https://doi.org/10.2352/EI.2026.38.16.AVM-103

Automatic Registration of Subsea LiDAR Point Clouds

Josie G. Clapp (jgc2932@rit.edu), Byron K. Eng (bkecis@rit.edu), Carl Salvaggio (cnscpi@rit.edu); Rochester Institute of Technology, College of Science, Chester F. Carlson Center for Imaging Science, 54 Lomb Memorial Drive, Rochester, NY 14623, USA

Abstract

Accurate registration of subsea Light Detection and Ranging (LiDAR) point clouds is critical for offshore metrology, where millimeter-level errors can significantly impact operational cost and risk. This study evaluates automated registration methods for static scan positions acquired by Kraken Robotics systems. Two approaches were implemented in Open3D: a hierarchical tree-based Iterative Closest Point (ICP) method and a pose-graph multiway registration framework. The methods were tested on multiple real subsea datasets containing 19–33 high-density scans per subsea scene and on synthetic datasets generated with Digital Imaging and Remote Sensing Image Generation (DIRSIG) to enable ground-truth evaluation.

Results show that multiway registration provides improved global consistency, lower adjacent-scan Root Mean Square Error RMSE, and reduced processing time compared to tree-based ICP. Ground-truth analysis demonstrated sub–sampling-level performance, corresponding to an expected 1–4 mm alignment accuracy for real datasets. Global coarse registration provided no measurable benefit for well-initialized static surveys.

The final analysis demonstrates that multiway registration enables accurate, efficient, and fully automated subsea LiDAR alignment, reducing manual effort and improving metrology reliability.

`Code for this project can be found here: https://github.com/josieclapp7/subsea_registration`

Introduction

The growing global demand for energy, coupled with the gradual depletion of accessible hydrocarbon resources, has driven the development of oil and gas fields into deeper waters and more challenging environments [1]. As subsea infrastructure expands and becomes more sophisticated, ensuring the reliability, efficiency, and safety of these systems has become a critical priority for operators worldwide. These industrial demands have placed greater emphasis on accurate underwater measurements and detailed 3D reconstructions, which support both routine monitoring and large-scale intervention campaigns [1], [2].

Within this context, subsea metrology, the precise and traceable 3D measurement of spatial relationships between underwater structures, has become essential for planning, executing, and validating offshore operations [2]. Because offshore work is extremely costly, any reduction in uncertainty or error directly translates to significant savings in time, materials, and vessel operations. As a result, there is a growing need for efficient, reliable, and cost-optimized metrology procedures that can deliver high-accuracy data [1].

Modern LiDAR 3D point clouds play a central role in these procedures, enabling detailed geometric reconstructions of subsea assets and the surrounding seabed. Light Detection and Ranging (LiDAR) technology, in particular, has become widely adopted for subsea inspection and monitoring due to its ability to generate dense, high-resolution 3D datasets at depths of hundreds or even thousands of meters [1]. For example, the company 3D at Depth, now Kraken Robotics (KR), has been a leader in subsea LiDAR surveys, producing model precision within 5mm accuracy [3].

However, no single LiDAR scan can capture an entire subsea scene, meaning that multiple point clouds must be combined into a common coordinate frame to create a complete scene reconstruction. This process, known as point cloud registration, is a critical step in producing accurate metrology for clients [4], [5]. An inaccurate 3D reconstruction can result in substantial financial losses, potentially costing the company hundreds of thousands of dollars in wasted time and materials if maintenance operations are carried out based on an inaccurate metrology.

Given these challenges, there is a clear need for a fast, automatic, and precise system for registering subsea LiDAR point clouds.

Background

KR's workflow relies on a remotely-operated vehicle (ROV), equipped with LiDAR sensors mounted on a pan-tilt unit. The ROV scans the seafloor in overlapping sections to produce multiple partially overlapping 3D point clouds for each scan position (SP). Figure 1 demonstrates a typical SP's scope. A SP produces eight or more scans, and a full metrology may involve two or more SPs. This leads to challenges due to the time-consuming process of point cloud registration.

The time bottleneck occurs during the point cloud registration process of aligning these overlapping scans into a unified 3D model. Currently, this alignment requires significant manual effort where operators visually inspect and register point clouds using proprietary software. This process is time-consuming, labor-intensive, and prone to human error. Overall, the project aims to test and optimize an automatic, robust registration method that could dramatically streamline subsea metrology workflows, thereby reducing costs, improving reliability by speeding up the registration process, and not relying on human intervention. Overall, the goal is to compare the performance of different methods in terms of accuracy, robustness to noise/outliers, and reduction in manual intervention, while also optimizing the methods for the specific requirements of static subsea scans.

Point Cloud Registration

Point cloud registration involves estimating the transformation matrix that aligns two or more point cloud scans [4]. By

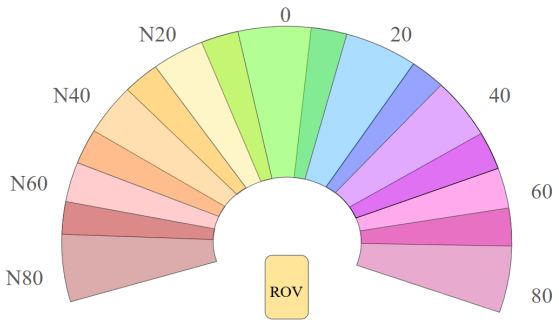

Figure 1. *Demonstration of how a scan position's (SP) data are collected. The ROV is placed in a static position onto the seabed, and the pan and tilt changes the direction of the section collected. Each section is collected in 20 degree increments, with a significant overlap to ensure that registration of the adjacent point clouds can be done with certainty. Numbers in the image represent the degree at which the point cloud was collected.*

determining this transformation, individual partial scans of the same 3D scene or object can be accurately merged into a unified and complete 3D representation. This process can be particularly challenging in the context of the overlapping scan geometry used by Kraken Robotics static scanning systems (SPs). Even a small registration error in one scan can propagate and accumulate across successive alignments, leading to significant overall inaccuracies [5]. Typically, the registration workflow consists of an initial coarse alignment, followed by a finer refinement step to improve accuracy.

Several factors contribute to the difficulty of registration, including noise from both the scanner and the subsea environment (especially evident in subsea environments with to poor visibility and floating sediment), as well as movement of the ROV, which prevents scans from being captured from identical viewpoints. To mitigate these issues, KR designs its scans with intentional overlap between adjacent point clouds, ensuring sufficient common features for accurate alignment.

