

Improving Self-Consistency in Underwater Mapping through Laser-Based Loop Closure

Thomas Hitchcox, *Student Member, IEEE*, and James Richard Forbes, *Member, IEEE*

Abstract—Accurate, self-consistent bathymetric maps are needed to monitor changes in subsea environments and infrastructure. These maps are increasingly collected by underwater vehicles, and mapping requires an accurate vehicle navigation solution. Commercial off-the-shelf (COTS) navigation solutions for underwater vehicles often rely on external acoustic sensors for localization, however survey-grade acoustic sensors are expensive to deploy and limit the range of the vehicle. Techniques from the field of simultaneous localization and mapping, particularly loop closures, can improve the quality of the navigation solution over dead-reckoning, but are difficult to integrate into COTS navigation systems. This work presents a method to improve the self-consistency of bathymetric maps by smoothly integrating loop-closure measurements into the state estimate produced by a commercial subsea navigation system. Integration is done using a white-noise-on-acceleration motion prior, without access to raw sensor measurements or proprietary models. Improvements in map self-consistency are shown for both simulated and experimental datasets, including a 3D scan of an underwater shipwreck in Wiarton, Ontario, Canada.

Index Terms—Marine robotics, sensor fusion, SLAM, commercial off-the-shelf (COTS) systems.

I. INTRODUCTION

ACCURATE, self-consistent bathymetric maps are critical for assessing the health of subsea environments and infrastructure. Increasingly, these maps are collected by autonomous underwater vehicles (AUVs) using a variety of on-board sensors, including cameras [1–3], sonar [4–6], and laser scanners [7, 8]. Since the map is constructed using the estimated AUV trajectory, long-term navigation accuracy is a prerequisite for building accurate maps.

The standard navigation solution for commercial AUVs is a commercial off-the-shelf (COTS) inertial navigation system (INS), with acoustic aiding from a Doppler velocity log (DVL). The dead-reckoned precision of these systems is measured by drift rate as a percent of distance traveled, with high-quality DVL-INS systems achieving a drift rate of as low as 0.01%. However, without localizing measurements the precision of the state estimate will deteriorate without bound, impacting long-term accuracy.

Manuscript received 30 May 2022; revised 21 September 2022; accepted 22 November, 2022. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada and in part by Voyis Imaging Inc. through the Collaborative Research and Development program. The work of Thomas Hitchcox was supported by the McGill Engineering Doctoral Award program. This paper was recommended for publication by Associate Editor Maurice Fallon and Editor Francois Chaumette upon evaluation of the reviewers' comments.

T. Hitchcox (corresponding author) and J. R. Forbes are with the Department of Mechanical Engineering, McGill University, Montreal, QC H3A 0C3, Canada. thomas.hitchcox@mail.mcgill.ca, james.richard.forbes@mcgill.ca.




Fig. 1: A point cloud map of an actual shipwreck collected in Wiarton, Ontario, Canada, where colour represents relative depth. The top map is generated using the state estimate produced by a commercial off-the-shelf (COTS) Doppler velocity log-aided inertial measurement system (DVL-INS). The bottom map is generated using the proposed method, which conditions the DVL-INS estimate on loop-closure measurements *without access to the raw DVL-INS sensor measurements*. Note the improvement in map *self-consistency* when loop-closure measurements are included.

Since GPS signals attenuate rapidly in water, AUV localization is primarily done using acoustics [9]. For example, long baseline (LBL) arrays are acoustic beacons installed on the seafloor that trilaterate the position of an AUV, much like an “acoustic GPS.” Short baseline (SBL) and ultrashort baseline (USBL) systems are affixed to a surface vessel, and measure the acoustic range and bearing of an underwater vehicle. These sensors are frequently deployed in a commercial setting, and have been used to aid AUV navigation in the literature, for example [10].

Acoustic positioning systems enable accurate and precise AUV trajectory estimates, however they are expensive to deploy and limit the mission domain of the vehicle. For example, LBL systems are time-consuming to install and calibrate, while SBL and USBL systems require the presence of a large surface vessel. In addition, acoustic positioning

systems produce measurements with limited precision, which may lead to small irregularities in a composite map built from several overlapping measurements of the same area. This in turn may make it difficult to assess relative distances and deformation, or other measurements critical to subsea safety.

A. Motivation

Loop closures play a central role in many simultaneous localization and mapping (SLAM) algorithms, whereby a vehicle returns to and is able to recognize a previously explored region of the map. Loop-closure measurements effectively “reset” any navigation drift accumulated throughout the loop [11], resulting in navigation solutions that are both more accurate and more precise than dead-reckoning, without the need for external localizing measurements. Multiple loop closures over time lead to bounded navigation drift and a more *self-consistent* map estimate, whereby the resulting map is free of irregularities and “double vision” effects produced by poorly aligned measurements, an example of which is shown in Fig. 1. This is not to be confused with the term *consistent*, which in the context of state estimation describes a solution for which the covariance bounds accurately reflect the error in the mean state estimate [12, Sec. 5.4.2].

Previous applications of SLAM for underwater mapping leverage loop-closure measurements to improve map self-consistency. However, these applications have largely been implemented on research platforms with access to raw sensor measurements and full knowledge of the state estimation algorithm. In contrast, commercial “strapdown” DVL-INS systems for subsea navigation produce a state *estimate*, and due to their proprietary nature rarely provide access to

- 1) *raw sensor measurements*, including interoceptive measurements \mathbf{u}_k , such as from an IMU, and exteroceptive measurements \mathbf{y}_ℓ , such as from a DVL;
- 2) a *process model* of the form

$$\tilde{\mathbf{x}}_k = \mathbf{f}_{k-1}(\mathbf{x}_{k-1}, \mathbf{u}_{k-1}), \quad (1)$$

which describes how the vehicle moves throughout time;

- 3) *sensor models* of the form

$$\tilde{\mathbf{y}}_\ell = \mathbf{g}_\ell(\mathbf{x}_\ell, \mathbf{v}_\ell), \quad (2)$$

which allow for predicted measurements; and

- 4) *sensor noise and bias specifications*, for example

$$\mathbf{u}(t) = \bar{\mathbf{u}}(t) + \boldsymbol{\beta}(t) + \mathbf{w}(t), \quad (3a)$$

$$\dot{\boldsymbol{\beta}}(t) \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_{\dot{\boldsymbol{\beta}}}\delta(t-t')), \quad (3b)$$

$$\mathbf{w}(t) \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_w\delta(t-t')), \quad (3c)$$

where \mathbf{u} is known to be corrupted by time-varying random walk bias $\boldsymbol{\beta}$ and Gaussian white noise \mathbf{w} , characterized by power spectral densities $\mathbf{Q}_{\dot{\boldsymbol{\beta}}}$ and \mathbf{Q}_w , respectively.

Commercial DVL-INS systems, for example the Sonardyne SPRINT-Nav 500 [13], are effectively “black boxes,” and their lack of transparency makes it difficult to incorporate loop-closure measurements using conventional state estimation tools [14, 15], as illustrated by the factor graph [16] in Fig. 2.




Fig. 2: A factor graph depicting the output from a commercial DVL-INS. Only the estimated trajectory $\hat{\mathbf{T}}_{0:K}$ and incomplete marginal covariance matrices $\hat{\mathbf{P}}_{0:K}$ are available, leading to the formation of prior factors $\mathbf{e}_{0:K}$. Without factors linking adjacent nodes, loop-closure corrections cannot propagate throughout the graph, and the trajectory cannot be updated.

B. Prior Work

The field of simultaneous localization and mapping has found ample application in the domain of subsea robotics. For example, [4] produced a self-consistent bathymetric map of two hydrothermal vents by aligning point cloud submaps generated using multibeam sonar. A distributed particle mapping algorithm was described in [5], where particle weighting was determined based on the innovation between multibeam sonar measurements and the existing map. However, the map resolution was limited by the selection of a grid cell size. This limitation was later addressed in [17], which adopted Gaussian processes as a map representation. The submap alignment approach was followed by [18], which demonstrated improvements in submap simplification and point cloud alignment in a harbour scanning application. Harbour scanning and surveillance was also the subject of [19], which used a feature-based approach to align point clouds collected from imaging sonar. Recent research has focused on more structured environments, for example ship hull inspection [1, 6, 20–22] and subsea infrastructure [7].

These studies generally had access to the information enumerated in Section I-A, and as a result were able to incorporate loop-closure measurements using conventional state estimation techniques. For example, [4] applied loop-closure measurements within an extended Kalman filtering framework, and enjoyed access to raw navigation sensor measurements as well as a vehicle process model. Individual particles in [5] and [17] were propagated forward using DVL measurements and a constant-velocity motion model. The research platform used in the related ship hull inspection studies [1, 6, 20–22] produced raw DVL, IMU, and depth sensor measurements, while the platform in [3] had access to a variety of raw sensor measurements including stereo vision and profiling sonar. These applications used conventional pose-graph SLAM to incorporate loop-closure measurements produced by various exteroceptive sensors.

C. Contribution

This work describes a novel approach to underwater mapping using a high-resolution laser line scanner and the output of a commercial DVL-INS navigation system. First, this work develops a robust laser-based front-end algorithm to produce high-precision loop-closure measurements by aligning point cloud scans collected in challenging underwater

environments. Next, this work shows how to cleanly fuse loop-closure measurements into the output of a survey-grade COTS DVL-INS system via factor graph optimization. As these commercial systems are typically “black boxes” which only provide a navigation *estimate*, the proposed approach shows how to systematically incorporate loop-closure measurements *without access to raw sensor measurements or other information typically required in state estimation tasks*. In contrast to previous approaches, the proposed methodology also enforces a *smoothness* requirement on the posterior trajectory estimate. This eliminates discontinuities often encountered in dead-reckoned trajectory estimates, and is critical for accurate feature detection in laser submaps. In summary, the proposed methodology describes a robust and comprehensive system for high-precision, self-consistent underwater mapping using COTS navigation systems. Improvements to both map self-consistency and the relative accuracy of the trajectory estimate are rigorously evaluated in simulation and on an actual underwater mapping dataset.

D. Paper Organization

This paper is organized as follows. Section II contains preliminary information on conventions used, state estimation on matrix Lie groups, and batch state estimation. Section III introduces the methodology, including the formulation of loop-closure measurements from laser scan data and the construction of the batch optimization problem. Section IV contains results on simulated and field datasets. The paper concludes in Section V with a review of the findings and opportunities for future work.

II. PRELIMINARIES

A. Reference Frames and Navigation Conventions

This section discusses the conventions for attitude and displacement used in this paper. A three-dimensional dextral reference frame \mathcal{F}_a is composed of three orthonormal physical basis vectors. The position of physical point z relative to point w , denoted by \vec{r}^{zw} , is resolved in reference frame \mathcal{F}_a as \mathbf{r}_a^{zw} and in reference frame \mathcal{F}_b as \mathbf{r}_b^{zw} . These quantities are related via $\mathbf{r}_a^{zw} = \mathbf{C}_{ab}\mathbf{r}_b^{zw}$, where \mathbf{C}_{ab} is a direction cosine matrix, $\mathbf{C} \in SO(3) = \{\mathbf{C} \in \mathbb{R}^{3 \times 3} \mid \mathbf{C}\mathbf{C}^T = \mathbf{I}, \det \mathbf{C} = +1\}$ [23]. Time-varying quantities are indicated by the subscript $(\cdot)_k$, for example $\mathbf{r}_a^{z_k w}$ describes the position of moving point z at time t_k . In this work, point z is affixed to the vehicle, while w denotes the stationary point in the world. Body frame \mathcal{F}_b rotates with the vehicle, while local geodetic frame \mathcal{F}_a remains stationary. Both \mathcal{F}_b and \mathcal{F}_a are north-east-down (NED), in agreement with maritime convention.

B. Matrix Lie Groups

The attitude and position of a vehicle at time t_k , collectively referred to as the vehicle’s “pose,” may be conveniently represented in 3D space as an element of matrix Lie group $SE(3)$ [23, Sec. 7.1.1],

$$\mathbf{T}_{ab_k}^{z_k w} = \begin{bmatrix} \mathbf{C}_{ab_k} & \mathbf{r}_a^{z_k w} \\ \mathbf{0} & 1 \end{bmatrix} \in SE(3), \quad (4)$$

with $SE(3) = \{\mathbf{T} \in \mathbb{R}^{4 \times 4} \mid \mathbf{C} \in SO(3), \mathbf{r} \in \mathbb{R}^3\}$. Associated with every matrix Lie group is a matrix Lie algebra, defined as the tangent space at the group identity [24]. For $SE(3)$, this is $\mathfrak{se}(3) \triangleq T_1 SE(3)$. For estimation problems involving matrix Lie groups, the matrix Lie algebra is a convenient space to represent perturbations and uncertainty. An element of $\mathfrak{se}(3)$ is given by [25, Sec. 2.3]

$$\boldsymbol{\xi}^\wedge = \begin{bmatrix} \boldsymbol{\phi} \\ \boldsymbol{\rho} \end{bmatrix}^\wedge = \begin{bmatrix} 0 & -\phi_3 & \phi_2 & \rho_1 \\ \phi_3 & 0 & -\phi_1 & \rho_2 \\ -\phi_2 & \phi_1 & 0 & \rho_3 \\ 0 & 0 & 0 & 0 \end{bmatrix} \in \mathfrak{se}(3), \quad (5)$$

where $(\cdot)^\wedge : \mathbb{R}^6 \rightarrow \mathfrak{se}(3)$ is an isometric operator. The inverse of this operator is $(\cdot)^\vee : \mathfrak{se}(3) \rightarrow \mathbb{R}^6$, such that $(\boldsymbol{\xi}^\wedge)^\vee = \boldsymbol{\xi}$. A Lie group and Lie algebra are related through the exponential map, which for matrix Lie groups is the matrix exponential,

$$\mathbf{T} = \exp(\boldsymbol{\xi}^\wedge). \quad (6)$$

The matrix logarithm is used to return to the Lie algebra via

$$\boldsymbol{\xi}^\wedge = \log(\mathbf{T}). \quad (7)$$

Elements of the matrix Lie algebra are combined according to the Baker-Campbell-Hausdorff (BCH) equation,

$$\boldsymbol{\gamma}^\wedge = \log(\exp(\boldsymbol{\xi}^\wedge)\exp(\boldsymbol{\eta}^\wedge)). \quad (8)$$

An approximation to the BCH equation for $\boldsymbol{\xi} \gg \boldsymbol{\eta}$ is

$$\boldsymbol{\gamma}^\wedge \approx (\boldsymbol{\xi} + \mathbf{J}^r(\boldsymbol{\xi})^{-1}\boldsymbol{\eta})^\wedge, \quad (9)$$

where \mathbf{J}^r is the right Jacobian of $SE(3)$ [23, Sec. 7.1.5].

