

Stable Multi-Drone GNSS Tracking System for Marine Robots

Shuo Wen^{1,2†*} Edwin Meriaux^{1,2,3†*} Mariana Sosa Guzmán^{1,2*}
 Zhizun Wang^{1,2} Junming Shi^{1,2} Gregory Dudek^{1,2}

Abstract—Stable and accurate tracking is essential for marine robotics, yet Global Navigation Satellite System (GNSS) signals vanish immediately below the sea surface. Traditional alternatives suffer from error accumulation, high computational demands, or infrastructure dependence. In this work, we present a multi-drone GNSS-based tracking system for surface and near-surface marine robots. Our approach combines efficient visual detection, lightweight multi-object tracking, GNSS-based triangulation, and a confidence-weighted Extended Kalman Filter (EKF) to provide stable GNSS estimation in real time. We further introduce a cross-drone tracking ID alignment algorithm that enforces global consistency across views, enabling robust multi-robot tracking with cooperative aerial coverage. We validate our system in diversified complex settings to show the accuracy and robustness of the proposed algorithm.

I. INTRODUCTION

While satellite-based positioning is widely accepted for surface marine robots, its effectiveness diminishes once the robots descend even a very short distance below the ocean surface or if the antenna is wet with salt water. This fundamental limitation prevents marine robots from maintaining continuous GNSS coverage, creating a critical gap in localization capability. As a result, robust alternatives are essential to ensure accurate and persistent positioning in real-world marine environments.

A variety of onboard localization methods have been proposed, including inertial sensors combined with DVLs [1, 2]. SLAM-based methods have been applied to improve autonomy, but are computationally expensive and unstable in dynamic environments [3, 4]. Acoustic techniques such as Long Baseline (LBL) and Ultra Short Baseline (USBL) localization rely on fixed beacons [5, 6]. While effective under certain conditions, these methods suffer from issues such as accumulated drift, reliance on infrastructure, and high computational cost.

To mitigate these drawbacks, offboard localization transfers sensing and computation to external platforms, reducing the hardware and energy burden on the underwater robot. This opens the possibility for aerial-underwater cooperation, with drones serving as mobile positioning beacons. Prior work has demonstrated the feasibility of using a single drone for this purpose [7], but relying on a single viewpoint makes the system susceptible to occlusions, limited coverage, and frequent tracking loss, while also creating a single point of failure. Our system instead employs multiple drones,




Fig. 1: GNSS trajectory estimates of the same marine robot across two trials from Categories 1 and 2 in open-water bodies. The color gradient represents localization error, with brighter colors indicating larger errors. White arrows mark the direction of motion along each trajectory. The mean error is 0.842 m for Mission 1 and 1.282 m for Mission 2.

ensuring overlapping coverage and redundancy: if one drone temporarily loses sight of the target, the others can continue providing accurate observations.

In this paper, we describe and evaluate a system for estimating the GNSS position of mobile vehicles, focused specifically on marine robots, as long as they remain close enough to the surface to remain visually observable. Building on prior work done in aerial localization [7], we extend and adapt it to a multi-drone, multi-robot setting, providing a robust, scalable, and low-cost solution for surface-level marine tracking. An illustration of the setup is shown in Fig. 2. To ensure stable ID assignment in long-range tracking tasks, we introduce a hybrid matching strategy that combines weighted similarity scores from both image and GNSS domains. We further develop an ID alignment algorithm to synchronize robot tracking IDs across different drones. Finally, by integrating the Extended Kalman Filter with confidence-based estimation aggregation, our system achieves an average relative localization error of less than 1 m in the most ideal trials, and only 1.7 m in the more challenging cases. Sample estimated data can be seen in Fig. 1, where 2 sample tracking missions are shown. The framework applies to both submerged and surface vehicles, as well as non-instrumented objects such as humans in search and rescue scenarios or marine objects. Experiments are conducted using two marine robots [8] from Independent Robotics, and the implementation code can be found at https://github.com/stevvwen/stable_multidrone with the data found at [9].

† Corresponding authors * Co-first authors. ¹ McGill University ² MILA-Quebec AI Institute ³ Université Paris-Saclay, CentraleSupélec. Emails: shuo.wen@mail.mcgill.ca; edwin.meriaux@mail.mcgill.ca

II. BACKGROUND

Accurate localization of marine robots, marine creatures, and even people in a Human Robot Interaction (HRI) contexts is crucial for many applications, including ocean exploration [10], environmental monitoring [11], underwater infrastructure inspection [12], marine research [13], and search and rescue operations [14, 15].

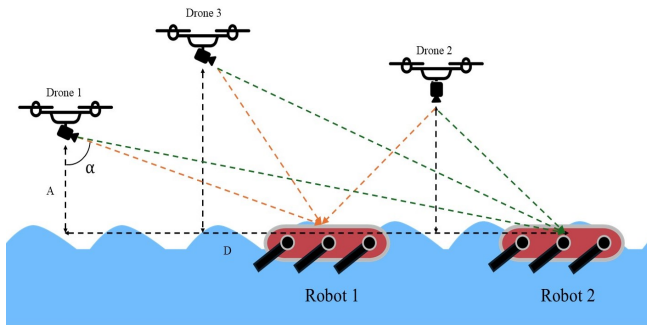


Fig. 2: Tracking system with three aerial drones collaboratively localizing two marine robots operating near the water surface.