The project explored two complementary approaches to the registration process: Iterative Closest Point (ICP) algorithm [6], and multiway registration [7]. Current research on registration has an emphasis on point clouds of objects and structures with distinct features, but the subsea point clouds may only have the seabed as a common feature, thus lacking artificial ground objects and structural features in the overlap [5]. These two algorithms were tested in their ability to register the subsea point clouds.

Iterative Closest Point (ICP) Registration

The Iterative Closest Point (ICP) algorithm is a well-established and widely used method for aligning 3D point clouds, making it an ideal starting point for registering subsea assets [4]. It is cited across a large body of research as the dominant geometry-based registration method due to its rigorous mathematical foundation and ability to converge on an optimal solution without requiring training data [6]. These properties are especially advantageous in subsea contexts, where acquiring large, labeled datasets is infeasible, and scene conditions vary significantly between surveys.

One of the main advantages of ICP is that it relies purely on geometric or color-based information rather than learned models, allowing it to generalize well to unknown subsea scenes [4]. Since data from subsea surveys can only be obtained through specific proprietary platforms (such as KR's metrology data), methods that do not rely on prior training are ideal. Additionally, the mathematical rigor of ICP guarantees convergence toward a local minimum under reasonable assumptions, providing reliability when working with noisy or incomplete underwater data.

Mathematically, the ICP algorithm iterates in two steps

- Establish correspondences: Identify pairs of points (p,q) where p comes from the target point cloud and q is the corresponding point in the source cloud after applying the current estimate of the transformation **T**. This set of matched pairs is denoted as K.[4]
- Refine the Transformation: Update the transformation T by minimizing an objective function E(**T**) , which measures the alignment error on all correspondences in K.

However, ICP is sensitive to issues common in subsea point clouds, such as noise, outliers, and partial overlap, which can increase computational cost and affect convergence. This is even more evident underwater, where subsea visibility can be poor, and significant noise is added to the data. Although KR does have a very effective cleaning method for removing outliers and subsea noise, it struggles at removing solid objects such as marine life like fish or "shadows" in the lidar data.

The ICP algorithm aligns two point clouds by starting with an initial guess of their relative transformation and iteratively refining it. In each iteration, ICP finds pairs of corresponding points (typically nearest neighbors) and minimizes an error metric that measures their alignment difference until convergence [4], [6], [8]. The methods that could be especially helpful in subsea registration are as follows. Point-to-Point ICP (P2P) minimizes the Euclidean distance between matched points; Point-to-Plane ICP (P2PL) reduces the orthogonal distance to the surface normals of the target cloud, improving results on smooth subsea structures such as pipelines.

All of these versions of ICP can be tested and implemented using the open-source library Open3D [9]. For this reason, it is an effective starting point to understand different algorithms and their effects on the registration.

ICP performs best when the initial transformation guess is reasonably close to the true alignment [6]. Without this, the algorithm can easily converge to an incorrect local minimum. Early tests using an identity matrix as the initial guess produced poor results, confirming that an initialization step is necessary [4], [8]. To overcome this, the project (and Open3D documentation) incorporates Global Registration (GR) before running ICP. Global methods, such as RANSAC-based registration using Fast Point Feature Histograms (FPFH), can automatically compute a coarse alignment between point clouds, even when their relative poses are unknown (as seen in [8]) [10].

The study by Zhu et. al shows that Global + ICP (GRICP) provides the most accurate and stable results, especially in the presence of noise [8] . However, it is also computationally intensive, with a high time cost and variable runtime performance. Therefore, while GRICP is the most accurate, Colored ICP (CICP) offers the best cost-performance ratio, achieving near-comparable accuracy with lower computational demands. How-

ever, CICP is irrelevant for the data that KR collects, as there is no color information assigned to the points in the point clouds.

The Root Mean Square Error (RMSE) between corresponding point pairs is used to assess registration quality [8]. For each target point, the nearest source point is identified, and its Euclidean distance is computed. The median of these distances represents the RMSE value, where lower values indicate better alignment.

Testing multiple variants (P2P, P2PL, and GRICP) enables evaluation of the accuracy–efficiency tradeoff, supporting selection of the most practical approach for field deployment [8].

Multiway Registration

The multiway Registration algorithm, implemented in Open3D, is derived from the work by Sungjoon et. al [11]. This method was originally developed for indoor scene reconstruction, where multiple overlapping point clouds are combined into a single globally consistent model. The algorithm's main objective is to "match and register different views globally" and to provide a fully automatic reconstruction pipeline that achieves accuracy comparable to manually assisted methods [7]. Unlike traditional pairwise registration, which aligns point clouds sequentially, multiway registration optimizes all alignments jointly, reducing the accumulation of local errors.

A key innovation of this approach is its use of line processes, an estimation technique that introduces penalty functions to suppress the effect of erroneous alignments during global optimization [11]. This feature allows the method to remain effective even when incorrect geometric matches notably outnumber correct ones. The optimization is solved through least-squares minimization, thus ensuring smooth convergence toward a globally consistent solution. Another advantage of this approach is its robustness to sensor noise and loop closure errors, which are common in 3D data acquisition [11]. Although the method was designed for indoor scenes, the main ideas (robust outlier rejection, global optimization, and noise resistance) make it well-suited for adapting to subsea environments, where sensor noise, turbidity, and limited visibility can degrade pairwise registration accuracy.

While multiway registration does not increase processing speed (in fact, it might even decrease it according to what voxel sizes are used in the calculations–see Methods, Altered Parameters section for a description of voxel size), it significantly enhances model accuracy and consistency, making it a valuable tool for applications where precise reconstruction takes priority over real-time performance [11]. This tradeoff aligns with the needs of subsea asset inspection and mapping, where datasets are typically processed offline, and global accuracy is essential.

According to the multiway theory, ICP is insufficient for cluttered environments, and the multiway approach can detect misalignments in the registration from these algorithms, and also determine how to handle outliers in the cloud [7].

Summary of method selection

Each of the registration methods considered (ICP, multiway Registration) offers advantages for handling the challenges of subsea point cloud data.