Errors on matrix Lie groups are defined multiplicatively. This work uses the left-invariant error definition,

$$\delta\mathbf{T} = \mathbf{T}^{-1}\tilde{\mathbf{T}}, \quad (10)$$

where \mathbf{T} is the current state estimate and $\tilde{\mathbf{T}}$ is a state estimate generated from sensor measurements or prior information. The corresponding perturbation scheme is

$$\mathbf{T} = \tilde{\mathbf{T}}\exp(-\delta\boldsymbol{\xi}^\wedge), \quad (11)$$

with perturbation $\delta\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{P})$, $\mathbf{P} = \mathbb{E}[\delta\boldsymbol{\xi}\delta\boldsymbol{\xi}^T] \in \mathbb{R}^{6 \times 6}$. Note the negative sign in (11) ensures consistency with the left-invariant error definition (10). The state estimate is therefore defined by mean estimate $\tilde{\mathbf{T}}$ and covariance \mathbf{P} .

This work makes frequent use of the adjoint matrix $\mathbf{Ad}(\mathbf{T})$, which maps perturbations about the group identity to other group elements [24]. Formally,

$$\mathbf{Ad}(\mathbf{T})\delta\boldsymbol{\xi} \triangleq \left(\mathbf{T}\delta\boldsymbol{\xi}^\wedge\mathbf{T}^{-1}\right)^\vee. \quad (12)$$

For $SE(3)$, the adjoint matrix is [25]

$$\mathbf{Ad}(\mathbf{T}) = \begin{bmatrix} \mathbf{C} & \mathbf{0} \\ \mathbf{r}^\times \mathbf{C} & \mathbf{C} \end{bmatrix}, \quad (13)$$

where $(\cdot)^\times$ is the skew-symmetric operator [23, Sec. 7.1.2]. The adjoint matrix is represented in the matrix Lie algebra as

$$(\mathbf{ad}(\boldsymbol{\xi}_1^\wedge)\boldsymbol{\xi}_2)^\wedge \triangleq [\boldsymbol{\xi}_1^\wedge, \boldsymbol{\xi}_2^\wedge] = \boldsymbol{\xi}_1^\wedge\boldsymbol{\xi}_2^\wedge - \boldsymbol{\xi}_2^\wedge\boldsymbol{\xi}_1^\wedge, \quad (14)$$

where $[\cdot, \cdot]$ is the Lie bracket [26, Sec. 10.2.6]. For $\mathfrak{se}(3)$,

$$\mathbf{ad}(\boldsymbol{\xi}^\wedge) = \begin{bmatrix} \boldsymbol{\phi}^\times & \mathbf{0} \\ \mathbf{r}^\times & \boldsymbol{\phi}^\times \end{bmatrix}. \quad (15)$$

C. Gaussian Processes

A continuous-time Gaussian process (GP) may be viewed as a distribution over functions,

$$\mathbf{f}(t) \sim \mathcal{GP}(\boldsymbol{\mu}(t), \boldsymbol{\Sigma}(t, t')), \quad (16)$$

where $\boldsymbol{\mu}(t)$ is the mean function, and $\boldsymbol{\Sigma}(t, t')$ is the covariance function [27, Sec. 2.2][23, Sec. 2.3]. For any finite collection of time steps $t_{0:K}$, $\mathbf{f}(t_{0:K})$ follows a joint Gaussian distribution. The covariance function determines how individual function samples $\mathbf{f}_i(t)$ covary over time. For example, a GP for which the covariance over time is large will be smoother than a GP for which the covariance over time is small. This work uses the zero-mean white noise GP, given by [23, Sec. 2.3]

$$\mathbf{w}(t) \sim \mathcal{GP}(\mathbf{0}, \mathcal{Q}\delta(t - t')), \quad (17)$$

where \mathcal{Q} is a power spectral density matrix and $\delta(\cdot)$ is the Dirac delta function.

D. The White-Noise-On-Acceleration Motion Prior

The white-noise-on-acceleration (WNOA) motion prior may be summarized by the following set of stochastic differential equations [28],

$$\dot{\mathbf{T}}(t) = \mathbf{T}(t)\boldsymbol{\varpi}_b(t)^\wedge, \quad (18a)$$

$$\dot{\boldsymbol{\varpi}}_b(t) \sim \mathcal{GP}(\mathbf{0}, \mathcal{Q}\delta(t - t')). \quad (18b)$$

Equation (18a) describes the continuous-time state kinematics for $SE(3)$, with $\boldsymbol{\varpi}_b$ the *generalized velocity*, such that $T\boldsymbol{\varpi}_b^\wedge \in \mathfrak{se}(3)$, with T a time increment. The subscript $(\cdot)_b$ has been included to emphasize that $\boldsymbol{\varpi}$ is a body-frame quantity. The time rate of change of $\boldsymbol{\varpi}$ is distributed according to the zero-mean white noise Gaussian process in (18b), with power spectral density \mathcal{Q} . Note that \mathcal{Q} is a hyperparameter that needs to be tuned. This motion prior helps to enforce smoothness in the posterior state estimate, as $E[\dot{\boldsymbol{\varpi}}] = \mathbf{0}$. In discrete time, this implies $E[\boldsymbol{\varpi}_k] = \boldsymbol{\varpi}_{k-1}$. The white noise assumption also preserves sparsity in the upcoming batch problem [29].

The WNOA assumption is reasonable in the context of subsea navigation, as AUV kinematics evolve slowly over time. With the inclusion of $\boldsymbol{\varpi}$, the augmented navigation state becomes the ordered pair

$$\mathbf{X} = (\mathbf{T}, \boldsymbol{\varpi}) \in SE(3) \times \mathbb{R}^6. \quad (19)$$

E. Batch State Estimation

Given a set of exteroceptive measurements $\{\mathbf{y}_\ell\}_{\ell=1}^L$, interoceptive measurements $\{\mathbf{u}_k\}_{k=0}^{K-1}$, and prior estimate $\mathbf{Y}_0 = \bar{\mathbf{Y}}_0 \exp(-\delta\boldsymbol{\eta}_0^\wedge)$, $\mathbf{S}_0 = E[\delta\boldsymbol{\eta}_0 \delta\boldsymbol{\eta}_0^\top]$, the standard approach to batch estimation is to produce a maximum a posteriori (MAP) solution, given by

$$\hat{\mathbf{X}} = \arg \max_{\mathbf{X}} p(\mathbf{X} | \mathbf{y}_{1:L}, \mathbf{u}_{0:K-1}, \mathbf{Y}_0). \quad (20)$$

Under the Markov assumption, the joint probability in (20) may be factored as

$$\hat{\mathbf{X}} = \arg \max_{\mathbf{X}} \prod_{\ell=1}^L p(\mathbf{y}_\ell | \mathbf{X}_\ell) \prod_{k=1}^K p(\mathbf{X}_k | \mathbf{X}_{k-1}, \mathbf{u}_{k-1}) p(\mathbf{X}_0 | \mathbf{Y}_0). \quad (21)$$

Taking the negative log likelihood of (21) results in a nonlinear least-squares problem of the form

$$\hat{\mathbf{X}} = \arg \min_{\mathbf{X}} J(\mathbf{X}), \quad (22)$$

where the objective function $J(\mathbf{X})$ is given by

$$J(\mathbf{X}) = \frac{1}{2} \sum_{\ell=1}^L \|\mathbf{e}_\ell(\bar{\mathbf{y}}_\ell, \mathbf{g}_\ell(\bar{\mathbf{X}}_\ell, \mathbf{0}))\|_{\mathbf{R}_\ell^{-1}}^2 + \frac{1}{2} \|\mathbf{e}_0(\bar{\mathbf{Y}}_0, \bar{\mathbf{X}}_0)\|_{\mathbf{S}_0^{-1}}^2 + \frac{1}{2} \sum_{k=1}^K \|\mathbf{e}_k(\mathbf{f}_{k-1}(\bar{\mathbf{X}}_{k-1}, \bar{\mathbf{u}}_{k-1}), \bar{\mathbf{X}}_k)\|_{\mathbf{Q}_k^{-1}}^2. \quad (23)$$

In (23), \mathbf{e}_k , \mathbf{e}_ℓ , and \mathbf{e}_0 are the interoceptive, exteroceptive, and prior errors, respectively, while \mathbf{f}_{k-1} and \mathbf{g}_ℓ represent the nonlinear process and measurement models, respectively. The notation $\|\mathbf{e}\|_{\boldsymbol{\Sigma}^{-1}}^2 = \mathbf{e}^\top \boldsymbol{\Sigma}^{-1} \mathbf{e}$ denotes the squared Mahalanobis distance, and \mathbf{Q}_k and \mathbf{R}_ℓ represent the discrete-time covariance on the interoceptive and exteroceptive errors, respectively. To minimize (22), (23) is repeatedly linearized about the current state estimate $\bar{\mathbf{X}}$, and the local minimizing solution found using, for example, Gauss-Newton or Levenberg-Marquardt.

III. METHODOLOGY

This section describes the primary contributions of this paper, namely the formulation of laser-based loop-closure measurements and the smooth incorporation of these measurements into a COTS DVL-INS trajectory estimate. An overview of the upcoming methodology is shown in Fig. 3.

A. Loop Closures from Subsea Point Cloud Scans

To correct for drift in the DVL-INS trajectory estimate, loop-closure measurements are obtained by aligning sections of the point cloud scan collected using a Voyis Insight Pro underwater laser scanner. The raw laser profiles are first filtered and registered to the trajectory estimate to produce a 3D point cloud. Loop-closure opportunities are identified at path crossings, and alignment is performed using a multi-step point cloud alignment algorithm.

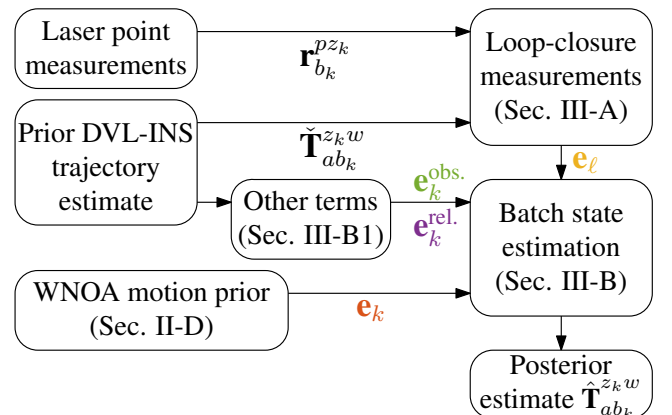


Fig. 3: A visual overview of Section III. Note the colour of the error terms is consistent with the factor graph of Fig. 6.

1) Point Cloud Generation

The Voyis Insight Pro underwater laser scanner, pictured in Fig. 4, records 2D profile measurements of the seabed at a frequency of 20 Hz. To construct a 3D point cloud, individual laser profiles are registered to the prior DVL-INS trajectory estimate $\check{\mathbf{T}}_{ab_k}^{z_k w}$ via

$$\begin{bmatrix} \mathbf{r}_a^{pw} \\ 1 \end{bmatrix} = \check{\mathbf{T}}_{ab_k}^{z_k w} \mathbf{T}_{bl}^{sz} \begin{bmatrix} \mathbf{r}_{l_k}^{psk} \\ 1 \end{bmatrix}, \quad (24)$$

where $\mathbf{r}_{l_k}^{psk} \in \mathbb{R}^3$ is a laser measurement of point p at time t_k resolved in the sensor frame, and $\mathbf{T}_{bl}^{sz} \in SE(3)$ is a static extrinsics matrix. Where necessary, the DVL-INS trajectory is interpolated according to [30, Sec. 2.4]

$$\check{\mathbf{T}}_j = \check{\mathbf{T}}_i \exp\left(\alpha \log\left(\check{\mathbf{T}}_i^{-1} \check{\mathbf{T}}_k\right)\right), \quad (25a)$$

$$\alpha = \frac{t_j - t_i}{t_k - t_i}, \quad (25b)$$

where $\check{\mathbf{T}}_j = \check{\mathbf{T}}_{ab_j}^{z_j w}$, and $t_i < t_j < t_k$. The result of these operations is a filtered point cloud \mathcal{P} resolved in the local geodetic frame, $\mathcal{P} = \{\mathbf{r}_a^{p_i w}\}_{i=1}^N$.

2) Point Cloud Alignment

The objective of point cloud alignment is to combine partially overlapping scans of the same 3D object or scene. In the context of SLAM, point cloud alignment is often performed to estimate the relative pose between two or more observations, for example to reduce odometry drift [31] or to bound navigation drift over time by closing large loops in the trajectory [32]. More formally, the problem of point cloud alignment may be expressed as

$$\mathbf{T}_{12}^* = \arg \min_{\mathbf{T} \in SE(3)} \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^M b_{ij} \cdot w_{ij} \cdot \|\mathbf{e}_{ij}(\mathbf{T}_{12}, \mathbf{r}_{b_2}^{p_i z_2}, \mathbf{r}_{b_1}^{p_j z_1})\|_{\Sigma_{ij}^{-1}}^2, \quad (26)$$

where the pose $\mathbf{T}_{12}^* = (\mathbf{T}_{b_1 b_2}^{z_2 z_1})^*$ optimally aligns source cloud $\mathcal{S} = \{\mathbf{r}_{b_2}^{p_i z_2}\}_{i=1}^N$ to target cloud $\mathcal{T} = \{\mathbf{r}_{b_1}^{p_j z_1}\}_{j=1}^M$. The Boolean value $b_{ij} = \{0, 1\}$ assumes a value of 1 if (p_i, p_j) represents an inlier correspondence, while $w_{ij} \in [0, 1]$ is a correspondence weight, often computed using a robust cost function [33]. \mathbf{T}_{12}^* is optimal in the sense that it minimizes the sum of squared weighted errors, often a combination of point-to-point and point-to-plane errors [34, 35], with associated error covariance $\Sigma_{ij}(\mathbf{R}_i, \mathbf{R}_j)$. \mathbf{R}_i and \mathbf{R}_j represent the covariance on the point measurements $\mathbf{r}_{b_2}^{p_i z_2}$ and $\mathbf{r}_{b_1}^{p_j z_1}$, respectively, with $\mathbf{R}_i = E[\delta \mathbf{r}_i \delta \mathbf{r}_i^T]$, $\delta \mathbf{r}_i = \mathbf{r}_{b_2}^{p_i z_2} - \check{\mathbf{r}}_{b_2}^{p_i z_2}$. A depiction of the point cloud alignment problem is shown in Fig. 5.



Fig. 4: An Insight Pro underwater laser scanner developed by Voyis Imaging Inc. The beam emitter is on the left, while the camera is on the right. 3D point clouds are produced by triangulating the laser beam. The baseline between the emitter and the camera is approximately 1 m.




