignment accuracy is evaluated in the matrix Lie algebra $\mathfrak{se}(3)$ as

$$\delta \xi_i = \log((\hat{\mathbf{T}}_{b_1 b_2}^{z_2 z_1})^{-1} (\mathbf{T}_{b_1 b_2}^{z_2 z_1})_i)^\vee = \begin{bmatrix} \delta \xi_i^\phi^\top & \delta \xi_i^\rho^\top \end{bmatrix}^\top, \quad (10)$$

where $\mathbf{T}_{b_2 b_1}^{z_2 z_1}$ is the relative transform obtained from the GPS-aided DVL-INS solution.


D. Full Pipeline Evaluation

The full pipeline of the algorithm is shown in Fig 7. The fine alignment is performed based on [23] and [24] with the prior pose initialized by the coarse alignment algorithm. Fine alignment produces a posterior estimate of the relative transform between the two submaps, $\hat{\mathbf{T}}_{b_1 b_2}^{z_2 z_1}$. Since ground truth information is not available, the alignment accuracy is evaluated using the self-consistency error metric of [9].


V. RESULTS AND DISCUSSION

A. Keypoint Evaluation

The first set of experiments investigates which keypoints are most robust to rotation and point cloud noise. The relative repeatability of each keypoint in all underwater scenes is shown in Fig 8. Solid lines represent the mean value, and the banded areas represent a 95% confidence level. ISS keypoints have the highest mean repeatability with the lowest variance for rotated point clouds, followed by Harris3D and Lowe. ISS, Harris3D, Lowe, Tomasi, Harris6D and SIFT



(a) Rotation



(b) Point cloud noise

Fig. 8: Keypoint repeatability results

TABLE III: Keypoint detection time analysis in milliseconds. The abbreviated column names are H3D: Harris3D, CU: Curvature, and H6D: Harris6D.

ISS	H3D	Low	Tomasi	CU	H6D	SIFT	SUSAN
16.3	13.9	13.7	12.2	4.02	24.1	1213	2.51

are rotation invariant. Curvature and SUSAN vary with rotation as the repeatability decreases with the amount of rotation. However, the Curvature method shows the highest robustness to the point cloud noise followed by SUSAN and SIFT. Harris6D is the most susceptible to point cloud noise. Further, Table IV shows that ISS and SIFT are the most versatile keypoints that can perform with various descriptors.

The average time taken to detect the keypoints is shown in Table III. SIFT takes the longest to detect at around 1.2s per submap, followed by Harris6D and ISS. SUSAN and Curvature take the least time to detect at around 0.002s and 0.004s, respectively. As SUSAN yields the lowest detection time with relatively high robustness to point cloud noise and rotations, it may be suitable for real-time applications.

B. Feature Evaluation

The features are evaluated with every detector/descriptor combination in all 118 loop-closure events classified into 3 different underwater scenes. For structured scenes, USC

TABLE IV: Percent success in feature matching for every descriptor pair

	PFH	PFHRGB	SHOT	CSHOT	3DSC	USC
ISS	100	95.8	100	100	40.7	99.2
H3D	94.9	83.9	84.7	75.4	21.2	98.3
Low	99.2	91.5	85.6	82.2	23.7	99.2
Curv	64.4	48.3	36.4	30.5	19.5	69.5
Tomasi	90.7	75.4	78.0	69.5	21.2	95.6
H6D	96.6	88.1	88.1	80.5	23.7	99.2
SUSAN	41.5	44.1	24.6	30.5	17.8	39.8
SIFT	96.6	89.8	83.9	82.2	53.4	100

TABLE V: Percent success in feature matching for different scenes

Feature	Structured	Semi-structured	Unstructured	Total
PFH	93.30	92.71	79.55	85.49
PFHRGB	88.84	86.46	68.75	77.12
SHOT	54.91	72.92	80.11	72.67
CSHOT	50.45	70.83	75.95	68.86
3DSC	51.34	17.19	21.40	27.65
USC	95.09	91.14	83.14	87.61

has the highest success rate in feature matching followed by PFH and PFHRGB as shown in Table V. USC also has the highest success rate for the unstructured environments, while all other descriptors except 3DSC show relatively similar performance. Overall, USC is the most robust keypoint in different underwater scenes with the highest success rate.

The alignment errors are plotted and compared in Fig 9. The rotation and translation errors exhibit similar patterns for each descriptor. 3DSC seems to perform the best with coarse alignment, but it has the lowest recall value shown in Fig 9c, indicating that the number of inliers is too low to accurately perform coarse alignment. Thus, USC performs the best in terms of the alignment accuracy, followed by CSHOT.

Feature matching has a greater tendency to fail in unstructured environments, as shown in Table V, yet successful alignments in unstructured environments produce the lowest alignment error. This means that once the correspondences are found, the coarse alignment succeeds with very low transformation error. Features tend to have higher chances of finding matches in structured scenes, but with higher alignment error. This may be because the coarse alignment algorithm is unable to find a sufficient number of inliers.