The ICP algorithm provides a reliable foundation for subsea asset registration due to its mathematical rigor, independence from training data, and effectiveness in rigid-body alignment when sufficient geometric overlap and a good initial guess are available. It performs well for well-defined structures like pipelines or manifolds, even in the presence of noise [6].

Multiway Registration extends ICP to large-scale or multiview reconstructions by jointly optimizing all pairwise alignments, reducing cumulative drift, and maintaining global consistency. Its robustness to sensor noise and ability to handle partial overlaps make it well-suited for reconstructing complex subsea environments, though at a higher computational cost [7].

Error Evaluation

The primary error metric used in Open3D and other registration studies [6], [7], citezhu2023registration, [10], [12] to evaluate registration quality is the inlier root mean squared error (RMSE) [9]. This value represents the square root of the mean squared distance between all pairs of "corresponding" points in the source and target point clouds.

The root mean squared error (RMSE) between the corresponding 3D points in a target cloud and a source cloud can be written as:

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left[(x_{pi}-x_{2i})^2+(y_{qi}-y_{2i})^2+(z_{qi}-z_{2i})^2\right]} \quad (1)$$

where (x_{pi}, y_{qi}, z_{qi}) is the i-th point in the target cloud, (x_{2i}, y_{2i}, z_{2i}) is the corresponding i-th point in the source cloud, and n is the total number of pairs of points that are compared [13].

Although RMSE is a useful overall indicator of alignment quality, it also has several limitations. For example, ICP may assign correspondences between points that are spatially close but do not represent the same physical location in the actual scene. In these cases, ICP cannot reduce the RMSE any further because the algorithm has already paired each point with its nearest neighbor, even if that neighbor is not the true physical match. This could produce deceptive results. They are either higher because the points are aligned correctly in physical space but still have distance between them, or deceptively low because it only measures the distances between nearest neighbors, which could be very close, although the physical surfaces are misaligned.

A ground-truth reference scene is essential for objectively evaluating the true accuracy of point cloud registration methods. When ground truth is available, alignment errors can be measured directly against the known scene geometry, rather than relying solely on internal metrics such as ICP-derived RMSE, which can be misleading. To enable this form of evaluation, Rochester Institute of Technology's Digital Imaging and Remote Sensing Image Generation (DIRSIG) software is employed [14]. DIRSIG is a physics-based simulation framework that uses path tracing to generate realistic imaging data, including electro-optical, infrared, and LiDAR measurements. By modeling fully synthetic underwater environments and simulating LiDAR acquisition, DIRSIG enables the creation of ground-truth point clouds with precisely known geometry. Controlled, artificial perturbations can then be applied to these point clouds, allowing the accuracy of the ICP and multiway registration algorithms to be quantitatively assessed against a known reference.

Methods

Data

This project uses subsea datasets provided by Kraken Robotics (KR) in the .e57 format, which will be converted into .pcd format for use in the open source library Open3D [9]. The datasets will contain cleaned and preprocessed scans from previous missions, removing noise and outliers by KR's established cleaning methods that remove most outliers and noise associated with turbidity. KR's LiDAR scans do not include any color data, so a CICP method is not useful. The point clouds are high-density (more than 2.5 million points per scan) and each point cloud includes either seabed or subsea structural data (Figure 2).

As seen in Figure 1, the data is collected by sweeping the LiDAR with a pan-tilt unit. However, Figure 1 only represents the "base" scan, whereas the other point clouds were collected at the same tilt, with varying pans. Sometimes there are duplicates of a particular scan pan and tilt angle that are collected multiple times, with varying power of the LiDAR, so that different details in the scene can be collected (Figure 3). There are also "additional" scans, where a pan is repeated, and the tilt angle is changed, so that tall structures or pipelines can be scanned in their entirety. Figure 4 shows an additional scan of a structure, where the tilt angle was higher than the base angle to get the complete structure in the scan.

The number of scans varies per SP, but there are around 20 scans per scene that need to be registered. Each scan has corresponding pan and tilt angle information as to where the LiDAR was aiming, as well as the varying power of the laser.

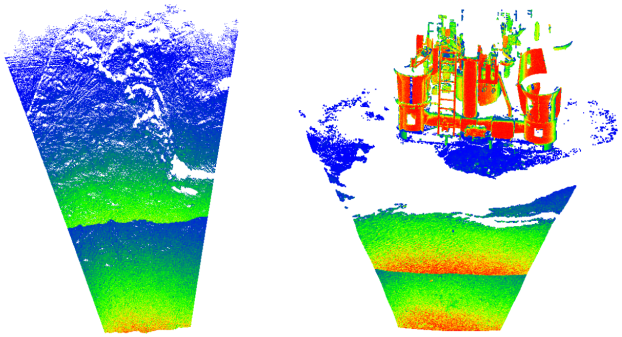

(a) (b)

Figure 2. *Scans of (a) seabed and (b) with points artificially colored by intensity of the pixels return (red higher intensity, blue lower). Green points are considered more reliable than the red or blue points.*

ICP

The mathematical underpinnings of the ICP algorithm are explained in detail in the Open3D documentation [9].

Briefly summarized, the point-to-point algorithm

$$E(T) = \sum_{(p,q)\in K} \|p - Tq\|^2 \tag{2}$$

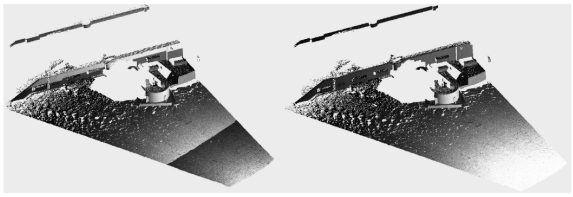

(a) (b)

Figure 3. *Point clouds collected at lower power (a) and a higher power (b)*

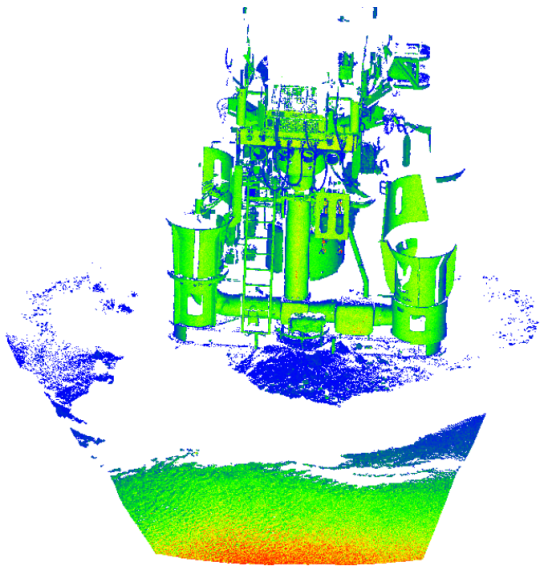

Figure 4. *Additional scan with a higher tilt angle to capture the top of the structure as seen in Figure 2(b)*

where (p,q) are corresponding points in the set K, and T is a transformation. The summed distance between each pair of corresponding points are minimized. When point-to-point was used, the registration_icp function called with the parameter TransformationPointToPoint. It "provides functions to compute the residuals and Jacobian matrices of the point-to-point ICP objective."