Fig. 5: Generating a loop-closure measurement by aligning source submap \mathcal{S} to target submap \mathcal{T} . The vehicle trajectory appears as a dashed line. Submaps $\mathcal{T} = \{\mathbf{r}_{b_1}^{p_j z_1}\}_{j=1}^M$ and $\mathcal{S} = \{\mathbf{r}_{b_2}^{p_i z_2}\}_{i=1}^N$ are constructed from the point measurements at vehicle poses $\mathbf{T}_{ab_1}^{z_1 w}$ and $\mathbf{T}_{ab_2}^{z_2 w}$, respectively. Point cloud alignment produces the loop-closure measurement $\mathbf{T}_{b_1 b_2}^{z_2 z_1}$.

In this work, loop-closure locations are identified at simple path crossings on the (x, y) plane, at time stamps t_{ℓ_1} and t_{ℓ_2} , with $t_{\ell_1} < t_{\ell_2}$. Ordinarily, cross-covariance information would be used to determine the search region, and thus the required size of the submaps to construct, using a squared Mahalanobis distance test [36], however this information is absent from the DVL-INS trajectory estimate. Instead, given the inherently low drift rate of the DVL-INS [13], the source and target clouds are constructed using a simple (x, y) distance threshold, e.g.

$$\mathbf{r}_a^{pw} \in \mathcal{T} \mid \|\begin{bmatrix} 1 & \mathbf{0} \end{bmatrix} (\mathbf{r}_a^{pw} - \check{\mathbf{r}}_a^{z_{\ell_1} w})\|_2 \leq \delta r^*. \quad (27)$$

In this work a constant value of $\delta r^* = 5$ m appears to work well, however a gradually increasing threshold related to the length of the trajectory could also be used.

To provide a body-frame relative pose measurement (26), the point measurements are first resolved in the body frames,

$$\begin{bmatrix} \mathbf{r}_{b_\ell}^{p z_\ell} \\ 1 \end{bmatrix} = (\check{\mathbf{T}}_{ab_\ell}^{z_\ell w})^{-1} \begin{bmatrix} \mathbf{r}_a^{pw} \\ 1 \end{bmatrix}. \quad (28)$$

The point clouds are preprocessed by downsampling to a 5 cm grid, which reduces the amount of point data by approximately a factor of 10 while still preserving high-frequency features of the scanned object. Normal vectors are then estimated using the 40 nearest Euclidean neighbours. To account for cases of large navigation drift between observations, the TEASER++ coarse alignment algorithm [37] is used to initialize an iterative closest point (ICP)-based fine alignment algorithm. To run TEASER++, PPFH feature descriptors [38] are computed at 3D SIFT keypoints [39], and a set of putative correspondences is formed from the 10 nearest neighbour matches in 33-dimensional PPFH space. Keypoints and descriptors are computed using the Point Cloud Library (PCL) v1.9 [40]. Default values are used for all TEASER++ parameters.

The combination of SIFT keypoints and PPFH descriptors was selected for this application following an alignment study on the shipwreck field dataset introduced in Section IV-D.

In this dataset, a vehicle makes eight passes over a small shipwreck, producing eight point cloud submaps and 28 unique submap pairs. 27 of the 28 pairs were then aligned using TEASER++ and different detector/descriptor combinations, with one submap pair excluded due to insufficient overlap. The study includes three keypoint detectors and two 3D feature descriptors. The keypoint detectors are SIFT, ISS [41], and Harris 3D keypoints [42], while the feature descriptors are FPFH and SHOT [43]. These detectors and descriptors were included in the study both due to their prevalence in the point cloud alignment literature and the availability of an open-source implementation in PCL v1.9. SIFT and Harris 3D keypoint parameters were tuned slightly to obtain several hundred keypoints in each submap, while default values were used for ISS keypoints. For fairness, both FPFH and SHOT descriptors used the same search radius value of 0.25 m.

Each submap pair was then aligned by TEASER++ using each of the six detector/descriptor combinations. The results are given in Table I, which lists summary statistics on attitude errors $\|\delta\phi\|$ and position errors $\|\delta\rho\|$ in the format 50% · 90% · MAX. Pose errors were computed between each TEASER++ relative pose estimate $\tilde{\mathbf{T}}_i$ and the ground-truth relative pose \mathbf{T}_i , computed from a well-initialized ICP alignment, as

$$\delta\xi_i = \log(\mathbf{T}_i^{-1}\tilde{\mathbf{T}}_i)^\vee. \quad (29)$$

Examining Table I, the combination of SIFT keypoints and SHOT descriptors (second row) delivers the lowest median attitude error (0.42 deg), as well as the lowest 50% and 90% position errors. However, this combination produced at least one outlier measurement from the 27 submap pairs, while the SIFT+FPFH combination (first row) produced zero outliers. In addition, the SIFT+FPFH combination produced reasonable median and 90% errors. Note the extremely large position errors in Table I are due to failed alignments producing a 180 deg “flip” of the (relatively flat) point cloud submaps. The submaps are measured at a range of approximately 7 m, thus “flipped” alignments produce a relative body-frame position error of more than twice this value.

As the objective of a coarse alignment algorithm is to robustly initialize ICP as close to ground-truth as possible,

TABLE I: Summary statistics from the keypoint detector and descriptor alignment study, reported in the format 50% · 90% · MAX. Attitude errors $\|\delta\phi\|$ and position errors $\|\delta\rho\|$ are computed according to (29). The lowest value in each column is indicated in bold font. Si = SIFT, I = ISS, H = Harris3D, F = FPFH, and So = SHOT.

KP	D	$\ \delta\phi\ $ [deg]			$\ \delta\rho\ $ [m]		
Si	F	0.44	· 1.05	· 1.54	0.04	· 0.09	· 0.16
	So	0.42	· 0.86	· 92.59	0.03	· 0.07	· 3.33
I	F	0.49	· 0.84	· 52.34	0.04	· 0.09	· 2.31
	So	0.61	· 176.28	· 179.72	0.05	· 20.12	· 22.56
H	F	0.83	· 21.14	· 178.44	0.08	· 3.05	· 22.48
	So	0.83	· 52.32	· 164.84	0.06	· 4.22	· 18.28

TABLE II: ICP preprocessing and alignment parameters

Stage	Configuration	Description
Preprocessing	VoxelGrid	Downsample to 5 cm grid
	Normals	40 nearest neighbours
Keypoints	3D SIFT	PCL v1.9 implementation
Descriptors	FPFH	PCL v1.9 implementation
Coarse align.	TEASER++	10 matches, default params.
ICP data assn.	KDTree	Single nearest neighbour
ICP error min.	Mixed	Pt-Pt if $v(p_j) < 3e-2$
Outlier reject.	FRMSD	Default params. from [49]
Termination	Diff.	$\ \delta\phi_i\ _2 < 1e-2$ rad, and $\ \delta\rho_i\ _2 < 1e-3$ m
	Counter	20 iterations max

the SIFT+FPFH combination was selected for this application. Note that TEASER++ was chosen for the coarse alignment algorithm as it has been shown in extensive point cloud alignment studies [37] to outperform other coarse alignment methods, for example FGR [44] and RANSAC [45].

For the fine alignment step, this work uses the Weighted Optimal Linear Attitude and Translation Estimator (WOLATE) algorithm [46] within an ICP-based alignment scheme. Alignment errors are formulated between each point in the source cloud and their single nearest neighbour in the target cloud. A combination of point-to-point and point-to-plane errors are used, with the surface variation $v(p_j)$ [47] of the target points determining the type of error used for each association. Following the study in [48], the Fractional Root Mean Squared Distance (FRMSD) robust cost function [49] is used for outlier rejection when aligning structured scans, such as shipwrecks. The algorithm terminates when the pose differential $\delta\xi_i$ between two successive iterations falls below a threshold, or when a maximum number of iterations is reached. Following the recommendations in [50] for best practices when reporting ICP algorithms, the preprocessing steps and relevant parameters are summarized in Table II.

3) Loop-Closure Measurement Model

Point cloud alignment yields the loop-closure measurement

$$\Xi_{\ell_1\ell_2} \triangleq \mathbf{T}_{b_{\ell_1}b_{\ell_2}}^{z_{\ell_2}z_{\ell_1}} = \left(\mathbf{T}_{ab_{\ell_1}}^{z_{\ell_1}w}\right)^{-1} \mathbf{T}_{ab_{\ell_2}}^{z_{\ell_2}w} \in SE(3), \quad (30)$$

and, given the perturbation scheme (11), the noise model is

$$\Xi_{\ell_1\ell_2} = \mathbf{g}_\ell(\check{\mathbf{T}}_{\ell_1}, \check{\mathbf{T}}_{\ell_2}, \delta\xi_\Xi) \quad (31a)$$

$$= \check{\Xi}_{\ell_1\ell_2} \exp(-\delta\xi_\Xi^\wedge), \quad (31b)$$

$$\delta\xi_\Xi \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_\Xi), \quad (31c)$$

where the shorthand $\check{\mathbf{T}}_{\ell_i} = \check{\mathbf{T}}_{ab_{\ell_i}}^{z_{\ell_i}w}$, $i = 1, 2$ is used for readability. The covariance \mathbf{R}_Ξ on the loop-closure measurement may be obtained from the point cloud alignment algorithm in a number of ways, for example the linearization-based approach in [51].

B. Updating the Trajectory

1) Formulating the Objective Function

The objective is now to condition the prior DVL-INS trajectory estimate on the newly available loop-closure measurements. This is accomplished through nonlinear batch state estimation, described in Section II-E. This section describes how the error terms in the batch problem are formulated, and Fig. 6 shows the resulting factor graph.

First, the prior, process, and measurement errors must be defined. Given the augmented navigation state (19), errors must be defined for the $SE(3)$ pose and for the generalized velocity. Using both the left-invariant error definition (10) and the constant velocity WNOA motion prior, the prior error is

$$\mathbf{e}_0 = \begin{bmatrix} \mathbf{e}_0^\xi \\ \mathbf{e}_0^\varpi \end{bmatrix} = \begin{bmatrix} \log(\mathbf{T}_0^{-1} \mathbf{Y}_0)^\vee \\ \varpi_0 - \psi_0 \end{bmatrix}, \quad (32)$$

where (\mathbf{Y}_0, ψ_0) is the prior estimate on the first navigation state. The process errors take the form

$$\mathbf{e}_k = \begin{bmatrix} \mathbf{e}_k^\xi \\ \mathbf{e}_k^\varpi \end{bmatrix} = \begin{bmatrix} \log(\mathbf{T}_k^{-1} \tilde{\mathbf{T}}_k)^\vee \\ \varpi_k - \varpi_{k-1} \end{bmatrix}, \quad (33)$$

where the predicted pose at time t_k ,

$$\tilde{\mathbf{T}}_k = \mathbf{f}_{k-1}(\mathbf{T}_{k-1}, \varpi_{k-1}) = \mathbf{T}_{k-1} \exp(T \varpi_{k-1}^\wedge), \quad (34)$$

arises from a forward Euler discretization of the continuous-time $SE(3)$ kinematics (18a) over an integration period of $T = t_k - t_{k-1}$. The loop-closure errors are

$$\mathbf{e}_\ell = \mathbf{T}_{\ell_2}^{-1} \tilde{\mathbf{T}}_{\ell_2} = \mathbf{T}_{\ell_2}^{-1} \mathbf{T}_{\ell_1} \Xi_{\ell_1 \ell_2}, \quad (35)$$

where \mathbf{T}_{ℓ_1} and \mathbf{T}_{ℓ_2} are the two poses involved in loop-closure measurement ℓ . Additionally, it was discovered in testing that including a relative pose constraint between each subsequent pair of poses helped the loop-closure correction to propagate throughout the trajectory. The relative pose errors take the same form as the loop-closure errors,

$$\mathbf{e}_k^{\text{rel.}} = \mathbf{T}_k^{-1} \tilde{\mathbf{T}}_k = \mathbf{T}_k^{-1} \mathbf{T}_{k-1} \Xi_{k-1,k}, \quad (36)$$

where the relative pose measurements are taken directly from the initializing solution,

$$\Xi_{k-1,k} = \tilde{\mathbf{T}}_{k-1}^{-1} \tilde{\mathbf{T}}_k. \quad (37)$$

Finally, since roll, pitch, and depth are directly observable AUV states [5], errors are included of the form

$$\mathbf{e}_k^{\text{obs.}} = \begin{bmatrix} e_k^{\phi_1} \\ e_k^{\phi_2} \\ e_k^z \end{bmatrix} = \mathbf{D} \mathbf{E}_k \log(\mathbf{T}_k^{-1} \tilde{\mathbf{T}}_k)^\vee, \quad (38)$$

where

$$\mathbf{D} = \begin{bmatrix} \mathbf{1} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{1} \end{bmatrix} \in \mathbb{R}^{3 \times 6}, \quad \mathbf{E}_k = \begin{bmatrix} \mathbf{1} \\ \mathbf{C}_{abk} \mathbf{J}^\ell(\mathbf{e}_k^\phi) \end{bmatrix}, \quad (39)$$

where \mathbf{J}^ℓ is the left Jacobian of $SO(3)$ [23, Sec. 7.1.3], and

$$\mathbf{e}_k^\xi = [(\mathbf{e}_k^\phi)^\top \quad (\mathbf{e}_k^\rho)^\top]^\top. \quad (40)$$

The least-squares objective function (23) is then augmented as

$$J_{\text{aug.}}(\mathbf{X}) = J(\mathbf{X}) + \frac{1}{2} \sum_{k=1}^K \left(\|\mathbf{e}_k^{\text{rel.}}\|_{\mathbf{R}_{\text{rel.}}}^2 + \|\mathbf{e}_k^{\text{obs.}}\|_{\mathbf{R}_{\text{obs.}}}^2 \right), \quad (41)$$




Fig. 6: The factor graph corresponding to the batch state estimation problem. The formation of WNOA factors \mathbf{e}_k and relative pose factors $\mathbf{e}_k^{\text{rel.}}$ allow corrections from loop-closure factor \mathbf{e}_ℓ to propagate throughout the graph.

where $\mathbf{R}_{\text{rel.}}$ and $\mathbf{R}_{\text{obs.}}$ are considered to be additional hyperparameters. Together, the relative pose errors (36) promote loop-closure propagation, while the WNOA errors (33) promote smoothing. The hyperparameters \mathbf{Q} and $\mathbf{R}_{\text{rel.}}$ may be tuned to control the smoothness of the posterior solution, while $\mathbf{R}_{\text{obs.}}$ is tuned to ensure the posterior does not stray too far in observable dimensions.

The batch estimation problem is visualized in the factor graph of Fig. 6. Note that, in contrast to the initial factor graph in Fig. 2, there are now factors linking adjacent nodes. This will allow corrections from the loop-closure measurements to propagate throughout the pose graph, as required.