A. Localization for Marine Robotics

Marine robotic localization has traditionally focused on enabling a robot to estimate its own position underwater. Earlier approaches relied on onboard sensing, balancing trade-offs between cost, complexity, robustness, and accuracy. A common method combines inertial sensors with DVLs, which estimate velocity relative to water but accumulate drift over time [1, 2, 16]. To mitigate this, acoustic methods such as Long Baseline (LBL) and Ultra Short Baseline (USBL) triangulate signals from beacons or transponders [5, 6, 17]. Extensions include model-aided localization for gliders, using kinematic models and single-beacon timing within an EKF [18], while newer frameworks like 3D-BLUE [19] and low-power 3D localization systems [20] improve scalability and efficiency via piezo-electric backscatter.

To reduce dependence on external infrastructure, Simultaneous Localization and Mapping (SLAM) has been explored with methods such as EKF, SEIF, and FastSLAM [21–23]. A notable example is Sonar Visual Inertial SLAM, fusing stereo vision, IMU, and sonar for GNSS-free mapping of underwater structures [24]. While SLAM enhances autonomy, challenges remain in computation and robustness at large-scale or highly dynamic environments [3, 4, 25].

Beyond onboard approaches, hybrid strategies have also emerged, including cross-domain localization aligning visual data [26, 27], whether it is aerial or water-based [28], cooperative bearing-only localization integrating inertial and depth sensors [29], and multi-robot frameworks where agents share mapping and exploration tasks to improve efficiency [14].

Offboard methods shift sensing and computation to external systems, making them cheaper, lighter, and less power-demanding for the robot. A representative example is the

use of aerial platforms. To mitigate drift, underwater robots need to surface periodically for positional recalibration, for which inexpensive drone-mounted GNSS modules offer a practical solution. UAV-based localization has been demonstrated with onboard object detection [30] and cooperative systems like Sunflower, which combines aerial laser sensing with underwater communication [31]. While accurate and robust, Sunflower requires specialized laser infrastructure, which limits its scalability in resource-constrained settings.

Offboard localization generally occurs in two modes: the robot may be informed of its estimated position, or remain uninformed while an external system tracks it. The latter is the focus of the paper and is particularly useful in shallow-water scenarios such as aquaculture inspections [32], flood responses, and more generally the tracking of near-surface objects.

B. Multi Camera Multi Object Tracking (MCMOT)

MCMOT overcomes the limitations of single-camera systems, such as occlusions and restricted views, by combining observations from distributed cameras for broader coverage and more reliable localization. Early work used centralized fusion of multi-camera streams [33], but scalability and latency issues led to distributed alternatives, including overlapping camera networks for activity monitoring [34]. Subsequent research explored diverse fusion strategies, from feature fusion for crowded scenes [35] and geometry-based triangulation for 3D drone tracking [36], to consensus-based estimation [37], decentralized inference [38], and real-time association frameworks [39], each trading off robustness, scalability, and efficiency. A particularly relevant direction is decentralized tracking, where cameras share only compact state information rather than raw data [40], removing communication bottlenecks and enabling real-time scalability. Building on these insights, our work adopts a distributed MCMOT framework for multi-drone object localization, leveraging geometric consistency across views to achieve robust and scalable tracking in marine environments.

C. Tracking While Moving

Tracking under motion has attracted increasing attention as cameras and platforms (e.g., vehicles, drones, robots) are now frequently deployed in dynamic rather than static settings. This shift introduces challenges such as ego-motion, viewpoint changes, and frame misalignment. In autonomous driving, significant effort has been devoted to mitigating identity switches through techniques such as Global Association Graphs [41] and linker models [42]. In the aerial domain, target association across multiple UAV views has been achieved by fusing appearance, geometric, and distribution features, while collaborative systems leverage temporal consistency to align robot frames without prior pose information [43, 44]. The marine environment presents additional difficulties due to complex image formation effects, further complicating robust tracking [25].

III. METHODOLOGY

The proposed method to observe and estimate the location of the marine robots consists of three main stages: Data Acquisition, Visual Localization, and Estimation.

A. Data Acquisition and Preprocessing

We conducted freshwater field trials with our marine robots to collect data for training the vision model and validating our methods. A simple GNSS receiver, mounted at the robot's center, recorded ground-truth signals. Meanwhile, three drones flew at different altitudes, capturing videos from multiple angles. The experiment setup is illustrated in Fig. 2. To test robustness across applications, we performed trials under varying experimental conditions.