The point to plane algorithm

$$E(T) = \sum_{(p,q)\in K} \big((p - Tq)\cdot n_p\big)^2 \tag{3}$$

As used in the TransformationPointToPlane Class. (p,q) are again corresponding points in K, and T is the transform, and n_p is the normals that were estimated on the point cloud. The registration_icp function is called with TransformationPointToPlane. "Internally, this class implements functions to compute the residuals and Jacobian matrices of the point-to-plane ICP objective."

ICP Use Case

The ICP registration workflow uses a hierarchical, tree-based strategy to align all scans within a scan position. Prior to registration, the input scans are divided into base scans and additional scans. The base tilt angle is identified as the most frequently occurring tilt among all scans in the dataset; all scans acquired at this tilt are treated as the base set, and scans acquired at other tilt angles are considered additional.

Multiple scans may share the same pan–tilt position but differ in power level. These duplicates are registered by eigen-based ICP refinement, using the identity matrix as a starting transformation, producing a single consolidated point cloud for each unique pan–tilt combination.

Next, each additional scan is registered to the corresponding base scan of matching pan angle using the same ICP pipeline.

Once all pan–tilt positions have a single consolidated cloud, the base scans are registered together using a tree-style adjacent-pair strategy. Adjacent base scans are registered in pairs, merged, and iteratively combined in successive levels of the tree until a single point cloud remains (Figure 6). This final cloud represents the fully registered scan position.

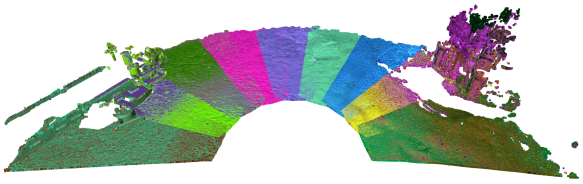

Figure 5. *Point cloud positions in a full scan position without being registered*

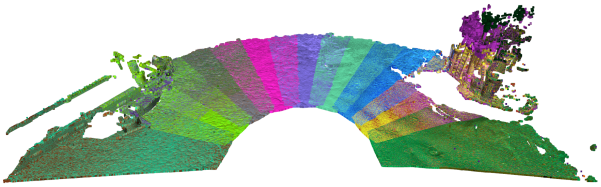

Figure 6. *Final registered point cloud using multiway algorithm*

During registration, the RMSE and standard deviation of each ICP refinement step are collected and averaged to quantify the overall registration error. These metrics allow direct comparison of the tree-registration approach across multiple ICP variants (point-to-point, point-to-plane, and G+ICP). 7 for a workflow diagram of how the ICP algorithm works.

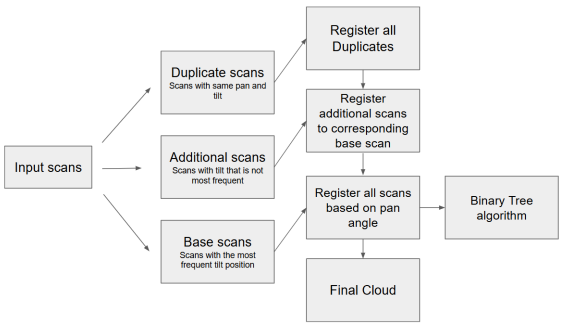

Figure 7. *Workflow chart depicting the workflow of the ICP algorithm.*

Multiway

The Multiway registration method represents an extension of the traditional ICP framework by incorporating pose-graph optimization to jointly estimate the transformations of all scans within a dataset. In this approach, each point cloud is represented as a node in a pose graph, and edges encode the relative transformations obtained through pairwise registration [7].

The goal of Multiway is to find the set of global rigid transformations $\{\mathbf{T}_i\}$ that are undertaken on a set of point clouds $\{\mathbf{P}_i\}$, so that the transformed point clouds $\mathbf{T}_i\mathbf{P}_i$ are aligned in a common coordinate frame (global space) [4]. This is the goal of most registration algorithms.

The pose graph is a large part in Multiway registration, and has two parts; the pose and the node. The **pose**, also thought of as the transformation $\mathbf{T}_i$ that transforms the point cloud, also called the **node**, $\mathbf{P}_i$ into global coordinates. The set of transformations $\{\mathbf{T}_i\}$ is what is to be optimized.

Multiway Use Case

Unlike the tree-based ICP method, the multiway framework does not merge or consolidate intermediate point clouds during processing. Instead, all scans remain distinct until the final optimization stage. Global consistency is achieved by constructing and optimizing a pose graph that encodes relative geometric constraints between scans (Figure 8).

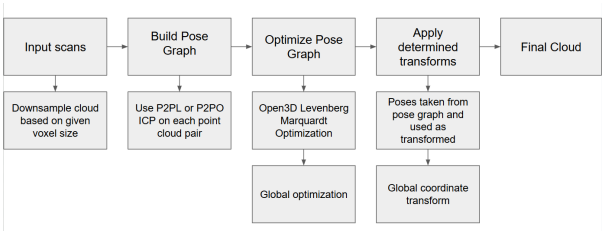

Figure 8. *Workflow chart depicting the workflow of the Multiway algorithm.*

Each scan is represented as a node in the pose graph, storing a rigid transformation that maps the scan into a common reference frame. Edges in the graph represent pairwise registration constraints obtained via ICP and include both the estimated transformation and an associated information matrix that reflects registration confidence.