2) Minimizing the Objective Function

To minimize (41), the estimation errors are repeatedly linearized about the current navigation state estimate $\bar{\mathbf{X}}$. Perturbing the navigation state as

$$\mathbf{T} = \bar{\mathbf{T}} \exp(-\delta \xi^\wedge), \quad (42a)$$

$$\varpi = \bar{\varpi} + \delta \varpi, \quad (42b)$$

$$\delta \mathbf{x} = [\delta \xi^\top \quad \delta \varpi^\top]^\top, \quad (42c)$$

the prior error (32) is linearized as

$$\mathbf{e}_0 = \bar{\mathbf{e}}_0 + \mathbf{F}_0^0 \delta \mathbf{x}_0 + \mathbf{M}_0 \delta \mathbf{y}_0, \quad (43)$$

where $\delta \mathbf{y}_0 = [\delta \eta_0^\top \quad \delta \psi_0^\top]^\top$, $\mathbf{S}_0 = E[\delta \mathbf{y}_0 \delta \mathbf{y}_0^\top]$, and where the prior Jacobians are

$$\mathbf{F}_0^0 = \text{blkdiag}(\mathbf{J}^\ell(\bar{\mathbf{e}}_0^\xi)^{-1}, \mathbf{1}), \quad (44a)$$

$$\mathbf{M}_0 = \text{blkdiag}(-\mathbf{J}^\top(\bar{\mathbf{e}}_0^\xi)^{-1}, -\mathbf{1}), \quad (44b)$$

with \mathbf{J}^ℓ being the left Jacobian of $SE(3)$. Note that detailed derivations of the work appearing in this section are available in the supplementary material in [52]. The discrete-time process errors (33) are linearized as

$$\mathbf{e}_k = \bar{\mathbf{e}}_k + \mathbf{F}_{k-1}^k \delta \mathbf{x}_{k-1} + \mathbf{F}_k^k \delta \mathbf{x}_k, \quad (45)$$

where the process error Jacobians are given by

$$\mathbf{F}_{k-1}^k = \begin{bmatrix} \mathbf{U}_{k-1} & \mathbf{V}_{k-1} \\ \mathbf{0} & -\mathbf{1} \end{bmatrix}, \quad (46a)$$

$$\mathbf{U}_{k-1} = -\mathbf{J}^\top(\bar{\mathbf{e}}_k^\xi)^{-1} \mathbf{Ad}(\exp(-T \varpi_{k-1}^\wedge)),$$

$$\mathbf{V}_{k-1} = T \mathbf{J}^\top(\bar{\mathbf{e}}_k^\xi)^{-1} \mathbf{J}^\top(T \varpi_{k-1}^\wedge),$$

$$\mathbf{F}_k^k = \begin{bmatrix} \mathbf{J}^\ell(\bar{\mathbf{e}}_k^\xi)^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} \end{bmatrix}. \quad (46b)$$

The loop-closure errors (35) are linearized as

$$\mathbf{e}_\ell = \bar{\mathbf{e}}_\ell + \mathbf{H}_{\ell_1}^\ell \delta \xi_{\ell_1} + \mathbf{H}_{\ell_2}^\ell \delta \xi_{\ell_2} + \mathbf{M}_\ell \delta \xi_\Xi, \quad (47)$$

with corresponding Jacobians

$$\mathbf{H}_{\ell_1}^\ell = -\mathbf{J}^r(\bar{\mathbf{e}}_\ell)^{-1} \mathbf{Ad}(\bar{\Xi}_{\ell_1 \ell_2}^{-1}), \quad (48a)$$

$$\mathbf{H}_{\ell_2}^\ell = \mathbf{J}^\ell(\bar{\mathbf{e}}_\ell)^{-1}, \quad (48b)$$

$$\mathbf{M}_\ell = -\mathbf{J}^r(\bar{\mathbf{e}}_\ell)^{-1}. \quad (48c)$$

The relative pose errors (36) are linearized in the same manner. Finally, errors on the observable states (38) are linearized by approximating

$$\mathbf{C}_{ab_k} = \bar{\mathbf{C}}_{ab_k} \exp(-\delta \phi_k^\times) \approx \bar{\mathbf{C}}_{ab_k} (\mathbf{1} - \delta \phi_k^\times). \quad (49)$$

Assuming $\mathbf{e}_k^{\rho} \rightarrow \mathbf{0}$ as the optimization proceeds, this yields

$$\mathbf{e}_k^{\text{obs.}} = \bar{\mathbf{e}}_k^{\text{obs.}} + \mathbf{H}_k^{\text{obs.}} \delta \xi_k, \quad (50)$$

$$\mathbf{H}_k^{\text{obs.}} = \mathbf{DE}_k(\bar{\mathbf{C}}_{ab_k}, \bar{\mathbf{e}}_k^\phi) \mathbf{J}^\ell(\bar{\mathbf{e}}_k^\xi)^{-1}. \quad (51)$$

The final step is to determine the covariance on the discrete-time WNOA process errors. This is done by discretizing the power spectral density \mathcal{Q} via [53, (4.110)]

$$\mathbf{Q}_k = \int_{t_{k-1}}^{t_k} \mathbf{A}(t_k, s) \mathbf{L}(s) \mathcal{Q}(s) (\mathbf{A}(t_k, s) \mathbf{L}(s))^T ds, \quad (52)$$

where \mathbf{A}, \mathbf{L} characterize the continuous-time error kinematics, which for the WNOA motion prior take the form

$$\delta \dot{\mathbf{x}}(t) = \underbrace{\begin{bmatrix} -\mathbf{ad}(\bar{\boldsymbol{\omega}}^\wedge) & -\mathbf{1} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}}_{\mathbf{A}} \delta \mathbf{x}(t) + \underbrace{\begin{bmatrix} \mathbf{0} \\ \mathbf{1} \end{bmatrix}}_{\mathbf{L}} \delta \mathbf{w}(t), \quad (53)$$

with $\delta \mathbf{w}(t) \sim \mathcal{GP}(\mathbf{0}, \mathcal{Q}(t-t'))$. The exact solution to (52) may be obtained via the matrix exponential [54], however to avoid this expensive operation this work makes use of a third-order approximation in \mathbf{A} [53, (4.119)],

$$\begin{aligned} \mathbf{Q}_k &\approx T \boldsymbol{\Upsilon} + \frac{T^2}{2} (\mathbf{A} \boldsymbol{\Upsilon} + \boldsymbol{\Upsilon} \mathbf{A}^T) \\ &+ \frac{T^3}{6} \left(\mathbf{A}^2 \boldsymbol{\Upsilon} + 2 \mathbf{A} \boldsymbol{\Upsilon} \mathbf{A}^T + \boldsymbol{\Upsilon} (\mathbf{A}^T)^2 \right) \\ &+ \frac{T^4}{24} \left(\mathbf{A}^3 \boldsymbol{\Upsilon} + 3 \mathbf{A}^2 \boldsymbol{\Upsilon} \mathbf{A}^T + 3 \mathbf{A} \boldsymbol{\Upsilon} (\mathbf{A}^T)^2 + \boldsymbol{\Upsilon} (\mathbf{A}^T)^3 \right), \end{aligned} \quad (54)$$

where $\boldsymbol{\Upsilon} = \mathbf{L} \mathcal{Q} \mathbf{L}^T$. Finally, the minimizing solution for a single iteration of Gauss-Newton is

$$\delta \mathbf{x}_* = \begin{bmatrix} \delta \xi_* \\ \delta \boldsymbol{\omega}_* \end{bmatrix} = -(\boldsymbol{\Gamma}^T \mathbf{W} \boldsymbol{\Gamma})^{-1} \boldsymbol{\Gamma}^T \mathbf{W} \mathbf{e}. \quad (55)$$

Jacobian $\boldsymbol{\Gamma}$ is given by

$$\boldsymbol{\Gamma} = \begin{bmatrix} \mathbf{F}^T & \mathbf{H}^T & \mathbf{H}^{\text{rel.}T} & \mathbf{H}^{\text{obs.}T} \end{bmatrix}^T, \quad (56)$$

$$\mathbf{F} = \begin{bmatrix} \mathbf{F}_0^0 & & & & & \\ \mathbf{F}_0^1 & \mathbf{F}_1^1 & & & & \\ & & \ddots & & & \\ & & & \ddots & & \\ & & & & \mathbf{F}_{K-1}^K & \mathbf{F}_K^K \end{bmatrix}, \quad (57)$$

$$\mathbf{H} = \begin{bmatrix} \mathbf{H}_{\ell_1}^1 & & & & & \\ & & & & \mathbf{H}_{\ell_2}^1 & \\ & & & & & \vdots \\ & & \mathbf{H}_{\ell_1}^L & & & \\ & & & & & \mathbf{H}_{\ell_2}^L \end{bmatrix}, \quad (58)$$

$$\mathbf{H}^{\text{rel.}} = \begin{bmatrix} \mathbf{H}_0^{\text{rel.,1}} & \mathbf{H}_1^{\text{rel.,1}} & & & & \\ & & \ddots & & & \\ & & & \ddots & & \\ & & & & \mathbf{H}_{K-1}^{\text{rel.,K}} & \mathbf{H}_K^{\text{rel.,K}} \end{bmatrix}, \quad (59)$$

$$\mathbf{H}^{\text{obs.}} = \begin{bmatrix} \mathbf{0} & \mathbf{H}_1^{\text{obs.}} & & & & \\ & & \ddots & & & \\ & & & & & \\ & & & & & \mathbf{H}_K^{\text{obs.}} \end{bmatrix}, \quad (60)$$

and weighting matrix $\mathbf{W} = \boldsymbol{\Sigma}^{-1}$ is described by

$$\boldsymbol{\Sigma} = \text{blkdiag}(\mathbf{M}_0 \mathbf{S}_0 \mathbf{M}_0^T, \mathbf{Q}_{1:K}, \mathbf{R}_{1:L}, \mathbf{R}_{1:K}^{\text{rel.}}, \mathbf{R}_{1:K}^{\text{obs.}}), \quad (61)$$

where, for the loop-closure errors,

$$\mathbf{R}_\ell = \mathbf{M}_\ell \mathbf{R}_\Xi \mathbf{M}_\ell^T. \quad (62)$$

The column matrix of errors is simply

$$\mathbf{e} = \begin{bmatrix} \mathbf{e}_0^T & \mathbf{e}_{1:K}^T & \mathbf{e}_{1:L}^T & (\mathbf{e}_{1:K}^{\text{rel.}})^T & (\mathbf{e}_{1:K}^{\text{obs.}})^T \end{bmatrix}^T. \quad (63)$$

Finally, in accordance with the perturbation scheme (42a, 42b), the state update is given by

$$\mathbf{T} \leftarrow \mathbf{T} \exp(-\delta \xi_*^\wedge), \quad (64a)$$

$$\boldsymbol{\omega} \leftarrow \boldsymbol{\omega} + \delta \boldsymbol{\omega}_*. \quad (64b)$$

3) Rejecting False Loop-Closure Measurements

Measurement outliers are inevitable in real-world robotics problems, and a robust implementation of the proposed methodology requires a method to identify and reject false loop-closure measurements. Many approaches exist in the literature for rejecting loop-closure measurement outliers, for example switchable constraints [55], expectation-maximization [56], and graduated non-convexity [57].

This application uses a recently developed adaptive robust cost function (RCF) to reject false loop-closure measurements, owing to its ability to handle multivariate, mixed-unit error definitions, such as loop-closure errors (35), in a statistically sound manner [58]. The RCF assigns a weight $w_\ell(\epsilon_\ell(\mathbf{e}_\ell)) \in (0, 1]$ to loop-closure error \mathbf{e}_ℓ according to the Mahalanobis distance associated with the error,

$$\epsilon_\ell(\mathbf{e}_\ell) = \|\mathbf{e}_\ell\|_{\boldsymbol{\Sigma}_\ell^{-1}} \in \mathbb{R}_{\geq 0}, \quad (65)$$

where, ordinarily, the covariance $\boldsymbol{\Sigma}_\ell$ on the (relative) loop-closure measurement would be the relative uncertainty computed between the two vehicle poses involved in the measurement [59]. Since the DVL-INS output does not contain the joint covariance information required to properly compute $\boldsymbol{\Sigma}_\ell$, a constant value is used here,

$$\boldsymbol{\Sigma}_\ell = \text{blkdiag}(\sigma_{\phi_{\text{out}}}^2 \mathbf{1}, \sigma_{\rho_{\text{out}}}^2 \mathbf{1}), \quad (66)$$

with $\sigma_{\phi_{\text{out}}} = 1$ deg and $\sigma_{\rho_{\text{out}}} = 1$ m. These 1σ values reflect the low heading uncertainty of the survey-grade DVL-INS used in the field experiments, as well as the static search bound used for the loop-closure detection method (27).

IV. RESULTS

A. Assessing the Quality of the State Estimate

The methodology described in Section III conditions an existing state estimate on newly available loop-closure measurements. Since loop closures provide relative constraints between poses, as shown in Fig. 6, it is expected that this approach will

- 1) reduce *relative pose errors* throughout the trajectory; and
- 2) produce a more self-consistent point cloud map, as measured by a reduction in the *point disparity error* in overlapping regions.

1) Measuring Errors in the Estimated Trajectory

A pose-based relative error metric based on [60] is used to measure the accuracy of the estimated trajectory. Let $\hat{\mathbf{T}}_k \in SE(3)$ represent the estimated pose at time t_k . The pose at time t_k relative to the pose at time t_ℓ is

$$\delta\hat{\mathbf{T}}_{\ell k} = \hat{\mathbf{T}}_\ell^{-1}\hat{\mathbf{T}}_k, \quad (67)$$

where $\hat{\mathbf{T}}_\ell$ is taken to be the earliest pose involved in any loop-closure measurement. The relative pose error may then be expressed as

$$\mathbf{E}_k^{\text{rel.}} = \delta\mathbf{T}_{\ell k}^{-1}\delta\hat{\mathbf{T}}_{\ell k} = \begin{bmatrix} \delta\mathbf{C}_k & \delta\mathbf{r}_k \\ \mathbf{0} & \mathbf{1} \end{bmatrix}, \quad (68)$$

where $\delta\mathbf{T}_{\ell k}$ is (67) evaluated using the ground-truth trajectory. Relative pose errors on $SE(3)$ may then be expressed in the Lie algebra as

$$\delta\mathbf{e}_k^{\text{rel.}} = \log\left(\mathbf{E}_k^{\text{rel.}}\right)^\vee = [(\delta\phi_k^{\text{rel.}})^\top (\delta\rho_k^{\text{rel.}})^\top]^\top. \quad (69)$$

However, since the AUV trajectories studied in this work are largely planar, the accuracy of the trajectory estimate is reported as the norm of the relative displacement error $\delta\mathbf{r}_k$ projected on the (x, y) plane,

$$e_k^{\text{rel.}} = \|\begin{bmatrix} \mathbf{1} & \mathbf{0} \end{bmatrix} \delta\mathbf{r}_k\|_2. \quad (70)$$

Assessing performance using a relative metric such as (68) avoids the problem of aligning the estimated and ground-truth trajectories, which would be required if attempting to provide an absolute performance metric [60]. In addition to being intuitive and easy to visualize, the relative planar displacement metric (70) allows for direct comparison to other navigation solutions within the subsea industry, where position drift is often reported on the (x, y) plane as a percentage of distance traveled. For estimates incorporating multiple loop closures along the length of the trajectory, the relative displacement error (70) is expected to remain bounded over time.