Despite USC delivering the highest performance in terms of the robustness to environments and the alignment accuracy, USC is the second-slowest in processing time. Thus, USC may not be the most adequate descriptor for real-time applications. As shown in Table VI, SHOT and CSHOT outperform in overall processing by a significant margin, and thus may be considered as candidates for real-time applications. PFH and PFHRGB can also be considered as candidates for increased robustness at the expense of longer processing times. Unfortunately, the experimental results produced no evidence that adding texture information to the point cloud scan improves alignment robustness. This may be due to the greyish, monochromatic colours in the images caused by low illumination underwater. This can possibly be improved with the development of a better colour levelling algorithm.




Fig. 9: Coarse alignment errors for ISS keypoints, grouped by category of underwater environment.

TABLE VI: Feature processing time analysis in milliseconds. The abbreviated column names are CA: coarse alignment, and FA: fine alignment.

Feature	Extraction	Matching	CA	FA	Total
PFH	124.7	5.619	50.22	266.0	446.5
PFHRGB	205.7	7.638	48.64	269.1	531.1
SHOT	14.97	9.686	56.92	219.3	300.9
CSHOT	17.04	23.38	58.11	203.7	302.2
3DSC	712.1	32.68	53.86	224.5	1023
USC	513.7	33.06	51.04	227.2	825.0

C. Full Pipeline Evaluation

The full point cloud alignment pipeline illustrated in Fig 10 is run with a combination of ISS keypoints and USC descriptors, as this combination exhibits the highest robustness to environments changes and the highest accuracy in the coarse alignment algorithm. Fig 10 shows the initial alignment errors, the errors after coarse alignment, and the errors after fine alignment. The errors after coarse alignment are low enough to properly initialize the ICP algorithm [23], and align the two submaps together.

VI. CONCLUSION

This paper evaluates the effectiveness of different detector/descriptor pairs for point cloud data collected in novel and challenging underwater environments. Additionally, a novel method has been proposed for colourizing subsea point cloud scans with images collected from an underwater colour camera. This allows descriptors incorporating colour information to be included as part of the study.

The study shows that ISS, Lowe, and Harris3D outperform other keypoints in terms of point cloud noise and rotation, while SUSAN outperforms others in terms of processing time. The USC feature outperforms other features in matching rate and alignment accuracy followed by PFH and PFHRGB. Interestingly, features with colour information do not outperform their geometric counterparts, possibly due to low illumination and contrast in the underwater environment. The evaluation of newly developed machine learning-based detectors and descriptors is an interesting avenue for future research.

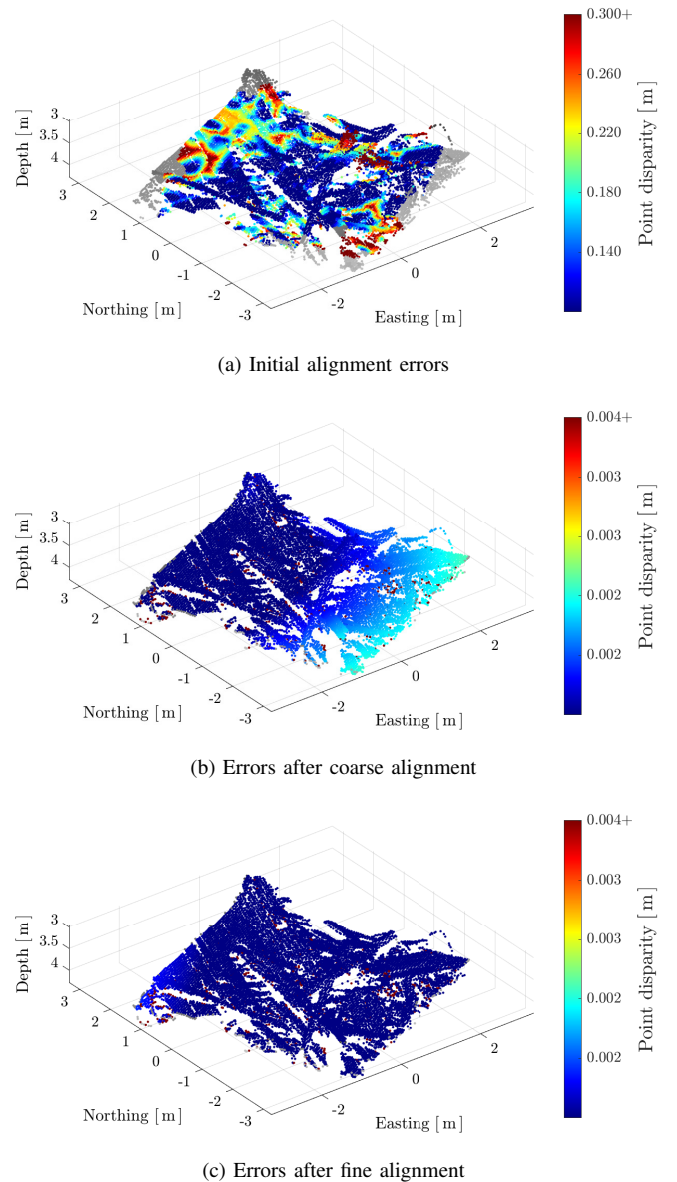


Fig. 10: Consistency-based errors [9] are computed on a loop closure from Shipwreck B data with ISS keypoints and USC descriptors.