First, pairwise registration is performed only between spatially adjacent scans, reflecting the acquisition order of the dataset. For each neighboring scan pair, a two-stage coarse-to-fine ICP procedure is applied. The coarse stage uses a large correspondence distance and an identity initialization to obtain a rough alignment, while the fine stage refines this estimate using a smaller correspondence threshold initialized from the coarse result. P2PL (and P2P for comparison) is employed to improve convergence accuracy.

Following fine registration, each candidate edge is evaluated using fitness and inlier RMSE metrics. Only scan pairs that satisfy predefined quality thresholds are incorporated into the pose graph. This gating step prevents poorly constrained or non-overlapping scan pairs from introducing erroneous constraints into the global optimization.

Accepted pairwise registrations are added as edges to the pose graph. Sequential neighbor edges are treated as odometry

constraints. Even when an odometry edge is rejected, the corresponding scan node is retained to preserve graph structure.

Once the pose graph is constructed, a global pose-graph optimization is performed using a Levenberg–Marquardt solver. This optimization computes a set of mutually consistent transformations that minimize the overall registration error while pruning weak constraints.

Registration accuracy is assessed after global optimization by computing the RMSE between adjacent transformed scans. The resulting RMSE values are averaged, and the standard deviation is reported to quantify alignment consistency across the dataset.

Altered Parameters

Several algorithmic parameters play a critical role in determining the accuracy, robustness, and computational performance of point cloud registration. Among the most influential are the maximum correspondence distance, the minimum correspondence distance, and the voxel size used during preprocessing and feature computation.

The maximum correspondence distance defines the radius within which the algorithm searches for a matching point in the target cloud for each point in the source cloud [9]. It restricts the neighborhood over which nearest-neighbor queries are performed. A larger correspondence distance increases the likelihood of identifying a match, which can be beneficial when the initial misalignment between scans is large. However, an overly large threshold may introduce false correspondences, causing points from unrelated surfaces to be incorrectly paired. This can reduce stability, introduce bias into the transformation estimate, and ultimately degrade registration accuracy. Conversely, a smaller correspondence distance enforces stricter matching criteria and helps prevent incorrect associations, but may fail to identify valid correspondences when point clouds are offset or contain substantial noise.

The minimum correspondence distance establishes a lower bound on acceptable point-to-point distances during correspondence formation. Although not always explicitly parameterized in every implementation, this threshold filters out matches that occur at extremely small distances, which may arise from duplicate points, quantization artifacts, or sensor noise. By discarding these near-zero correspondences, the algorithm reduces the influence of degenerate or unstable matches that can skew the estimated transformation. This parameter is particularly important when working with dense or oversampled point clouds, where multiple points may project onto nearly identical locations.

The voxel size is a critical parameter in the downsampling stage and shapes the geometric resolution of the point cloud used for both global registration and local ICP refinement [9]. During voxel downsampling, the point cloud is subdivided into a uniform 3D grid of cubes (voxels), and a single representative point is retained for each voxel. Larger voxel sizes yield a coarser representation, reducing computational cost and improving robustness against noise, but may also remove fine geometric details necessary for accurate alignment. Smaller voxel sizes preserve greater surface fidelity but increase computational burden and sensitivity to noise. In feature-based methods such as FPFH, the voxel size directly influences the locality of the descriptors and the reliability of feature matching. Thus, selecting an appropriate voxel size involves a trade-off between accuracy, robustness, and computational efficiency.

To optimize each method for subsea adjacent scans,each of these parameters is altered and tested for time and accuracy within each registration method.

Comparisons

Each method needs to be compared with the others to determine the best accuracy and performance efficiency. The RMSE and DIRSIG model are used to understand the accuracy of each method.

RMSE

As stated above, RMSE quantifies the average Euclidean distance between corresponding points in two aligned point clouds [9]. Lower RMSE values indicate a closer geometric fit between the registered surfaces. In addition to the mean RMSE, the standard deviation of RMSE provides information about the consistency and reliability of the registration. A low standard deviation suggests that the algorithm performs uniformly across all pairwise registrations or all iterations of a multiway pipeline. Conversely, a high standard deviation indicates variability in registration quality, which may be caused by issues such as uneven overlap, poor initial alignment, or environmental structure (*e.g.*, planar vs. highly textured surfaces). Reporting both RMSE and its standard deviation allows for a more comprehensive evaluation of registration performance.

DIRSIG

To enable an accurate comparison against ground truth, a reference scene with known geometry must first be established. To this end, a synthetic subsea environment was created using the DIRSIG LiDAR simulation tools. An open-source seabed model was combined with 3D ship models placed on the seafloor, resulting in a scene that reasonably approximates a typical KR survey environment, including realistic seafloor texture and representative subsea structures (Figure 9).

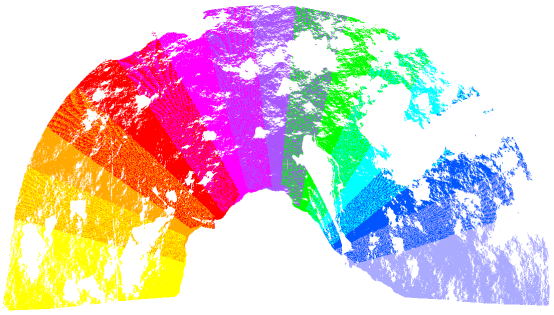

Figure 9. *Artificial point cloud colored by scan. Note the seabed and ships within the point cloud.*

LiDAR data were then generated using DIRSIG's simulated LiDAR sensor, configured with a 532 nm wavelength and collecting data in 30° × 30° swaths at 20° angular increments. This acquisition geometry closely mirrors real-world subsea LiDAR scanning procedures. The resulting point clouds closely resemble those produced by 3DD systems and contain sufficient overlap between adjacent swaths to support the evaluation of both ICP- and

multiway-based registration methods.

When comparing registration methods, a DIRSIG model can provide a controlled, repeatable environment with precise knowledge of the true geometry and sensor configuration. DIRSIG simulation has its own known scene that allows for ground-truth comparison. Since the underlying scene geometry is known exactly, registration error can be assessed in absolute terms rather than relative ones. This error metric is expected to be the best one in terms of actually determining the registration method with the best accuracy.