2) Measuring Self-Consistency in the Point Cloud Map

Performance is also assessed by evaluating the self-consistency in overlapping regions of the point cloud map. A point cloud map generated from an accurate trajectory estimate is expected to be well-aligned, or “crisp.” In contrast, a map produced using a drifting trajectory estimate will see “double vision” effects in overlapping regions due to poorly aligned scans. To assess self-consistency in the point cloud map, this work uses the *point disparity* metric from [61]. For point clouds $\mathcal{S} = \{\mathbf{r}_a^{p_i w}\}_{i=1}^N$ and $\mathcal{T} = \{\mathbf{r}_a^{p_j w}\}_{j=1}^M$, this metric is

$$e_j^{\text{rel.}} = \|\mathbf{r}_a^{p_j w} - \mathbf{r}_a^{p_i w}\|_2, \quad (71)$$

where point p_j in \mathcal{T} is the nearest Euclidean neighbour to point p_i in \mathcal{S} . Note (71) is only computed within the intersection of \mathcal{S} and \mathcal{T} to prevent cloud size from biasing the metric. The point disparity metric is relative, and may be computed without access to ground-truth information [61], making it especially important for field trials where a ground-truth map is not available. Note the point disparity metric is susceptible to map-to-map error, whereby an erroneous group of two or more well-aligned submaps would produce low disparity errors, despite separation from the true submap group. The results in the following studies were visually checked to ensure the absence of this error, though extending (71) to account for map-to-map error is an interesting avenue for future research.

B. Hyperparameter Values

The methodology described in Section III involves three sets of hyperparameters. These are

- 1) \mathcal{Q} , the PSD on the white noise Gaussian process;
- 2) $\mathbf{R}_k^{\text{rel.}}$, the covariance on the relative pose errors; and
- 3) $\mathbf{R}_k^{\text{obs.}}$, the covariance on the roll, pitch, and depth errors, all of which are assumed to be observable.

The hyperparameter sets take the form

$$\mathcal{Q} = \text{diag}(\mathcal{Q}_\omega \mathbf{1}, \mathcal{Q}_v \mathbf{1}), \quad (72a)$$

$$\mathbf{R}_k^{\text{rel.}} = \text{diag}(\sigma_\phi^2 \mathbf{1}, \sigma_\rho^2 \mathbf{1}), \quad (72b)$$

$$\mathbf{R}_k^{\text{obs.}} = \text{diag}(\sigma_{\text{rp}}^2 \mathbf{1}, \sigma_z^2), \quad (72c)$$

where \mathcal{Q}_ω , \mathcal{Q}_v are the power spectral densities on the body-centric angular and linear acceleration, respectively, and σ_ϕ^2 , σ_ρ^2 are the variances on the body-centric angular and linear displacement, respectively. σ_{rp}^2 is the variance on vehicle roll and pitch, and σ_z^2 is the variance on vehicle depth. While this hyperparameter structure is simple, it was found to work well for both simulated and field experiments, and generally makes physical sense. For example, [62] scales the values of a diagonal \mathcal{Q} matrix to reflect the nonholonomic constraints of an automobile. In contrast, AUVs are highly maneuverable, leading to the selection of isotropic hyperparameters in (72).

This section contains results from both simulated and field experiments. The hyperparameter values used to obtain each set of results are summarized in Table III. These values were hand-selected to produce good results without extensive tuning, and were modified based on the quality of the DVL-INS state estimate and the frequency at which DVL-INS data were available.

TABLE III: Hyperparameter values used in experiments

Hyperparameter	Unit	Simulated	Field
\mathcal{Q}_ω	$\text{rad}^2 \text{s}^{-3}$	1e-2	1e-2
\mathcal{Q}_v	$\text{m}^2 \text{s}^{-3}$	9e-4	1e-4
σ_ϕ	rad	1e-3	1e-6
σ_ρ	m	1e-4	1e-6
σ_{rp}	deg	5	5
σ_z	m	0.25	0.25

C. Simulation Results: AUV Area Inspection

Simulated output from a DVL-INS and a corresponding ground-truth trajectory were provided by industry collaborator Sonardyne. The DVL-INS output contains latitude and longitude, depth, and roll-pitch-yaw estimates for a simulated AUV deployment. At each time step, marginal variance estimates are available for the depth and heading states, and

a joint covariance estimate is available for planar position in the local geodetic frame. 30 min of DVL-INS output data is available at a frequency of 5 Hz. Note that the DVL-INS output has been heavily degraded by Sonardyne to better assess the ability of loop closures to mitigate navigation drift and does not reflect the performance of Sonardyne commercial products.

The ground-truth trajectory is shown in Fig. 7, where the vehicle starts with a single tie-line followed by multiple planar “lawnmower” passes over a large inspection area. Loop closures occur at the eight intersections between the lawnmower passes and the tie-line. The prior “INS” estimate is then conditioned on the loop-closure measurements using the methodology from Section III to produce a posterior “INS+LC” trajectory estimate. Both estimates are then compared to the ground-truth solution (“GT”) using the metrics from Section IV-A. The application and propagation of loop-closure measurements within the WNOA framework is expected to produce a more accurate navigation solution with a correspondingly more self-consistent point cloud map.




Fig. 7: The simulated dataset. A single tie-line is intersected on the left by multiple “lawnmower” inspection passes.

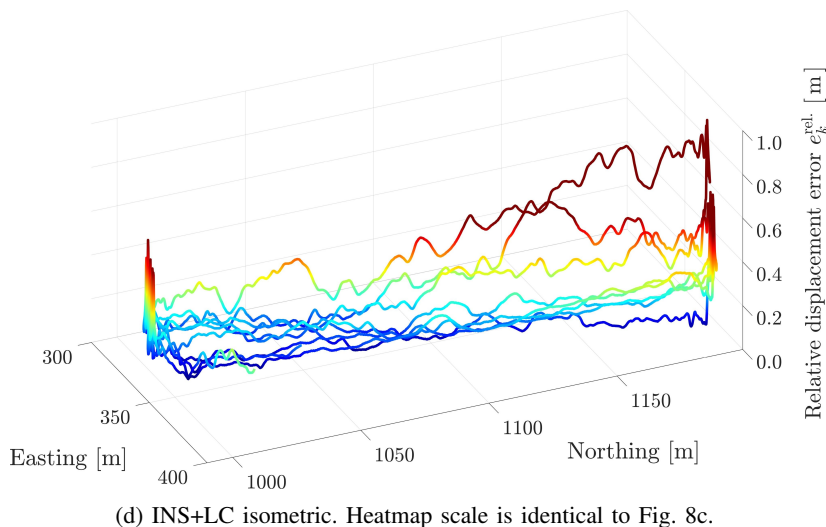





Fig. 8: Heatmaps representing the relative displacement error for the INS and INS+LC trajectories. The left column shows a top view, while the right column shows an isometric view to illustrate the smoothing effects of the WNOA terms. Note the series of spikes at the far ends of Fig. 8d, where the vehicle completes low-radius turns to initiate the next survey pass.

To illustrate the improvements in accuracy and the smoothing effect from the WNOA terms, Fig. 8 displays the relative displacement error e_k^{rel} as a “heatmap” for the prior INS and posterior INS+LC trajectories. From Fig. 8a it is clear that the prior estimate is not accurate, with relative displacement errors exceeding 0.5 m for much of the trajectory. Incorporating loop-closure measurements improves the accuracy of the trajectory estimate, as demonstrated by

the cooler colours throughout the heatmap in Fig. 8c. The improvement is particularly noticeable around the tie-line at the bottom of Fig. 8c, however for many of the passes the effects extend *hundreds of meters beyond the loop-closure location*, increasing the accuracy across the entire inspection area. Relative displacement errors of 0.5 m may seem inconsequential at this scale, but may prove critical for certain subsea activities such as jumper pipe installation.



(a) An underwater scene generated in the Stonefish AUV simulator [63]. From left to right: boat hull, dragon, propeller, and armadillo. The hull and propeller are available in Stonefish, while the dragon and armadillo are from the Stanford 3D Scanning Repository [64].

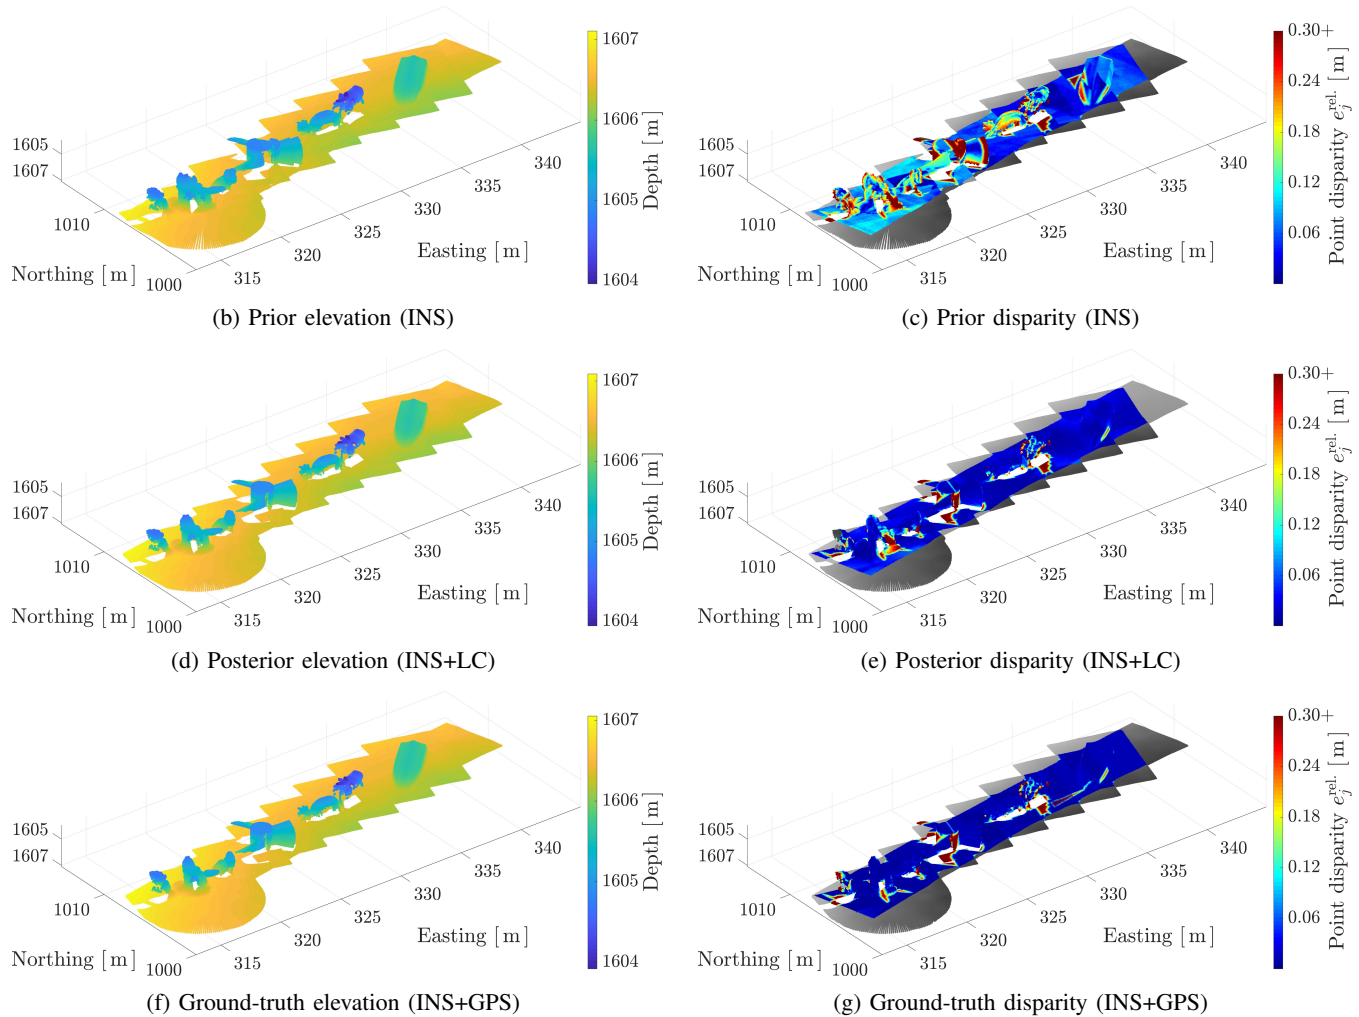




Fig. 9: Scanning a 3D scene in the Stonefish AUV simulator [63] to evaluate the quality of the point cloud map. The four models in Fig. 9a are positioned along the tie-line near the loop-closure locations, and are scanned multiple times as the AUV completes the trajectory. Errors in the trajectory estimate are easily seen in the point disparity maps in the right column. Note the areas of high point disparity in the ground-truth map Fig. 9g are due to occlusion of the range-bearing scanner. Scanning the 3D models at different orientations produces different occlusion patterns, leading to non-overlapping areas of the map. This in turn leads to large nearest-neighbour distances to points in other scans, and thus a high point disparity error.



(a) Empirical probability density functions (EPDF)



(b) Empirical cumulative distribution functions (ECDF)

Fig. 10: Distributions on the point disparity error for each of the three navigation solutions for the simulated AUV area inspection dataset. The INS+LC solution greatly improves on the prior INS solution, virtually eliminating errors beyond 6 cm. The relatively uniform error distribution for the INS solution is visible in the heatmap of Fig. 9c.

The smoothing effects from the WNOA motion prior are visible in Fig. 8b and 8d, which are, respectively, isometric views of Fig. 8a and 8c. Here, the relative displacement error is represented both by the heatmap and the plot elevation. The prior INS estimate shows visible step changes in the relative displacement error, which are characteristic of the correction step of a filter and likely represent the effects of DVL measurements within the DVL-INS estimation algorithm. In contrast, the posterior INS+LC estimate has been visibly smoothed due to the presence of the WNOA error terms. Trajectory smoothing is important in this context, as the point cloud map is generated by registering individual laser profiles to the trajectory estimate. A trajectory with step changes will

produce a map with step changes, which will surely impact front-end activities such as feature detection and point cloud alignment. A smooth, self-consistent map is also visually appealing, and will improve the accuracy of subsea metrology.