We use established data augmentation techniques from prior work [7] to enhance the robustness of our vision module training. We use generic augmentations (rotation, flipping, cropping, padding, and affine transformations) together with domain-specific ones: motion blur for fast movement of the drones, glass blur for water distortion, and brightness or color shifts for illumination variation. To maximize diversity, all augmentations are applied independently with assigned probabilities and executed in random order.

B. Marine Robot Detection and Tracking

For the detection backbone, we fine-tune YOLO v-11 [45] on the marine robot videos we collected. It is selected for its balance of accuracy, response speed, and lightweight design. This stepwise approach enables the model to build upon its initial knowledge, improving its ability to handle challenging underwater scenarios. The vision model is trained on a desktop equipped with an NVIDIA RTX 4090 GPU and an AMD Ryzen 9 9950X processor using the Ultralytics framework with default training and model hyperparameters. Input images are downsampled to 640×640 pixels to match the native training resolution of YOLO.




Fig. 3: Sample frame from the drone at the highest altitude, showing two tracked marine robots (ID:1 and ID:2) along with two drones at lower altitudes in view.

We integrate the trained vision module into ByteTrack [46], which we employ as our tracking module. ByteTrack distinguishes itself from other algorithms through its simplicity and computational efficiency, while maintaining reliable data association under uncertainty. This makes it particularly

suitable for our marine robotics tracking tasks. Each drone runs ByteTrack independently, with one drone's detections seen in Fig. 3. ID alignment is applied at the initial detection (see III-F). The detection score threshold is set to $\tau = 0.5$, with a lower bound of 0.1 for the second association.

C. GNSS Position Estimation

To estimate GNSS coordinates, we follow the computational approach of [7]. For an image of dimensions $W \times H$, where W and H denote width and height in pixels, the pixel coordinates (x_o, y_o) of the object center are first obtained by vision models such as YOLO as described earlier. For simplicity, we assume the camera has zero yaw and roll angles.

The first step is to compute the object's pixel displacement from the image center:

$$\Delta x = x_o - \frac{W}{2}, \quad \Delta y = y_o - \frac{H}{2}. \quad (1)$$

Using the camera focal length f and the sensor dimensions (S_x, S_y) , the object's angular displacement from the image center in the x and y directions can be expressed as

$$\theta_x = \tan^{-1}\left(\frac{\Delta x}{W} \cdot \frac{S_x}{f}\right), \quad \theta_y = \tan^{-1}\left(\frac{\Delta y}{H} \cdot \frac{S_y}{f}\right) \quad (2)$$

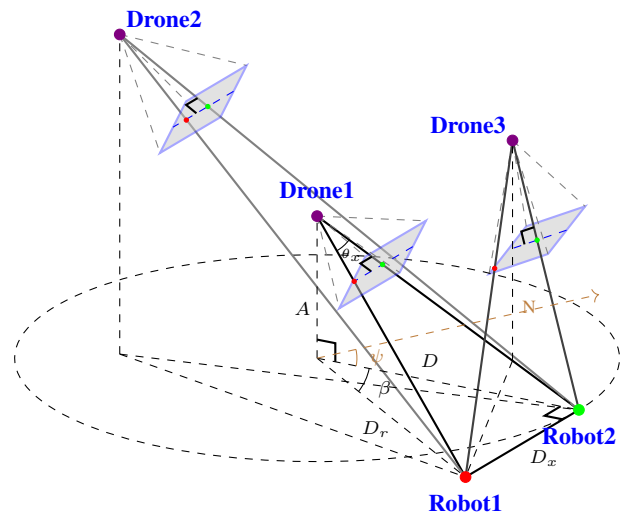


Fig. 4: Position estimation geometry. Image planes and projection rays illustrate the cameras' fields of view and triangulation. For simplicity, Robot 2 is centered in all views.

With the angular displacement between the drone to the object, the ground distance D to the object is

$$D = A \tan(\alpha + \theta_y). \quad (3)$$

where A is the drone altitude and α is the effective camera tilt relative to the vertical, combined with pitch.

The real-world distance corresponding to θ_x is given by:

$$D_x = \sqrt{A^2 + D^2} \tan(\theta_x). \quad (4)$$

The robot's angular displacement relative to the drone and the radial ground distance are given as:

$$\beta = \tan^{-1}\left(\frac{D_x}{D}\right), \quad D_r = \sqrt{D^2 + D_x^2} \quad (5)$$

Fig. 4 provides a 3D illustration of the computation on Drone 1 and Robot 1, augmented with additional information to depict the full geometry with multi-drone system.

The object's displacement is then resolved into the global GNSS frame, where the drone's heading ψ relative to true north and β yield the northward and eastward offset $\Delta N, \Delta E$:

$$\Delta N = D_r \cos(\psi + \beta), \quad \Delta E = D_r \sin(\psi + \beta) \quad (6)$$

The estimated GNSS latitude and longitude coordinates are then obtained by directly incorporating the conversion from linear offsets to angular shifts:

$$\phi^{\text{est}} = \phi^{\text{drone}} + \frac{\Delta N}{\text{FPD}_{\text{lat}}}, \quad \lambda^{\text{est}} = \lambda^{\text{drone}} + \frac{\Delta E}{\text{FPD}_{\text{lon}}} \quad (7)$$

where FPD_{lat} and FPD_{lon} denote the feet-per-degree conversion factors for latitude and longitude, respectively.