To quantitatively evaluate registration accuracy, known rigid-body transformations were first applied to the original DIRSIG point clouds to generate artificially perturbed datasets. The magnitude of the perturbations are very small, the translation matrices are very close to an identity matrix, but this is realistic alterations based on the movement of the scanner subsea. These transformations consisted of small translations and rotations representative of realistic ROV motion and sensor drift in a subsea environment, as they were taken from the final transforms applied to point clouds in some of the output of the registration methods. The applied transformations were treated as ground truth.

Both the ICP-based and multiway registration algorithms were then used to estimate transformations that returned each perturbed point cloud to its original pose. Let $\mathbf{T}_{\text{gt}}$ denote the ground-truth transformation applied to a given point cloud, and let $\mathbf{T}_{\text{est}}$ denote the transformation estimated by the registration algorithm. A residual error transformation was computed as

$$\mathbf{T}_{\text{err}} = \mathbf{T}_{\text{est}}^{-1}\mathbf{T}_{\text{gt}} \tag{4}$$

In the ideal case, perfect registration yields $\mathbf{T}_{\text{err}} = \mathbf{I}_{4\text{x}4}$, the 4x4 identity matrix.

The translation error was computed as the Euclidean norm of the translational component of the residual transform,

$$e_t = \|t_{\text{err}}\|_2 \tag{5}$$

where $\mathbf{t}_{\text{err}}$ is the upper right 3x1 translation vector of $\mathbf{T}_{\text{err}}$ [15].

The rotation error was computed from the rotational component of the residual transform. Let $\mathbf{R}_{\text{err}}$ denote the 3x3 rotation matrix extracted from $\mathbf{T}_{\text{err}}$. The angular rotation error was then computed as

$$e_r = \cos^{-1}\left(\frac{\text{tr}(\mathbf{R}_{err})}{2} - 1\right) \tag{6}$$

This results in the minimum rotation angle required to align the estimated orientation with the ground truth. Rotation error is expressed in degrees [16].

These per-cloud translation and rotation errors provide an absolute, physically interpretable measure of registration accuracy and allow direct comparison between the ICP-based and multiway registration approaches.

ICP and multiway registration used different anchor scans, resulting in different global coordinate frames. Since evaluation was performed using frame-invariant pose error metrics, this does not affect the reported accuracy.

Results and Discussion

Each registration method was evaluated on four scan positions to assess accuracy, robustness, and computational performance. Three scan positions (SP1–SP3) consist of cleaned subsea LiDAR datasets, while the fourth (SP4) includes residual noise that could not be removed through standard preprocessing. Each SP was run five times to ensure accurate averages. This setup enables comparison of algorithm performance under both nominal and degraded data conditions. Results are reported using adjacent-scan RMSE, standard deviation, and total processing time, with additional ground-truth validation performed using DIRSIG simulations.

Global Registration

As noted in prior work [8], ICP-based registration pipelines often benefit from an initial coarse alignment step, in which a global registration algorithm provides a rough estimate of the relative transformation before local refinement. This strategy is commonly employed when the relative poses of adjacent scans are unknown or when large initial misalignments are expected.

In the subsea datasets considered in this study, however, global registration provides no measurable benefit. The relative transformations between adjacent scans were small, as the initial scan poses were already well constrained by the LiDAR system's navigation data. As a result, the coarse global alignment produced transformations that were often less accurate than the identity initialization, leading to poorer intermediate alignments (Figure 10).

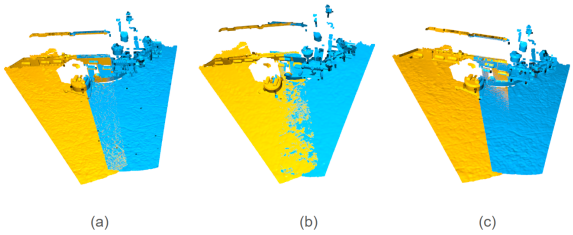

Figure 10. *Original cloud positions (a), global registration result with voxel size 0.04m (b), and global registration result with voxel size 0.025 (c).*

Although global registration improved upon the initial misalignment when the voxel size was decreased, the resulting transformations converged to solutions comparable to those obtained using an identity initialization alone, but at notably higher computational cost. Given the additional runtime and the lack of improvement in final registration accuracy, the use of global registration was determined to be unnecessary for this application.

These results indicate that, for static subsea scan positions with high-quality pose initialization, direct local ICP refinement is sufficient, and global registration does not offer a favorable accuracy–efficiency tradeoff.

Point-to-Point vs. Point-to-Plane

A comparison was performed between point-to-point (P2P) and point-to-plane (P2PL) variants of the ICP algorithm to evaluate their suitability for subsea point cloud registration. Quantitatively, both methods produced similar adjacent-scan RMSE values, with P2PL achieving a slightly lower average RMSE than P2P (3.12 [mm] vs 3.20 [mm], 2.19 [mm] vs 2.22 [mm]).

Despite the comparable RMSE values, visual inspection revealed significant differences in registration quality. In particular, P2P frequently produced visually incorrect alignments in regions dominated by planar seabed geometry, as shown in Figure 11. These misalignments were not reliably reflected in the RMSE metric, highlighting a limitation of nearest-neighbor error measures for evaluating true geometric consistency in largely flat environments.

The effect that P2P is less effective on seabed scans is likely caused by the largely planar, low-feature nature of the seabed. On planar surfaces, multiple relative poses produce similar point-to-point distances, making the optimization ill-conditioned. Because P2P does not use surface orientation, it lacks sufficient constraints, whereas point-to-plane ICP leverages surface normals to better constrain alignment in smooth environments.

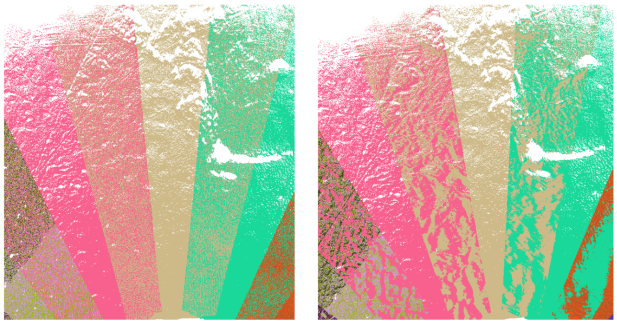

Figure 11. *Point-to-plane registration seabed scans (a) vs point-to-point registration seabed scans.*

In addition to reduced visual accuracy, the P2P implementation required approximately 50% more computation time than P2PL for the same scan position. Given its inferior visual performance and higher computational cost, P2P was deemed unsuitable for this application.