Solution	50%	1σ	2σ	3σ
INS	6.23 cm	8.32 cm	23.39 cm	56.15 cm
INS+LC	1.61 cm	2.27 cm	11.60 cm	44.70 cm
INS+GPS	0.81 cm	1.13 cm	11.48 cm	45.64 cm

Note the series of “spikes” in the relative displacement error of Fig. 8d, corresponding to the posterior trajectory estimate. These spikes occur at the beginning and end of each low-radius turn, suggesting that the smoothing effect of the WNOA terms may be erasing trajectory information in high-curvature regions. A geometry-based trajectory upsampling approach, for example based on scale-invariant density [65], is expected to resolve this, and will be explored as part of future work.

To evaluate the effects of the optimization on map quality, an underwater scene was constructed and scanned using the open-source AUV simulator Stonefish [63]. Raw laser profiles were collected along the ground-truth trajectory, and were then registered to the INS and INS+LC trajectories to produce, respectively, the prior and posterior point cloud maps. The 3D scene in Stonefish, as well as the resulting elevation and point disparity maps, are shown throughout Fig. 9.

Compared to the ground-truth disparity map Fig. 9g, the prior map Fig. 9c shows high point disparity errors throughout, indicating a self-inconsistent map. In contrast, the disparity errors are largely resolved by the INS+LC solution, which incorporates both loop-closure measurements and smoothing into the INS estimate. The improvements are quantified in Fig. 10 by plotting an empirical probability density function (EPDF) and an empirical cumulative distribution function (ECDF) of the disparity error for each of the three solutions. For a quantitative comparison, Table IV lists critical values drawn from the ECDF curves. The INS+LC solution improves on the INS solution by producing a larger fraction of points with a lower point disparity error. This is especially evident in Fig. 10b, where the INS+LC solution converges to the GT solution around 6 cm. Assuming the remaining 5% of errors lie in occluded regions, as explained in the caption of Fig. 9, the INS+LC solution has effectively eliminated point disparity errors beyond 6 cm. In contrast, 20% of disparity errors from the prior INS solution exceed 10 cm. “Double-vision” effects of this magnitude arising from poor scan alignment are sure to complicate inspection and metrology tasks, even within the small domain of this simulation. Following the methodology from Section III, the INS+LC solution has produced a smooth, crisp, self-consistent point cloud map from which relative distance measurements may accurately be drawn.

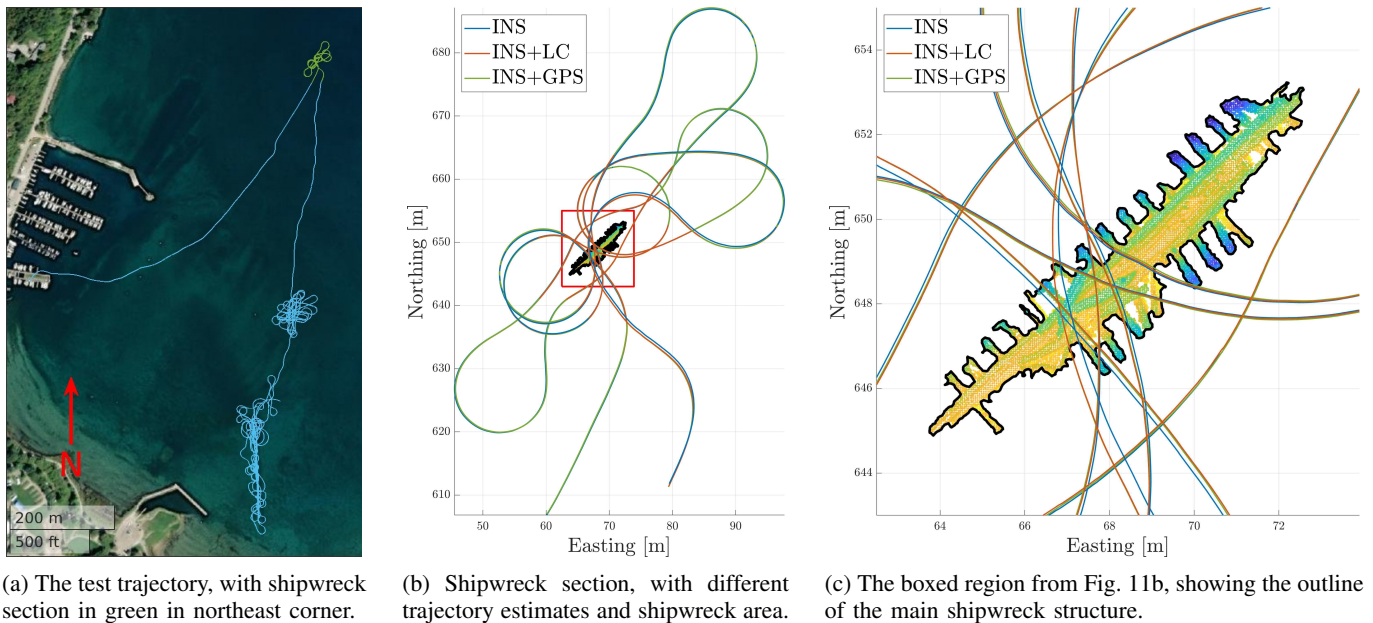


Fig. 11: Field deployment in Colpoy’s Bay, Warton, Ontario, Canada. The full trajectory is shown in Fig. 11a. The shipwreck section, shown in green in northeast corner of Fig. 11a, is 0.58 km long and took approximately 10.5 min to complete. The trajectory makes eight passes over the shipwreck area. The different navigation solutions are summarized in Table V.

D. Field Results: Warton Shipwreck

A field trial was conducted with Voyis Imaging Inc. in Colpoy’s Bay, Warton, Ontario, Canada. The bay is shallow and contains multiple shipwrecks and other manmade structures, making it an ideal test location for Voyis’s surface vessel. The full test trajectory is shown in blue in Fig. 11a. A section of the trajectory, highlighted in green in the northeast corner of Fig. 11a, makes eight passes over a small shipwreck. Fig. 11b shows this section in detail, and Fig. 11c shows the main shipwreck structure segmented from the lakebed. This section of the trajectory, which is approximately 580 m long and took 10.5 min to complete, is the focus of the field results.

The surface vessel was equipped with a Sonardyne SPRINT-Nav 500 DVL-aided INS, a u-blox ZED-F9P high precision GNSS module equipped with a u-blox ANN-MB series high precision multi-band antenna, and a Voyis Insight Pro underwater laser scanner. GNSS data were post-processed using the Canadian Spatial Reference System Precise Point Positioning (CSRS-PPP) application [66], which in a recent study was found to be capable of measuring 2D position with a precision of 2 cm (1σ) [67].

Three navigation solutions were generated from these data. The first solution is a dead-reckoned DVL-INS trajectory (“INS”), where the positioning precision of the SPRINT-Nav 500 has been manually degraded by Sonardyne from the nominal value of 0.02% of distance traveled (CEP50) [53, Sec. 4.9.1.2][13]. The DVL-INS output is available at 10 Hz. Note that this solution is representative of the state-of-the-art for high-grade commercial systems, and will be used to benchmark the proposed methodology. The second solution, referred to as “INS+LC,” applies the methodology of Section III to incorporate loop-closure measurements into

TABLE V: Understanding the different navigation solutions for the Warton shipwreck field dataset.

Solution	Description
INS	Dead-reckoned DVL-INS solution, with position precision manually degraded by Sonardyne.
INS+LC	The dead-reckoned DVL-INS trajectory estimate conditioned on loop-closure measurements. Raw sensor measurements from the DVL-INS are inaccessible , and GNSS data is not used as part of this solution.
INS+GPS	DVL-aided INS solution with GNSS correction.

the dead-reckoned DVL-INS estimate. Batch processing was performed offline in MATLAB, taking approximately 90 s to converge on a laptop with an E3-1505M v5 CPU and 16 GB of RAM. It is important to note that this solution is produced using the DVL-INS state *estimate*, *without access to the raw DVL-INS sensor measurements*. The third solution fuses the DVL-INS output with the GNSS data to form a ground-truth estimate (“INS+GPS”). The three navigation solutions are summarized in Table V, and the trajectory estimates are overlaid on the shipwreck area in Fig. 11b and 11c.

Incorporating loop-closure measurements produces a more accurate trajectory estimate, as measured by the relative displacement error (70). Fig. 12 shows the relative displacement error, measured against the ground-truth INS+GPS estimate, for the dead-reckoned INS trajectory and the INS+LC trajectory with an increasing number of loop




Fig. 12: Relative displacement errors for trajectory estimates incorporating an increasing number of loop closures. Loop-closure locations are marked as vertical dashed lines, and “INS+XLC” indicates the first X loop closures were used in generating the estimate. Incorporating loop closures bounds navigation drift over time. For numerical results, see Table VI.

closures. The loop-closure locations are marked with vertical dashed lines, with the first observation of the shipwreck occurring approximately 40s into the trajectory. The relative displacement drift in the INS trajectory estimate increases without bound, while the maximum displacement error decreases monotonically as more loop-closure measurements are applied. Even a single loop-closure measurement at the end of the trajectory is effective in bounding the relative displacement drift over time, as demonstrated by the cyan line in Fig. 12. From Table VI, which summarizes the maximum error and final error as a percent of distance traveled for the different solutions, the final drift error for the “INS+1LC (last)” solution is $6.82e-3\%$ of distance traveled. This particular solution suggests an order of magnitude improvement over state-of-the-art DVL-INS systems [13]. Importantly, the dashed yellow “INS+0LC” curve in Fig. 12 indicates that the posterior solution does not deviate far from the prior DVL-INS solution when loop-closure measurements are absent. This suggests that the proposed methodology may still be used to smooth the DVL-INS solution in the absence of loop-closure measurements, without sacrificing solution accuracy. For example, the maximum position drift error for the “INS-0LC” solution tabulated in Table VI is only 9 mm

TABLE VI: Summary of drift errors from Fig. 12, with values drawn *after* the first shipwreck observation at 40s.

Solution	Max drift [m]	Final %DT
INS	0.658	$10.98e-2$
INS+0LC	0.667	$10.82e-2$
INS+1LC	0.378	$5.99e-2$
INS+3LC	0.255	$4.12e-2$
INS+5LC	0.150	$2.25e-2$
INS+7LC	0.084	$6.83e-3$
INS+1LC (last)	0.224	$6.82e-3$

greater than the maximum drift observed in the “INS” solution. However, note the proposed methodology is intended to be used in a targeted fashion in situations where at least one loop-closure measurement is available.

In addition to the relative displacement errors summarized in Fig. 12, relative pose errors (69) are computed across the trajectory for the prior “INS” and posterior “INS+LC” solutions, with summary statistics given in Table VII. Relative attitude errors in Table VII are decomposed into body-centric roll, pitch, and yaw, while the rightmost column gives statistics on the Euclidean norm of the body-centric relative position errors. Interestingly, the proposed methodology has produced an increase in the relative body-centric pitch and roll errors, from median values of $1.1e-3$ deg and $1.5e-5$ deg, respectively, to $1.1e-1$ deg and $1.2e-1$ deg, respectively. Relative body-centric yaw errors remain largely unchanged by the proposed methodology, while trends in the relative body-centric position error generally follow the trend of the relative displacement error plotted in Fig. 12. For example, 90% of body-centric position errors for the prior “INS” solution fall below 0.630 m, while the corresponding value for the posterior “INS+LC” solution is 0.108 m.

An increase in roll and pitch errors may seem concerning, however the posterior errors remain low and bounded. Such errors were likely introduced in this field trial through a combination of small angular errors in the INS-laser extrinsics estimate and the relatively weak pitch and roll prior used in the optimization (see (38) and the value of hyperparameter σ_{rp} in Table III). The more important result is that relative body-centric position errors remain low and bounded when multiple loop-closure measurements are present.

Barring measurement outliers, finding that loop closures improve trajectory accuracy is not particularly surprising in a conventional state estimation context. However, these results have been achieved following the methodology of Section III,

TABLE VII: Summary statistics on relative pose errors (69) computed for the Warton field trial. Relative attitude errors $|\delta\phi_i^{\text{rel}}|$ have been broken down by component, while the rightmost column gives statistics on the norm of the relative body-centric position error. Cumulative statistics for each error are reported in the format 50% · 75% · 90%.

Solution	$ \delta\phi_1^{\text{rel}} $ [deg]	$ \delta\phi_2^{\text{rel}} $ [deg]	$ \delta\phi_3^{\text{rel}} $ [deg]	$\ \delta\rho^{\text{rel}}\ $ [m]
INS	$1.1e-3 \cdot 1.9e-3 \cdot 2.5e-3$	$1.5e-3 \cdot 2.4e-3 \cdot 3.5e-3$	$1.5e-2 \cdot 1.9e-2 \cdot 2.2e-2$	$0.482 \cdot 0.562 \cdot 0.630$
INS+LC	$1.1e-1 \cdot 2.7e-1 \cdot 5.3e-1$	$1.2e-1 \cdot 3.8e-1 \cdot 5.7e-1$	$1.5e-2 \cdot 2.0e-2 \cdot 2.2e-2$	$0.077 \cdot 0.096 \cdot 0.108$

without access to raw sensor measurements, a vehicle process model, exteroceptive sensor models, or sensor noise and bias specifications. The loop-closure corrections have instead been smoothly integrated into the DVL-INS estimate using the factor graph illustrated in Fig. 6, improving the accuracy of the trajectory estimate.

Incorporating loop-closure measurements produces a more self-consistent point cloud map, as measured by the point disparity error (71). Fig. 13 shows the point disparity in the shipwreck area as a heatmap, for each of the three navigation solutions. The disparity is computed for each of the eight passes over the wreck, and is the Euclidean distance from each point in one pass to its nearest neighbour in the remaining seven passes. A highly accurate trajectory estimate is expected to produce a tightly overlapping, crisp point cloud map from

the composite scans, with a low point disparity error.

From a qualitative evaluation of Fig. 13b, the dead-reckoned INS trajectory estimate has clearly produced a self-inconsistent point cloud map. Areas around the ribs of the shipwreck have point disparity errors of around 20 cm, while one of the passes shows relatively large errors on the seabed owing to drift in the depth dimension. In contrast, both the posterior and the ground-truth estimates have produced highly self-consistent maps, with low point disparity errors throughout.