Our formulation provides a direct method for estimating a marine robot's geolocation from aerial detections, based on the drone's navigation and imaging parameters.

D. Kalman Filter with Estimation Aggregation

Tracking underwater robots requires combining noisy measurements from multiple sensing platforms into a consistent and robust state estimate. The Kalman filter provides a principled framework for this task, as it recursively fuses new positional measurements with prior state estimates while accounting for uncertainty [47]. At each time step, multiple drones may detect the same marine robot, producing complementary position estimates. To consolidate these measurements, we compute a weighted average of the reported positions, where the weights are given by the normalized YOLO confidence scores of each detection. Formally, given GNSS coordinate estimates $e_i = (\phi_i^{\text{est}}, \lambda_i^{\text{est}})$ from each drone i with detection confidences c_i , the fused estimate e_{final} is defined as: $e_{\text{final}} = \frac{\sum_i c_i e_i}{\sum_i c_i}$.

We then pass e_{final} through the aforementioned Kalman filter to get a stable filtered estimation. While the standard Kalman filter assumes linear system and measurement models, underwater dynamics and sensing modalities are often nonlinear (e.g., bearing and range observations, drag-dependent motion). To address these nonlinearities, we employ an Extended Kalman Filter (EKF) [48], which linearizes the motion and observation models around the current estimate at each step. This extension yields more accurate and stable tracking performance in our marine environment, ensuring that fused multi-drone observations are effectively integrated into the state estimation process. We implement an EKF with a discrete time step of $\Delta t = 0.1\text{s}$, using a process noise acceleration variance of $q_{\text{acc}} = 1.0$ to model system dynamics uncertainty, and a measurement noise variance of $r_{\text{meas}} = 3.0$ to account for sensor observation errors.

E. Minimize Tracking ID Switches with Hybrid Matching

A core step in ByteTrack is computing the similarity score for object matching, where we follow the original design and use Intersection Over Union (IOU) as the primary metric [46]. We omit Re-ID feature, which was discussed in ByteTrack [46], since it is computationally expensive for edge devices and ineffective for our robots, which share almost identical appearances. While the IOU is efficient and simple, camera shake caused by air turbulence and the rapid motions of drones and swimming robots can heavily affect its performance. A missed detection in a single frame can cause serious ID switching. To improve consistency, we fuse IOU scores with normalized GNSS distance scores. Formally, for each tracked object $t \in \mathcal{T}$ and detection $d \in \mathcal{D}$, we compute the pairwise IOU score $I(t, d)$ in the image space. We then combine this score with the GNSS-space distance $L(t, d)$, where the filtered GNSS position of t is used as input instead of the raw estimation for stability. The final hybrid matching score is $s(t, d) = w_1 I(t, d) + w_2 L(t, d)$. We select the basic normalized Haversine distance [49] to be the distance function L . We choose $w_1 = 0.7, w_2 = 0.3$ to balance the reliability of IOU in the image space against the noise inherent in the GNSS space.

Algorithm 1: ID alignment algorithm

Input: Initialized tracks \mathcal{T} from the reference drone, thresholds μ, γ , drones $\{i\}$ with robot detections

Output: \mathcal{Z} : Set of tracking IDs for the drones

```

1  $\mathcal{Z} = \{\}$ 
2 for each drone  $i$  do
3    $\mathcal{D} = [d_j \in \text{drone } i.\text{detections}]$ 
4    $M = [0]^{|\mathcal{D}| \times |\mathcal{T}|}$ 
5   for  $d_j \in \mathcal{D}$  do
6     for  $t_k \in \mathcal{T}$  do
7        $M_{j,k} = L(t_k, d_j)$ 
8    $\text{ID}_{\text{matched}}, \text{ID}_{\text{unmatched}} = \text{LinearAssignment}(M, \mu)$ 
9    $\text{ID}_{\text{new}} = \text{Initialize\_ID}(\text{ID}_{\text{unmatched}}, \gamma)$ 
10   $\mathcal{Z}.\text{append}(\text{ID}_{\text{matched}} \cup \text{ID}_{\text{new}})$ 

```

F. Tracking with ID Consistency

Each drone runs ByteTrack independently while sharing a unified ID counting system. This design ensures that objects maintain consistent identities across views, but without proper alignment, the same robot could be assigned different IDs in different drone frames. To avoid this mismatch, we designate one drone as the reference and initialize its tracks t_i . For every other drone, we compute the pairwise GNSS distance $L(t_i, d_j)$ between the reference tracks and its detections d_j . The Hungarian method [50] is then applied to solve the assignment problem and enforce consistent ID mapping across all drones, subject to a maximum matching threshold μ . Unmatched detections are assigned new IDs only if their confidence score c_j exceeds a threshold γ .