ICP Tree Algorithm

Table 1 summarizes the performance of the ICP tree-based registration across all scan positions. For the cleaned datasets (SP1–SP3), the method achieved millimeter-level accuracy with relatively low variability. Performance (as determined with nearest neighbor RMSE) degraded only slightly in the noisy dataset (SP4), indicating robustness to residual noise at the cost of increased error. All runtime measurements were obtained on a workstation equipped with an AMD Ryzen 5 processor and 16 GB of system memory, with no GPU or hardware acceleration utilized.

Multiway Algorithm

Table 2 shows the results of the multiway registration approach. Compared to the ICP tree method, multiway registration achieved comparable or slightly lower average RMSE values while requiring less processing time. However, increased RMSE variability suggests greater sensitivity to individual pairwise constraints.

Table 1: Resulting data for different scan positions using the ICP tree algorithm after five iterations

Scan Position	Number of Clouds	Average Adjacent RMSE (mm)	Average Adjacent RMSE Stdev (mm)	Average Time (min)
SP1	23	3.12	0.37	7.57
SP2	32	2.20	1.01	15.15
SP3	33	2.48	1.23	10.40
SP4	19	3.20	0.21	9.61

Table 2: Resulting data for different scan positions using multiway algorithm after running five iterations

Scan Position	Average Adjacent RMSE (mm)	Average Adjacent RMSE Stdev (mm)	Average Time (min)
SP1	2.86	0.65	6.79
SP2	2.47	1.03	10.00
SP3	2.89	0.61	9.34
SP4	2.72	0.96	7.06

Comparisons

RMSE

Across all scan positions, the multiway registration method consistently outperformed the ICP-based approach in terms of average adjacent RMSE, RMSE variability (standard deviation), and total computation time. These trends were observed uniformly across both clean datasets and datasets containing residual noise that could not be removed through preprocessing.

In addition to improved quantitative performance, qualitative inspection revealed notable failure cases in the ICP results that were not captured by RMSE alone. In one scan position containing elongated pipe structures spanning the length of the scene, the ICP method produced two scans that were visibly misaligned and placed in incorrect global locations (Figures 12 and 13). Despite these large misregistrations, the corresponding RMSE values remained relatively low due to the reliance on nearest-neighbor distance metrics.

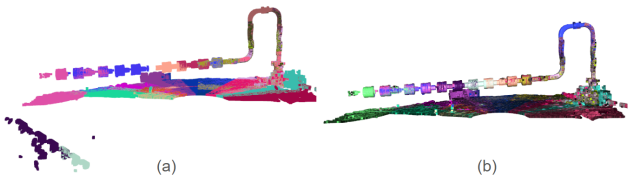

Figure 12. *ICP (a) vs multiway (b) algorithm on the same scan position (SP4). ICP had 2 clouds completely misaligned, even though the RMSE does not depict this error.*

The multiway method did not exhibit these failure modes

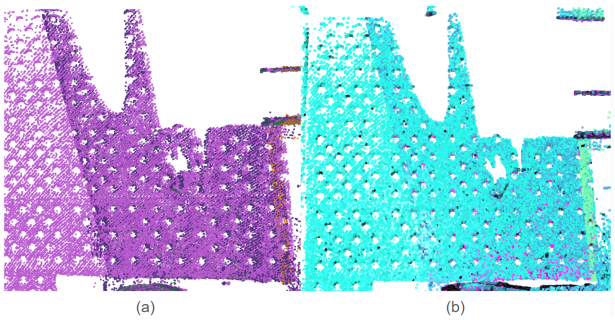

Figure 13. *ICP (a) and multiway (b) registration on the same scan position, showing a visual misalignment in the ICP but not in the multiway. See how in (a) the edge of the dark purple peeks out at the top right edge of the repeating circles, while in (b) the dark blue is evenly seen throughout the light blue.*

and maintained consistent alignment across all scans in the same scene. These results indicate that, while RMSE is a useful local error metric, it is insufficient for identifying global misregistration. Based on both quantitative metrics and qualitative assessment, the multiway approach demonstrates superior robustness and reliability compared to pairwise ICP for subsea point cloud registration.

DIRSIG

To validate registration accuracy against ground truth, experiments were performed using synthetic point clouds generated in DIRSIG. Known rigid transformations were applied to each cloud and subsequently estimated using the ICP and Multiway algorithms. The resulting rotation and translation errors, computed relative to the ground-truth transformations, are summarized in Tables 3 and 4.

Table 3: Errors associated with ICP algorithm and DIRSIG data

Cloud Number	Rotation Error (Deg)	Translation Error (m)
0	0.0000	0.0000
1	0.2007	0.01958
2	0.4053	0.03598
3	0.5506	0.04225
4	0.6867	0.05028
5	0.7943	0.05463
6	0.4552	0.03161
7	0.4539	0.03421
8	0.4548	0.03186

The DIRSIG point clouds have an average point spacing of approximately 10 cm, whereas the real subsea datasets have a spacing of approximately 5–10 mm. Because registration accuracy is fundamentally limited by sampling resolution, the absolute errors measured in the DIRSIG data cannot be directly interpreted at the millimeter scale. Instead, performance was evaluated relative to the sampling interval.

For the DIRSIG experiments, ICP produced translation errors ranging from 0–5.5 cm and Multiway from 0.5–3.5 cm.

Table 4: Errors associated with multiway algorithm and DIRSIG data

Cloud Number	Rotation Error (Deg)	Translation Error (m)
0	0.2865	0.007169
1	0.1452	0.005078
2	0.1304	0.015354
3	0.2003	0.017347
4	0.3390	0.023205
5	0.4700	0.028082
6	0.5877	0.033770
7	0.5847	0.034788
8	0.5945	0.034180

Given the average point spacing of approximately 10 cm, these correspond to roughly 0.3–0.6 and 0.1–0.35 times the sampling interval for ICP and Multiway, respectively. Both methods therefore achieve sub–point-spacing accuracy, with Multiway consistently producing smaller errors.

Registration accuracy in geometry-based methods is fundamentally limited by the spatial resolution of the data, since correspondences cannot be determined more precisely than the local point density and surface sampling allow. For ICP-based approaches, convergence errors typically scale with a fraction of the point spacing, assuming similar scene geometry, noise characteristics, and overlap conditions.