REFERENCES

- [1] S. Rahman, A. Q. Li, and I. Rekleitis, "SVIn2: An underwater SLAM system using sonar, visual, inertial, and depth sensor," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 1861–1868.
- [2] M. Castellón, A. Palomer, J. Forest, and P. Ridao, "State of the art of underwater active optical 3D scanners," *Sensors*, vol. 19, no. 23, 2019.
- [3] A. Palomer, P. Ridao, and D. Ribas, "Inspection of an underwater structure using point-cloud SLAM with an AUV and a laser scanner," *Journal of Field Robotics*, vol. 36, no. 8, pp. 1333–1344, 2019.
- [4] F. Tombari, S. Salti, and L. Di Stefano, "A combined texture-shape descriptor for enhanced 3D feature matching," in *IEEE International Conference on Image Processing (ICIP)*, IEEE, 2011, pp. 809–812.
- [5] S. Salti, F. Tombari, L. Di Stefano, and L. D. Stefano, "A performance evaluation of 3D keypoint detectors," in *International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT)*, IEEE, 2011, pp. 236–243.
- [6] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, 1996, pp. 303–312.
- [7] Y. Guo, M. Bennamoun, F. Sohel, M. Lu, J. Wan, and N. M. Kwok, "A comprehensive performance evaluation of 3D local feature descriptors," *International Journal of Computer Vision*, vol. 116, no. 1, pp. 66–89, 2016.
- [8] X.-F. Han, S.-J. Sun, X.-Y. Song, and G.-Q. Xiao, "3D point cloud descriptors in hand-crafted and deep learning age: State-of-the-art," 2018. arXiv: 1802.02297.
- [9] C. Roman and H. Singh, "A self-consistent bathymetric mapping algorithm," *Journal of Field Robotics*, vol. 24, no. 1-2, pp. 23–50, 2007.
- [10] R. B. Rusu and S. Cousins, "3D is here: Point cloud library (PCL)," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2011, pp. 1–4.
- [11] S. Katz, A. Tal, and R. Basri, "Direct visibility of point sets," *ACM Transactions on Graphics*, vol. 26, no. 3, 24–es, 2007.
- [12] P. Vechersky, M. Cox, P. Borges, and T. Lowe, "Colourising point clouds using independent cameras," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3575–3582, 2018.
- [13] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [14] I. Sipiran and B. Bustos, "Harris 3D: A robust extension of the Harris operator for interest point detection on 3D meshes," *Visual Computer*, vol. 27, pp. 963–976, 2011.
- [15] Y. Zhong, "Intrinsic shape signatures: A shape descriptor for 3D object recognition," in *IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, 2009, pp. 689–696.
- [16] S. M. Smith and J. M. Brady, "SUSAN—A new approach to low level image processing," *International Journal of Computer Vision*, vol. 23, no. 1, pp. 45–78, 1997.
- [17] A. Frome, D. Huber, R. Kolluri, T. Bülow, and J. Malik, "Recognizing objects in range data using regional point descriptors," in *European Conference on Computer Vision (ECCV)*, T. Pajdla and J. Matas, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 224–237.
- [18] R. B. Rusu, N. Blodow, Z. C. Marton, and M. Beetz, "Aligning point cloud views using persistent feature histograms," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3384–3391, 2008.
- [19] F. Tombari, S. Salti, and L. Di Stefano, "Unique shape context for 3D data description," in *Proceedings of the ACM Workshop on 3D Object Retrieval*, ser. 3DOR '10, New York, NY, USA: Association for Computing Machinery, 2010, pp. 57–62.
- [20] F. Tombari, S. Salti, and L. D. Stefano, "Unique signatures of histograms for local surface description," in *European Conference on Computer Vision (ECCV)*, Springer, 2010, pp. 356–369.
- [21] H. Yang, J. Shi, and L. Carlone, "Teaser: Fast and certifiable point cloud registration," *IEEE Transactions on Robotics*, vol. 37, no. 2, pp. 314–333, 2021.
- [22] F. Hidalgo and T. Bräunl, "Evaluation of several feature detectors/extractors on underwater images towards vSLAM," *Sensors*, vol. 20, no. 15, pp. 1–16, 2020.
- [23] D. Qian, G. Charland-Arcand, and J. R. Forbes, "TWOLATE: Total registration of point-clouds using a weighted optimal linear attitude and translation estimator," in *IEEE Conference on Control Technology and Applications (CCTA)*, 2020, pp. 43–48.
- [24] T. Hitchcox and J. R. Forbes, "Comparing robust cost functions for bathymetric point cloud registration," in *IEEE/OES Autonomous Underwater Vehicles (AUV)*, 2020, pp. 1–6.