The full procedure is summarized in Algorithm 1. In the experiments, we set $\mu = 30, \gamma = 0.6$.

IV. RESULTS

A. Test Methodology

We designed five categories of experiments to evaluate the accuracy and robustness of our tracking system under progressively more challenging conditions. **Category 1:** Two marine robots traversed long, mostly straight trajectories (up to 600m), while three drones tracked them overhead. These trials established a baseline with minimal changes in speed or heading. **Category 2:** The long-distance experiment was repeated with only two drones to assess the impact of reduced aerial coverage. **Category 3:** The experiment was repeated again with a single drone tracking both marine robots, providing insight into single-observer performance. **Category 4:** Two marine robots followed non-linear trajectories with three drones overhead, testing system performance under less structured motion. **Category 5:** A single marine robot executed a shorter trajectory with multiple sharp U-turns, challenging the system’s ability to handle rapid changes in both speed and direction. The results of these experiments can be seen in Table I. Each test category is marked with the average error and standard deviation.

TABLE I: Results Across All Test Categories

Test Categories	# Drones	Avg. Error	Error STD
Linear Test 1	3	0.942	0.743
Linear Test 2	2	1.042	0.845
Linear Test 3	1	1.111	0.845
Non-Linear	3	1.524	1.435
Hard Turns	3	1.732	1.481

B. Error Estimation

The GNSS estimation accuracy metric is defined as the spatial deviation between the true and estimated trajectories, both consisting of discrete data points sampled at different rates (10 Hz on the drones and 1 Hz on the GNSS receiver). Computing this error is not as simple as taking the haversine distance [49] between points with matching timestamps, as (i) the clocks on the drones and GNSS receiver are not perfectly synchronized, with GNSS receiver timestamps delayed by up to 0.02 s, and (ii) a systematic GNSS offset existed between the two systems due to differences in receivers, satellite constellations, and error models. We use a translation-only variant of the Iterative Closest Point (ICP) algorithm [51] to align two GNSS lines as seen in Fig. 5 where the black arrows indicate the shift. Both lines are first converted from latitude/longitude into a common local East–North coordinate frame using an equirectangular projection around their mean location. At each iteration, every point in the estimation line is paired with its nearest neighbour in the true line using a KD-tree. The best translation for the current matches is then computed as the difference between the centroids of the matched point sets, and the estimated line is shifted accordingly. This process

is repeated until the translation update becomes very small or a maximum number of iterations is reached, yielding aligned estimated GNSS values to properly compute the error between the two lines.




Fig. 5: Category 1 example showing ground truth and ICP-corrected trajectories, with black arrows indicating GNSS error corrections.

C. Analysis of Results

Across all categories in Table I, the system achieves sub-2m mean tracking error with bounded variability, indicating consistent performance under varied operating conditions. The effect of aerial coverage is clear in the linear trials: with three drones, the mean error is lowest at 0.94m (STD 0.74m), increasing to 1.04m (STD 0.85m) with two drones, and 1.11m (STD 0.85m) with a single drone. We can see the algorithm improves as more drones track the same marine robot, even though in the one-drone case, the marine robot is still accurately localized. Motion complexity degrades accuracy in a predictable but controlled manner. Non-linear trajectories tracked by three drones yield a mean error of 1.52m (STD 1.44 m), and the hardest regime with repeated sharp U-turns reaches 1.73 m (STD 1.48 m). These settings stress the image-space association and increase the likelihood of short detection dropouts or transient ID switches. Nonetheless, the hybrid matching that blends image IOU with GNSS-space proximity, together with EKF filtering, limits drift; and our translation-only ICP pre-alignment removes systematic inter-receiver offsets so that residuals primarily reflect instantaneous estimation uncertainty rather than bias.

D. Sources of Error

We identify several key sources of error that limit the accuracy of our results. **Estimation Uncertainty:** Although generally accurate, the vision-based tracking system faces difficulties in precise center localization, which introduces minor positional deviations. Additionally, the limited amount of data from the robot’s field operations restricted our ability to fine-tune YOLO, hindering its performance. The model

occasionally has extended detection failures, which severely disrupt trajectory tracking. **Sensor Noise:** The computation of our approach uses multiple variables that are prone to error. GNSS signals may be distorted by multipath interference or atmospheric delays; altitude estimates can drift with barometric pressure; and angular displacements are affected by lens distortion. Reliable computation further requires clock synchronization across different drones and the GNSS receiver, but timing imperfections add error. Environmental effects such as wind vibrations and temporary signal loss exacerbate these issues, jointly contributing to errors. Moreover, the accuracy of GNSS measurements and position estimation strongly impacts the performance of Algorithm 1. Inaccurate localization can cause ID misalignments across drones, which in turn hinder reliable aggregation of the estimations.




Fig. 6: ID switches under strong air turbulence. Both images are from the same frame at the end of the same trial in Category 1 experiments. (a) Hybrid matching (IOU + normalized GNSS distance, see Section III-E) maintains consistent IDs. (b) IOU-only matching suffers from severe ID switching.