Interestingly, the INS+LC solution produces a point cloud map that is *more self-consistent than the ground-truth estimate*. This is difficult to judge qualitatively from Fig. 13, however Fig. 14 shows the EPDF and ECDF of the disparity error for each of the three navigation solutions. Critical values from the ECDF are tabulated in Table VIII. In Fig. 14a,

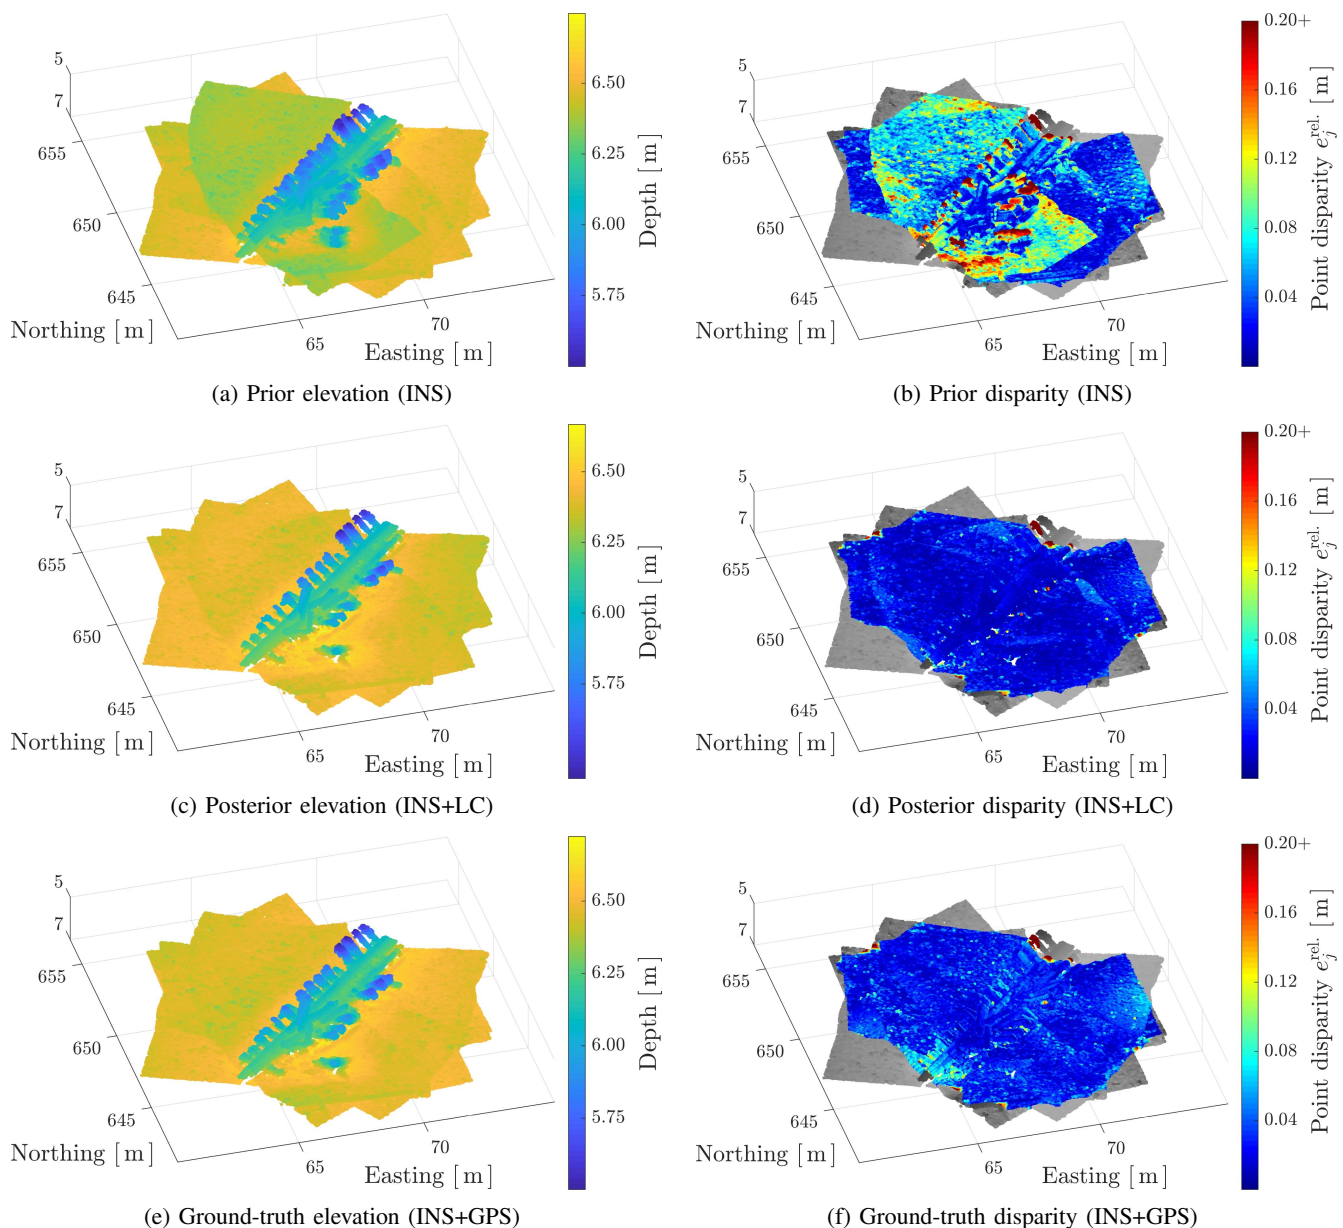




Fig. 13: Visualizing the point disparity error in the shipwreck area for the three navigation solutions. Left column: colour map indicates depth, and has been included for context. Right column: colour map indicates point disparity.



(a) Empirical probability density functions (EPDF)



(b) Empirical cumulative distribution functions (ECDF)

Fig. 14: Distributions on the point disparity error for each of the three navigation solutions for the Wiarnton shipwreck field dataset. The posterior INS+LC solution produces a more self-consistent point cloud map than the ground-truth INS+GPS solution, likely owing to a combination of residual navigation and extrinsics errors. Critical values from the ECDF are summarized in Table VIII.

the INS+LC curve peaks to the left of the INS+GPS curve, indicating a lower overall point disparity error and thus a more self-consistent point cloud map [5]. This is likely due to a combination of small estimation errors in the ground-truth solution and small errors in the scanner extrinsics estimate \mathbf{T}_{bl}^{sz} from (24). It should therefore come as no surprise that the INS+LC solution delivers a more self-consistent map, as the point disparity error is precisely what is minimized during point cloud alignment (26). For additional images of the shipwreck area generated using the prior and posterior navigation solutions, see [52].

TABLE VIII: Critical values from the ECDF of Fig. 14b. For example, for the INS solution, 95.45% (2σ) of point disparity errors are below 8.51 cm. Note the improvement in the INS+LC solution over the INS+GPS solution.

Solution	50%	1σ	2σ	3σ
INS	1.51 cm	2.24 cm	8.51 cm	18.89 cm
INS+LC	0.75 cm	1.00 cm	2.35 cm	7.23 cm
INS+GPS	1.08 cm	1.45 cm	3.49 cm	8.63 cm

Again, this improvement in map self-consistency has been achieved *without access* to the standard ingredients available in typical state estimation problems. Visualizing the point cloud map and the resulting disparity errors is a straightforward way to verify that the loop-closure measurements have been successfully applied, and that the updates have been smoothly propagated throughout the trajectory.

Compared to the GPS-aided solution, the improvement in map self-consistency that comes from leveraging loop-closure measurements may appear modest. However, an improvement on the order of centimeters may be consequential for certain subsea inspection tasks, such as measuring deformation in manmade structures. In this respect, the methodology of Section III offers a valuable addition to inspection and metrology work. This is especially true for dead-reckoned solutions, but remains true even when localizing measurements are available, for example LBL, USBL, or GPS measurements.

A Monte Carlo experiment was conducted on the Wiarnton field dataset to test the effectiveness of the loop-closure measurement outlier rejection method discussed in Section III-B3. To run the experiment, between one and five of the seven loop-closure measurements were randomly replaced by randomly generated measurements. Thirty Monte Carlo trials were conducted for each outlier corruption level, for a total of 150 trials. The number of trials at each corruption level was chosen to provide a representative statistical sample. Additionally, the experimental results obtained using 30 trials per corruption level were very similar to results obtained when using 20 and 25 trials per level.

The outlier measurements were generated to mimic the outliers experimentally observed in the detector/descriptor study in Section III-A2. Outlier position measurements were uniformly sampled so that $\|\mathbf{r}^{xy}\| \leq 5$ m and $r^z \in [[-0.5$ m, 0.5 m] \cup [13.5 m, 14.5 m]], with $\mathbf{r}^{\text{out}} = [(\mathbf{r}^{xy})^T \ r^z]^T$. This reflects both the planar search bound used to detect loop-closure candidates (27) as well as the range “flipping” effect discussed in Section III-A2. Outlier attitude measurements were uniformly sampled according to $\phi_j^{\text{out}} \in (-\pi, \pi)$ rad, $j = 1, 2, 3$. An outlier measurement $\Xi_{\ell_1 \ell_2}^{\text{out}}$ is then generated according

$$\Xi_{\ell_1 \ell_2}^{\text{out}} = \begin{bmatrix} \mathbf{C}^{\text{out}}(\phi^{\text{out}}) & \mathbf{r}^{\text{out}} \\ \mathbf{0} & 1 \end{bmatrix}. \quad (73)$$

Results from this experiment are summarized throughout Fig. 15. All 150 Monte Carlo trials are plotted in Fig. 15a, along with the ground-truth “INS+GPS” trajectory and the

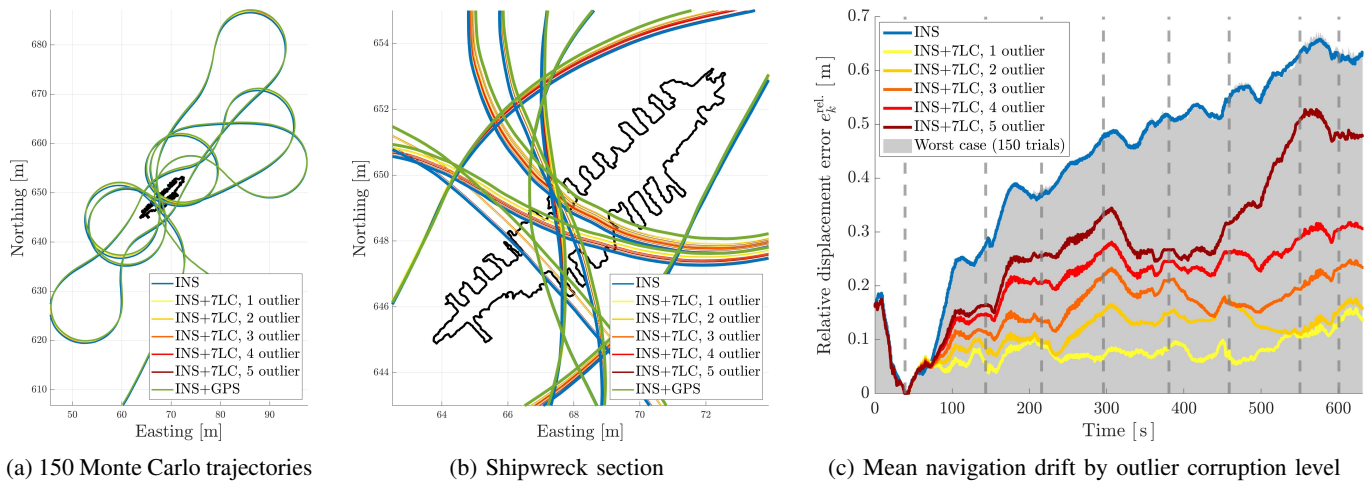


Fig. 15: Trajectory estimates and associated relative drift errors (70) for 150 Monte Carlo trials. For each trial, between 1 and 5 of the seven loop-closure measurements are replaced by outlier measurements (73), with 30 trials conducted per outlier corruption level. The outlier rejection method discussed in Section III-B3 is effective in rejecting false loop-closure measurements, with no failures visible in Fig. 15a. The zoom in Fig. 15b shows a graceful decay from the ground-truth “INS+GPS trajectory” to the prior “INS” estimate as more outlier measurements are added. This is confirmed by the relative displacement error plot in Fig. 15c, which simply follows the trend of measurement removal first seen in the ablation study of Fig. 12. The worst-case position drift at each time step measured across all 150 trials is shown as a grey patch in Fig. 15c. When compared against the relative displacement error from the prior “INS” estimate (blue line), the proposed methodology is seen to produce estimates that are, at any given time, *no worse than the prior estimate, even in instances with extreme outlier rates*.

prior “INS” trajectory estimate. No visible navigation failures are seen in Fig. 15a, implying the adaptive robust cost function is effective in rejecting false loop-closure measurements. A zoom of the shipwreck region in Fig. 15b shows the Monte Carlo trajectory samples gracefully decaying from the ground-truth solution to the prior estimate as more outliers are included. This behaviour is also seen in the relative displacement error plot of Fig. 15c, where the mean relative (x, y) navigation drift (70) is plotted for each outlier corruption level. The trend of higher outlier rates producing larger relative drift values is reminiscent of the ablation study summarized in Fig. 12, in which loop-closure measurements are simply removed from the solution. This provides sound evidence that the proposed outlier rejection algorithm is successful in identifying and removing false loop-closure measurements.

Finally, the grey patch in Fig. 15c shows the worst-case relative displacement error at each time step across all 150 Monte Carlo trials. Compared to the prior “INS” estimate (blue curve), it is clear that, for this dataset, the proposed methodology delivers worst-case posterior estimates that are, at any given time, *no worse than the prior estimate, even in instances with extreme outlier rates*.

V. CONCLUSION

This work presents a novel and comprehensive method for systematically conditioning the output of a COTS DVL-INS navigation system on loop-closure measurements for the purpose of improving the *self-consistency* [4, 60] of the resulting bathymetric map. The method relies on a combination of relative pose and white-noise-on-acceleration

[28] error terms to smoothly integrate the measurements in a batch state estimation framework.

The first contribution of this work is the development of a robust front-end algorithm for computing high-precision loop-closure measurements from 3D scans of challenging underwater environments. Second, loop-closure measurements are cleanly incorporated into an existing state estimate via a factor graph optimization framework, *without access* to raw sensor measurements, sensor models, or other information typically required in conventional state estimation problems.

The effectiveness of the proposed method was demonstrated for both simulated and field datasets using loop-closure measurements from an underwater laser scanner. The same simple hyperparameter structure was used for both studies, with good results. For the field results, conditioning the dead-reckoned DVL-INS estimate on loop-closure measurements produced a markedly more self-consistent point cloud map of an underwater shipwreck. Incorporating all seven loop-closure measurements resulted in a maximum relative position drift of 8.4 cm over a 576 m trajectory, with a final position error of $6.83e-3\%$ of distance traveled. This represents an order of magnitude improvement over unaided commercial DVL-INS systems. Additionally, the proposed methodology was demonstrated to be robust to false loop-closure measurements.

Future work will primarily focus on hyperparameter training through an expectation-maximization framework, for example [68] and [69]. The algorithm will be tested over longer trajectories with more varied terrain, including open seabed [8]. Finally, future work may also incorporate image information, in the form of conventional image descriptors and textured point cloud maps.