The real subsea datasets used in this study have a point spacing of approximately 5–10 mm, which is 10–20× finer than the DIRSIG data. Assuming similar fractional performance relative to sampling density, the normalized DIRSIG results correspond to an expected alignment accuracy of approximately 2–6 mm for ICP and 1–4 mm for Multiway on real subsea data. Under this interpretation, the Multiway method satisfies the ¡5 mm accuracy requirement and provides improved accuracy compared to pairwise ICP.

Overall, the DIRSIG experiments demonstrate that multiway registration improves accuracy and reduces drift relative to sequential ICP, while achieving sub–sampling-level performance that is expected to translate to millimeter-scale accuracy on the higher-density real datasets.

Conclusions

An automated registration framework for static subsea LiDAR scan positions was developed and evaluated using both real Kraken Robotics datasets and DIRSIG LiDAR simulations. Results demonstrate that pose-graph multiway registration provides superior global consistency, lower error, and faster execution compared to a hierarchical ICP approach. Unlike sequential ICP, the multiway method effectively mitigates error accumulation and prevents large-scale misregistration, even when local RMSE metrics appear acceptable.

Ground-truth evaluation established that multiway registration achieves sub–sampling-level accuracy, corresponding to an estimated 1–4 mm alignment error for high-density subsea data, meeting operational metrology requirements. Addition-

ally, global registration appears to provide no benefit for well-initialized static surveys, and point-to-plane ICP was identified as the preferred local estimator.

These findings indicate that multiway registration is a robust and practical solution for automated subsea metrology workflows. Adoption of this approach could reduce processing time, minimize operator involvement, and improve confidence in delivered measurements, ultimately lowering operational cost and risk.

References

[1] F. Menna *et al.*, "Towards real-time underwater photogrammetry for subsea metrology applications," in *OCEANS 2019 - Marseille*, pp. 1–10, June 2019.

[2] B. Castanier and M. Rausand, "Maintenance optimization for subsea oil pipelines," *International Journal of Pressure Vessels and Piping*, vol. 83, pp. 236–243, Apr. 2006.

[3] 3D at Depth, "SL4 Subsea LiDAR." `https://3datdepth.com/technology/sl4/`, 2025. Accessed: September 18, 2025.

[4] X. Huang, G. Mei, J. Zhang, and R. Abbas, "A comprehensive survey on point cloud registration," *arXiv*, Mar. 2021.

[5] S. Li, D. Su, F. Yang, H. Zhang, X. Wang, and Y. Guo, "Bathymetric LiDAR and multibeam echo-sounding data registration methodology employing a point cloud model," *Applied Ocean Research*, vol. 123, p. 103147, June 2022.

[6] S. Rusinkiewicz and M. Levoy, "Efficient variants of the ICP algorithm," in *Proceedings of the Third International Conference on 3-D Digital Imaging and Modeling*, pp. 145–152, May 2001.

[7] "Multiway registration – Open3D documentation," 2025. Accessed: September 20, 2025.

[8] Z. Zhu, S. Rowlinson, T. Chen, and A. Patching, "Exploring the impact of different registration methods and noise removal on the registration quality of point cloud models in the built environment," *Buildings*, vol. 13, no. 9, p. 2365, 2023.

[9] "Open3D – A modern library for 3D data processing." `https://www.open3d.org/`, 2025. Accessed: September 20, 2025.

[10] S. Choi, T. Kim, and W. Yu, "Performance evaluation of RANSAC family," in *Proceedings of the British Machine Vision Conference*, pp. 81.1–81.12, 2009.

[11] S. Choi, Q.-Y. Zhou, and V. Koltun, "Robust reconstruction of indoor scenes," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5556–5565, June 2015.

[12] B. Göbel, J. Huurdeman, A. Reiterer, and K. Möller, "Robot-based procedure for 3d reconstruction of abdominal organs using the iterative closest point and pose graph algorithms," *Journal of Imaging*, vol. 11, no. 2, 2025.

[13] The MathWorks, Inc., "pcregistericp - register two point clouds using icp algorithm." `https://www.mathworks.com/help/vision/ref/pcregistericp.html`, 2025. Accessed: 2025-12-05.

[14] Rochester Institute of Technology, "Digital Imaging and Remote Sensing Image Generation (DIRSIG)," 2025.

[15] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of RGB-D SLAM systems," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Rystems*, pp. 573–580, IEEE, 2012.

[16] Wikipedia contributors, "Axis–angle representation." `https://en.wikipedia.org/w/index.php?title=Axis%E2%80%93angle_representation&oldid=1316705280`, 2025. Wikipedia, accessed 2026-02-07.

JOIN US AT THE NEXT EI!

electronic IMAGING

Imaging across applications . . . Where industry and academia meet!

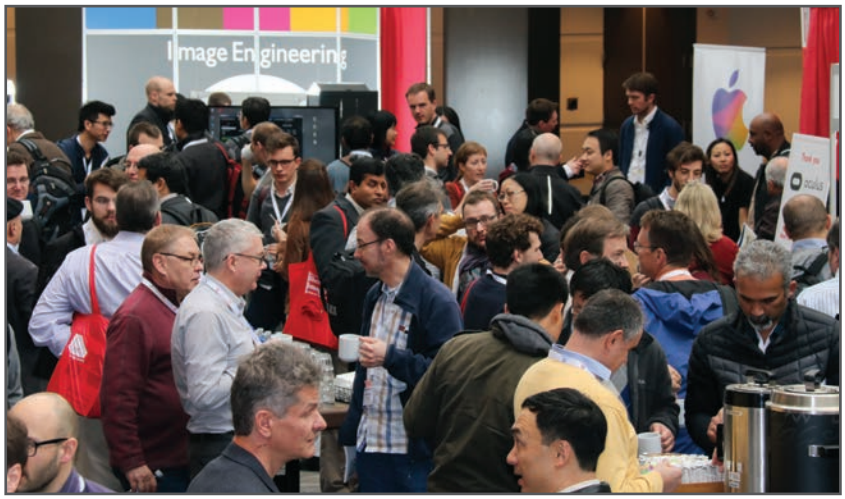

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org

ADVANCING IMAGING SCIENCE 75 1947 IS&T 2022 imaging.org