E. Evaluating ID Stability Under Air Turbulence

To demonstrate the stability of our hybrid matching, we evaluate the same tracking system but with different matching methods on the selected trials from Category, where it had many air turbulences, and then compare the results. Such turbulence caused rapid and severe camera shakes, often leading to errors in ID association during tracking. In these challenging trials, the IOU-based tracking system averages 1.33 ID switches per drone per 500 m, whereas our hybrid matching achieves zero switches. An example of an ID switch is shown in Fig. 6.

F. Hardware Testing

To evaluate real-time performance, we deploy the tracking pipeline on an NVIDIA Jetson Xavier (32GB), a compute unit well-suited for embedded drone applications [52]. In the same trial segments used for the results with the presence of two marine robots, the Xavier processes the full pipeline in about 0.176 seconds per frame, averaging over 10 runs on more than 9,000 frames, yielding a 5 Hz position update rate. This frequency enables multiple marine robots to receive location updates from the drones in real time.

V. DISCUSSION AND CONCLUSION

In this paper, we propose and examine a tracking system that employs one or more drones to track marine robots operating near the water’s surface. The system integrates a vision-based detection module, a tracking algorithm, and a data filtering technique to achieve reliable performance. To ensure ID stability during long-range trials and under challenging conditions, we introduce a hybrid matching method that computes weighted similarity scores in both GNSS and image space. In addition, we design an inter-drone matching algorithm to align tracking IDs across different drones. Finally, we develop a data aggregation strategy that fuses complementary estimations based on confidence, further enhancing the overall accuracy and robustness of the system.

The results presented in this paper indicate that multi-drone deployments strike an effective balance between accuracy and robustness for long, mission-scale operations. In contrast, single-drone setups remain suitable for applications where higher error tolerances are acceptable, such as surface inspections or diver guidance. The accuracy gains demonstrated with three drones are likely to improve further with larger fleets, though this hypothesis requires dedicated experimental validation.

To fully demonstrate the algorithm’s utility, financial cost must be considered. A single onboard marine localization system costs roughly the same as four drones used in our experiments, yet it forces robots to remain on the surface and cannot localize multiple robots simultaneously. Our system runs in real time, transmitting high-frequency localization data to marine robots and enabling broader application potential.

In the future, we anticipate additional reductions in both mean error and variance through the use of more extensively trained YOLO object detection models and by upgrading the state estimation pipeline from an EKF to a particle filter, which would better capture the nonlinearities of the system. We also aim to develop a voting-based decentralized approach for ID alignment to further improve consistency. Furthermore, it would be valuable to evaluate the algorithm’s robustness in the face of more intense weather conditions and cases where, for certain frames, some or all the drones suffer data loss. Finally, future work could explore applying the proposed algorithm to moving objects in the water, as well as tracking groups of objects or marine creatures, such as schools of fish, thereby demonstrating the broader applicability of this approach to marine research and conservation.

VI. ACKNOWLEDGEMENT

The authors gratefully thank Chloe Si, Ali, Sasha Nicolas Dolgopolly, Independent Robotics (IR), and the broader community of the McGill Mobile Robotics Lab (MRL) for their crucial assistance in field trials.