ACKNOWLEDGMENT

The authors would like to thank Ryan Wicks of Voyis for providing experimental data and guidance, and Martin Jørgensen of Sonardyne International for providing simulation data and helpful feedback.

REFERENCES

- [1] A. Kim and R. M. Eustice, "Real-time visual SLAM for autonomous underwater hull inspection using visual saliency," *IEEE Trans. Robot.*, vol. 29, no. 3, pp. 719–733, 2013.
- [2] S. Suresh, E. Westman, and M. Kaess, "Through-water stereo SLAM with refraction correction for AUV localization," *IEEE Robot. Autom. Lett. (RAL)*, vol. 4, no. 2, pp. 692–699, 2019.
- [3] S. Rahman, A. Q. Li, and I. Rekleitis, "SVIn2: An underwater SLAM system using sonar, visual, inertial, and depth sensor," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, IEEE, 2019, pp. 1861–1868.
- [4] C. Roman and H. Singh, "A self-consistent bathymetric mapping algorithm," *J. Field Robot.*, vol. 24, no. 1-2, pp. 23–50, 2007.
- [5] S. Barkby, S. B. Williams, O. Pizarro, and M. V. Jakuba, "A featureless approach to efficient bathymetric SLAM using distributed particle mapping," *J. Field Robot.*, vol. 28, no. 1, pp. 19–39, 2011.
- [6] P. Ozog, N. Carlevaris-Bianco, A. Kim, and R. M. Eustice, "Long-term mapping techniques for ship hull inspection and surveillance using an autonomous underwater vehicle," *J. Field Robot.*, vol. 33, no. 3, pp. 265–289, 2016.
- [7] A. Palomer, P. Ridaou, and D. Ribas, "Inspection of an underwater structure using point-cloud SLAM with an AUV and a laser scanner," *J. Field Robot.*, vol. 36, no. 8, pp. 1333–1344, 2019.
- [8] T. Hitchcox and J. R. Forbes, "A point cloud registration pipeline using Gaussian process regression for bathymetric SLAM," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, IEEE, 2020, pp. 4615–4622.
- [9] L. Paull, S. Saeedi, M. Seto, and H. Li, "AUV navigation and localization: A review," *IEEE J. Ocean. Eng.*, vol. 39, no. 1, pp. 131–149, 2014.
- [10] M. V. Jakuba, C. N. Roman, H. Singh, C. Murphy, C. Kunz, C. Willis, T. Sato, and R. A. Sohn, "Long-baseline acoustic navigation for under-ice autonomous underwater vehicle operations," *J. Field Robot.*, vol. 25, no. 11-12, pp. 861–879, 2008.
- [11] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Trans. Robot.*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [12] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software*. John Wiley & Sons, 2004.
- [13] *SPRINT-Nav datasheet*, SPRINT-Nav 500, Sonardyne, Jun. 2021. [Online]. Available: https://www.sonardyne.com/wp-content/uploads/2021/07/Sonardyne_8253_SPRINT_Nav.pdf.
- [14] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "g2o: A general framework for graph optimization," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, IEEE, 2011, pp. 3607–3613.
- [15] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, "iSAM2: Incremental smoothing and mapping using the Bayes tree," *Int. J. Robot. Res.*, vol. 31, no. 2, pp. 216–235, 2012.
- [16] F. Dellaert, M. Kaess, et al., "Factor graphs for robot perception," *Foundations Trends Robot.*, vol. 6, no. 1-2, pp. 1–139, 2017.
- [17] S. Barkby, S. B. Williams, O. Pizarro, and M. V. Jakuba, "Bathymetric particle filter SLAM using trajectory maps," *Int. J. Robot. Res.*, vol. 31, no. 12, pp. 1409–1430, 2012.
- [18] A. Palomer, P. Ridaou, and D. Ribas, "Multibeam 3D underwater SLAM with probabilistic registration," *Sensors*, vol. 16, no. 4, p. 560, 2016.
- [19] H. Johannsson, M. Kaess, B. Englot, F. Hover, and J. Leonard, "Imaging sonar-aided navigation for autonomous underwater harbor surveillance," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, IEEE, 2010, pp. 4396–4403.
- [20] A. Kim and R. M. Eustice, "Active visual SLAM for robotic area coverage: Theory and experiment," *Int. J. Robot. Res.*, vol. 34, no. 4-5, pp. 457–475, 2015.
- [21] P. V. Teixeira, M. Kaess, F. S. Hover, and J. J. Leonard, "Underwater inspection using sonar-based volumetric submaps," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, IEEE, 2016, pp. 4288–4295.
- [22] J. Li, M. Kaess, R. M. Eustice, and M. Johnson-Roberson, "Pose-graph SLAM using forward-looking sonar," *IEEE Robot. Autom. Lett. (RAL)*, vol. 3, no. 3, pp. 2330–2337, 2018.
- [23] T. D. Barfoot, *State Estimation for Robotics*. Cambridge University Press, 2017.
- [24] J. Sola, J. Deray, and D. Atchuthan, "A micro Lie theory for state estimation in robotics," *arXiv preprint arXiv:1812.01537*, 2018.
- [25] J. Arsenault, "Practical considerations and extensions of the invariant extended Kalman filtering framework," Master's thesis, Department of Mechanical Engineering, McGill University, 2019.
- [26] G. S. Chirikjian, *Stochastic Models, Information Theory, and Lie Groups, Volume 2: Analytic Methods and Modern Applications*. Springer Science & Business Media, 2011, vol. 2.
- [27] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT press, 2006.
- [28] S. Anderson and T. D. Barfoot, "Full STEAM ahead: Exactly sparse Gaussian process regression for batch continuous-time trajectory estimation on SE(3)," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, IEEE, 2015, pp. 157–164.
- [29] T. D. Barfoot, C. H. Tong, and S. Särkkä, "Batch continuous-time trajectory estimation as exactly sparse Gaussian process regression," in *Robot.: Sci. Syst. (RSS)*, Citeseer, vol. 10, 2014.
- [30] E. Eade, "Lie groups for computer vision," Tech. Rep., 2014.
- [31] H. Ye, Y. Chen, and M. Liu, "Tightly coupled 3D lidar inertial odometry and mapping," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, IEEE, 2019, pp. 3144–3150.
- [32] J. Lin and F. Zhang, "A fast, complete, point cloud based loop closure for LiDAR odometry and mapping," *arXiv preprint arXiv:1909.11811*, 2019.
- [33] P. Babin, P. Giguère, and F. Pomerleau, "Analysis of robust functions for registration algorithms," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, IEEE, 2019, pp. 1451–1457.
- [34] Y. Chen and G. Medioni, "Object modelling by registration of multiple range images," *Image Vis. Comput.*, vol. 10, no. 3, pp. 145–155, 1992.
- [35] K.-L. Low, "Linear least-squares optimization for point-to-plane ICP surface registration," *Chapel Hill, University of North Carolina*, vol. 4, no. 10, pp. 1–3, 2004.
- [36] J. Neira and J. D. Tardós, "Data association in stochastic mapping using the joint compatibility test," *Trans. Robot. Autom.*, vol. 17, no. 6, pp. 890–897, 2001.
- [37] H. Yang, J. Shi, and L. Carlone, "TEASER: Fast and certifiable point cloud registration," *IEEE Trans. Robot.*, vol. 37, no. 2, pp. 314–333, 2020.
- [38] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms (FPFH) for 3D registration," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, IEEE, 2009, pp. 3212–3217.
- [39] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, 2004.
- [40] R. B. Rusu and S. Cousins, "3D is here: Point cloud library (PCL)," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, IEEE, 2011, pp. 1–4.
- [41] Y. Zhong, "Intrinsic shape signatures: A shape descriptor for 3D object recognition," in *IEEE Int. Conf. Comput. Vis. (ICCV) Workshops*, IEEE, 2009, pp. 689–696.
- [42] I. Sipiran and B. Bustos, "Harris 3D: A robust extension of the Harris operator for interest point detection on 3D meshes," *Vis. Comput.*, vol. 27, no. 11, pp. 963–976, 2011.
- [43] F. Tombari, S. Salti, and L. Di Stefano, "Unique signatures of histograms for local surface description," in *Eur. Conf. Comput. Vis.*, Springer, 2010, pp. 356–369.
- [44] Q.-Y. Zhou, J. Park, and V. Koltun, "Fast global registration," in *Eur. Conf. Comput. Vis.*, Springer, 2016, pp. 766–782.
- [45] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [46] 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 *Proc. IEEE Conf. Control Technol. Appl.*, IEEE, 2020, pp. 43–48.
- [47] M. Pauly, M. Gross, and L. P. Kobbelt, "Efficient simplification of point-sampled surfaces," in *IEEE Vis.*, IEEE, 2002, pp. 163–170.
- [48] T. Hitchcox and J. R. Forbes, "Comparing robust cost functions for bathymetric point cloud registration," in *IEEE/OES Auton. Underwater Veh. Symp. (AUV)*, IEEE, 2020, pp. 1–6.
- [49] J. M. Phillips, R. Liu, and C. Tomasi, "Outlier robust ICP for minimizing fractional RMSD," in *Int. Conf. 3D Digital Imag. Model.*, IEEE, 2007, pp. 427–434.

- [50] F. Pomerleau, F. Colas, R. Siegwart, and S. Magnenat, "Comparing ICP variants on real-world data sets: Open-source library and experimental protocol," *Auton. Robots*, vol. 34, no. 3, pp. 133–148, 2013.
- [51] M. Brossard, S. Bonnabel, and A. Barrau, "A new approach to 3D ICP covariance estimation," *IEEE Robot. Autom. Lett. (RAL)*, vol. 5, no. 2, pp. 744–751, 2020.
- [52] T. Hitchcox and J. R. Forbes, "Improving self-consistency in underwater mapping through laser-based loop closure (extended)," *arXiv preprint arXiv:2301.02297*, 2023.
- [53] J. Farrell, *Aided Navigation: GPS with High Rate Sensors*. McGraw-Hill, Inc., 2008.
- [54] C. Van Loan, "Computing integrals involving the matrix exponential," *IEEE Trans. Autom. Control*, vol. 23, no. 3, pp. 395–404, 1978.
- [55] N. Sünderhauf and P. Protzel, "Switchable constraints for robust pose graph SLAM," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, IEEE, 2012, pp. 1879–1884.
- [56] G. H. Lee, F. Fraundorfer, and M. Pollefeys, "Robust pose-graph loop-closures with expectation-maximization," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, IEEE, 2013, pp. 556–563.
- [57] H. Yang, P. Antonante, V. Tzoumas, and L. Carlone, "Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection," *IEEE Robot. Autom. Lett. (RAL)*, vol. 5, no. 2, pp. 1127–1134, 2020.
- [58] T. Hitchcox and J. R. Forbes, "Mind the gap: Norm-aware adaptive robust loss for multivariate least-squares problems," *IEEE Robot. Autom. Lett. (RAL)*, vol. 7, no. 3, pp. 7116–7123, 2022.
- [59] J. G. Mangelson, M. Ghaffari, R. Vasudevan, and R. M. Eustice, "Characterizing the uncertainty of jointly distributed poses in the Lie algebra," *IEEE Trans. Robot.*, vol. 36, no. 5, pp. 1371–1388, 2020.
- [60] R. Kümmerle, B. Steder, C. Dornhege, M. Ruhnke, G. Grisetti, C. Stachniss, and A. Kleiner, "On measuring the accuracy of SLAM algorithms," *Auton. Robots*, vol. 27, no. 4, pp. 387–407, 2009.
- [61] C. Roman and H. Singh, "Consistency based error evaluation for deep sea bathymetric mapping with robotic vehicles," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, IEEE, 2006, pp. 3568–3574.
- [62] T. Y. Tang, D. J. Yoon, and T. D. Barfoot, "A white-noise-on-jerk motion prior for continuous-time trajectory estimation on SE(3)," *IEEE Robot. Autom. Lett. (RAL)*, vol. 4, no. 2, pp. 594–601, 2019.
- [63] P. Cieślak, "Stonefish: An advanced open-source simulation tool designed for marine robotics, with a ROS interface," in *OCEANS*, IEEE, 2019, pp. 1–6.
- [64] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *SIGGRAPH*, 1996, pp. 303–312.
- [65] G. Kurz, M. Holoch, and P. Biber, "Geometry-based graph pruning for lifelong SLAM," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, IEEE, 2021, pp. 3313–3320.
- [66] P. Tétreault, J. Kouba, P. Héroux, and P. Legree, "CSRS-PPP: An internet service for GPS user access to the Canadian Spatial Reference Frame," *Geomatica*, vol. 59, no. 1, pp. 17–28, 2005.
- [67] R. M. Alkan, S. Erol, I. M. Ozulu, and V. Ilci, "Accuracy comparison of post-processed PPP and real-time absolute positioning techniques," *Geomatics, Nat. Hazards Risk*, vol. 11, no. 1, pp. 178–190, 2020.
- [68] J. N. Wong, D. J. Yoon, A. P. Schoellig, and T. D. Barfoot, "A data-driven motion prior for continuous-time trajectory estimation on SE(3)," *IEEE Robot. Autom. Lett. (RAL)*, vol. 5, no. 2, pp. 1429–1436, 2020.
- [69] T. D. Barfoot, J. R. Forbes, and D. J. Yoon, "Exactly sparse Gaussian variational inference with application to derivative-free batch nonlinear state estimation," *Int. J. Robot. Res.*, vol. 39, no. 13, pp. 1473–1502, 2020.



Thomas Hitchcox received his B.Eng. and M.Eng. degrees in mechanical engineering in 2015 and 2018, respectively, from McGill University, Montreal, QC, Canada. He is currently a Ph.D. Candidate with the Department of Mechanical Engineering at McGill. His research interests include state estimation, computer vision, and robust algorithms for point cloud filtering and alignment.



James Richard Forbes James Richard Forbes received the B.A.Sc. degree in Mechanical Engineering (Honours, Co-op) from the University of Waterloo, Waterloo, ON, Canada in 2006, and the M.A.Sc. and Ph.D. degrees in Aerospace Science and Engineering from the University of Toronto Institute for Aerospace Studies (UTIAS), Toronto, ON, Canada in 2008 and 2011, respectively. James is currently an Associate Professor and William Dawson Scholar in the Department of Mechanical Engineering at McGill University, Montreal, QC, Canada. James is a Member of the Centre for Intelligent Machines (CIM), and a Member of the Group for Research in Decision Analysis (GERAD). James was awarded the McGill Association of Mechanical Engineers (MAME) Professor of the Year Award in 2016, the Engineering Class of 1944 Outstanding Teaching Award in 2018, and the Carrie M. Derick Award for Graduate Supervision and Teaching in 2020. The focus of James' research is navigation, guidance, and control of robotic systems.