REFERENCES

- [1] B. Garau, A. Alvarez, and G. Oliver, "Auv navigation through turbulent ocean environments supported by onboard h-adcp," in *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, 2006, pp. 3556–3561.
- [2] R. McEwen, H. Thomas, D. Weber, and F. Psota, "Performance of an auv navigation system at arctic latitudes," *IEEE Journal of Oceanic Engineering*, vol. 30, no. 2, pp. 443–454, 2005.
- [3] R. M. Eustice, O. Pizarro, and H. Singh, "Visually augmented navigation for autonomous underwater vehicles," *IEEE Journal of Oceanic Engineering*, vol. 33, no. 2, pp. 103–122, 2008.
- [4] P. Newman and J. Leonard, "Pure range-only sub-sea slam," in *2003 IEEE International Conference on Robotics and Automation (Cat. No.03CH37422)*, vol. 2, 2003, 1921–1926 vol.2.
- [5] B. Bingham and W. Seering, "Hypothesis grids: Improving long baseline navigation for autonomous underwater vehicles," *IEEE Journal of Oceanic Engineering*, vol. 31, no. 1, pp. 209–218, 2006.
- [6] N. Kussat, C. Chadwell, and R. Zimmerman, "Absolute positioning of an autonomous underwater vehicle using gps and acoustic measurements," *IEEE Journal of Oceanic Engineering*, vol. 30, no. 1, pp. 153–164, 2005.
- [7] S. Wen, E. Meriaux, M. S. Guzmán, C. Morissette, C. Si, B. Baghi, et al., "Scalable aerial gnss localization for marine robots," *arXiv preprint arXiv:2505.04095*, 2025.
- [8] G. Dudek, P. Giguere, C. Prahacs, S. Saunderson, J. Sattar, L.-A. Torres-Mendez, et al., "Aqua: An amphibious autonomous robot," *Computer*, vol. 40, no. 1, pp. 46–53, 2007.
- [9] E. Meriaux, S. Wen, D. Widhalm, Z. Wang, J. Shi, M. Sosa Guzmán, et al., "Amp2026: A multi-platform marine robotics dataset for tracking and mapping," *arXiv preprint arXiv:2603.04225*, 2026.
- [10] L. D. Barker, M. V. Jakuba, A. D. Bowen, C. R. German, T. Maksym, L. Mayer, et al., "Scientific challenges and present capabilities in underwater robotic vehicle design and navigation for oceanographic exploration under-ice," *Remote Sensing*, vol. 12, no. 16, p. 2588, 2020.
- [11] X. Su, I. Ullah, X. Liu, and D. Choi, "A review of underwater localization techniques, algorithms, and challenges," *Journal of Sensors*, vol. 2020, no. 1, p. 6403161, 2020.
- [12] A. C. Kapoutsis, S. A. Chatzichristofis, L. Doitsidis, J. B. De Sousa, J. Pinto, J. Braga, et al., "Real-time adaptive multi-robot exploration with application to underwater map construction," *Autonomous robots*, vol. 40, no. 6, pp. 987–1015, 2016.
- [13] L. Whitcomb, D. R. Yoerger, H. Singh, and J. Howland, "Advances in underwater robot vehicles for deep ocean exploration: Navigation, control, and survey operations," in *Robotics Research: The Ninth International Symposium*, Springer, 2000, pp. 439–448.
- [14] M. Xanthidis, B. Joshi, J. M. O’Kane, and I. Rekleitis, "Multi-robot exploration of underwater structures," *IFAC-PapersOnLine*, vol. 55, no. 31, pp. 395–400, 2022.
- [15] M. J. Islam, M. Fulton, and J. Sattar, "Toward a generic diver-following algorithm: Balancing robustness and efficiency in deep visual detection," *IEEE Robotics and Automation Letters*, vol. 4, no. 1, pp. 113–120, 2019.
- [16] B. He, H. Zhang, C. Li, S. Zhang, Y. Liang, and T. Yan, "Autonomous navigation for autonomous underwater vehicles based on information filters and active sensing," *Sensors*, vol. 11, no. 11, pp. 10958–10980, 2011.
- [17] D. Ribas, P. Ridao, A. Mallios, and N. Palomeras, "Delayed state information filter for USBL-aided AUV navigation," in *IEEE International Conference on Robotics and Automation*, 2012, pp. 4898–4903.
- [18] J. Sun, F. Hu, W. Jin, J. Wang, X. Wang, Y. Luo, et al., "Model-aided localization and navigation for underwater gliders using single-beacon travel-time differences," *Sensors*, vol. 20, no. 3, p. 893, 2020.
- [19] S. S. Afzal, W. Chen, and F. Adib, "3d-blue: Backscatter localization for underwater robotics," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 11040–11047.
- [20] S. Sharma, D. Van Passen, R. V. Prasad, and K. Chowdhury, "Low power, non-intrusive 3d localization for underwater mobile robots," *Communications Engineering*, vol. 4, no. 1, p. 93, 2025.
- [21] J. Salvi, Y. Petillo, S. Thomas, and J. Aulinas, "Visual slam for underwater vehicles using video velocity log and natural landmarks," in *OCEANS 2008*, 2008, pp. 1–6.
- [22] R. Eustice, H. Singh, and J. Leonard, "Exactly sparse delayed-state filters for view-based slam," *IEEE Transactions on Robotics*, vol. 22, no. 6, pp. 1100–1114, 2006.
- [23] S. Barkby, S. B. Williams, O. Pizarro, and M. V. Jakuba, "Bathymetric particle filter slam using trajectory maps," *The International journal of robotics research*, vol. 31, no. 12, pp. 1409–1430, 2012.
- [24] S. Rahman, A. Q. Li, and I. Rekleitis, "Sonar visual inertial slam of underwater structures," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 5190–5196.
- [25] T. Manderson, J. C. G. Higuera, R. Cheng, and G. Dudek, "Vision-based autonomous underwater swimming in dense coral for combined collision avoidance and target selection," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 1885–1891.
- [26] M. Boldo, M. De Marchi, E. Martini, S. Aldegheri, D. Quaglia, F. Fummi, et al., "Real-time multi-camera 3d

- human pose estimation at the edge for industrial applications,” *Expert Systems with Applications*, vol. 252, p. 124 089, 2024.
- [27] J. Lu, Z. Liang, T. Xie, F. Richter, S. Lin, S. Liu, et al., “Ctnet-x: Camera-to-robot pose estimation in real-world conditions using a single camera,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2025, pp. 1914–1920.
- [28] M. M. D. Santos, G. G. De Giacomo, P. L. J. Drews-Jr, and S. S. C. Botelho, “Cross-view and cross-domain underwater localization based on optical aerial and acoustic underwater images,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4969–4974, 2022.
- [29] H. Damron, A. Q. Li, and I. Rekleitis, “Underwater surveying via bearing only cooperative localization,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 3957–3963.
- [30] S. Sanyal, S. Bhushan, and K. Sivayazi, “Detection and location estimation of object in unmanned aerial vehicle using single camera and GPS,” in *2020 First International Conference on Power, Control and Computing Technologies (ICPC2T)*, 2020, pp. 73–78.
- [31] C. J. Carver, Q. Shao, S. Lensgraf, A. Sniffen, M. Perroni-Scharf, H. Gallant, et al., “Sunflower: Locating underwater robots from the air,” in *Proceedings of the 20th annual international conference on mobile systems, applications and services*, 2022, pp. 14–27.
- [32] A. Hogue, A. German, J. Zacher, and M. Jenkin, “Underwater 3d mapping: Experiences and lessons learned,” in *The 3rd Canadian Conference on Computer and Robot Vision (CRV’06)*, 2006, pp. 24–24.
- [33] Q. Cai and J. Aggarwal, “Tracking human motion in structured environments using a distributed-camera system,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 11, pp. 1241–1247, 1999.
- [34] R. Collins, A. Lipton, H. Fujiyoshi, and T. Kanade, “Algorithms for cooperative multisensor surveillance,” *Proceedings of the IEEE*, vol. 89, no. 10, pp. 1456–1477, 2001.
- [35] T. Teepe, P. Wolters, J. Gilg, F. Herzog, and G. Rigoll, “Lifting multi-view detection and tracking to the bird’s eye view,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 667–676.
- [36] J. Rosner, T. Krzeszowski, A. Świtoński, H. Josiński, W. Lindenheim-Locher, M. Zieliński, et al., “Multi-modal dataset for indoor 3d drone tracking,” *Scientific Data*, vol. 12, no. 1, p. 257, 2025.
- [37] A. T. Kamal, J. A. Farrell, and A. K. Roy-Chowdhury, “Information consensus for distributed multi-target tracking,” in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 2403–2410.
- [38] S. Casao, A. Naya, A. C. Murillo, and E. Montijano, “Distributed multi-target tracking in camera networks,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 1903–1909.
- [39] S. Yang, F. Ding, P. Li, and S. Hu, “Distributed multi-camera multi-target association for real-time tracking,” *Scientific Reports*, vol. 12, no. 1, p. 11 052, 2022.
- [40] M. Taj and A. Cavallaro, “Distributed and decentralized multicamera tracking,” *IEEE Signal Processing Magazine*, vol. 28, no. 3, pp. 46–58, 2011.
- [41] P. Nguyen, K. G. Quach, C. Nhan Duong, N. Le, X.-B. Nguyen, and K. Luu, “Multi-camera multiple 3d object tracking on the move for autonomous vehicles,” in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2022, pp. 2568–2577.
- [42] Y. Zhang, S. Wang, Q. Wang, C. Yan, and R. Fan, “Towards effective multi-moving-camera tracking: A new dataset and lightweight link model,” *arXiv preprint arXiv:2312.11035*, 2023.
- [43] P. Wu, Y. Li, and D. Xue, “Multi-target tracking with multiple unmanned aerial vehicles based on information fusion,” *Drones*, vol. 8, no. 12, p. 704, 2024.
- [44] M. B. Peterson, P. C. Lusk, A. Avila, and J. P. How, “Tcaff: Temporal consistency for robot frame alignment,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025, pp. 15 821–15 827.
- [45] R. Khanam and M. Hussain, “Yolov11: An overview of the key architectural enhancements,” *arXiv preprint arXiv:2410.17725*, 2024.
- [46] Y. Zhang, P. Sun, Y. Jiang, D. Yu, F. Weng, Z. Yuan, et al., “Bytetrack: Multi-object tracking by associating every detection box,” in *European conference on computer vision*, Springer, 2022, pp. 1–21.
- [47] B. Allotta, A. Caiti, L. Chisci, R. Costanzi, F. Di Corato, C. Fantacci, et al., “An unscented kalman filter based navigation algorithm for autonomous underwater vehicles,” *Mechatronics*, vol. 39, pp. 185–195, 2016.
- [48] M. I. Ribeiro, “Kalman and extended kalman filters: Concept, derivation and properties,” *Institute for Systems and Robotics*, vol. 43, pp. 3736–3741, 2004.
- [49] N. R. Chopde and M. Nichat, “Landmark based shortest path detection by using a* and haversine formula,” *International Journal of Innovative Research in Computer and Communication Engineering*, vol. 1, no. 2, pp. 298–302, 2013.
- [50] H. W. Kuhn, “The hungarian method for the assignment problem,” *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.
- [51] P. Besl and N. D. McKay, “A method for registration of 3-d shapes,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.
- [52] L. O. Rojas-Perez and J. Martínez-Carranza, “On-board processing for autonomous drone racing: An overview,” *Integration*, vol. 80, pp. 46–59, 2021.