2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Weighted Group-*k* Consistent Set Maximization for Outlier Rejection of Azimuth-Elevation Measurements

Kalliyan Velasco, Timothy W. McLain, and Joshua G. Mangelson

***Abstract*— Reliable localization in robotics requires robust handling of sensor outliers, particularly in environments where acoustic or bearing measurements are noisy. We propose a replicator-dynamics-based approach for weighted group-*k* consistent set maximization (rGkCM) to identify the densest subsets of mutually consistent measurements in hypergraphs. To complement existing range-based consistency metrics, we introduce a $k = 3$ azimuth-elevation consistency check for bearing measurements to static landmarks. Our method efficiently identifies cliques in weighted *k*-uniform hypergraphs, leveraging the fitness of nodes to guide both pruning and recovery. We evaluate rGkCM on simulated trajectories with varying outlier levels and demonstrate significant computational speedup over the heuristic unweighted GkCM (uGkCM) method while maintaining comparable accuracy. Finally, we validate the approach on a WAM-V autonomous surface vessel equipped with an acoustic beacon and GNSS ground truth, showing effective outlier rejection in a shallow, multipath-prone marina. Results indicate that rGkCM enables robust and efficient outlier rejection for real-world bearing-based localization tasks.**

I. Introduction

Reliable localization is fundamental to autonomous robotics, but in real-world environments sensor measurements are frequently corrupted by noise and outliers. Even a small fraction of outliers can lead to catastrophic estimation failures, making robust outlier rejection a critical component of modern simultaneous localization and mapping (SLAM).

Pairwise consistent set maximization (PCM) has emerged as a widely used tool for pruning spurious sensor observations in pose-graph SLAM, especially in multi-agent settings [1]. Weighted extensions of PCM have demonstrated improvements in runtime and accuracy by formulating the problem as a continuous optimization problem [2]. More recently, PCM has been generalized to the case of group-*k* consistent set maximization (GkCM) for scenarios where pairwise consistency is not sufficient for low degree-of-freedom (DoF) measurements, such as range [3]. However, the computational burden of solving GkCM, even heuristically, limits real-time deployment on robotic platforms.

This work was funded by the Naval Sea Systems Command (NAVSEA), Naval Surface Warfare Center - Panama City Division (NSWC-PCD) under the Naval Engineering Education Consortium (NEEC) Grant Program under award number N00174-23-1-0005. This work was also partially funded by the Office of Naval Research (ONR) under award number N00014-24-1-2301 and the Center for Autonomous Air Mobility and Sensing (CAAMS), a National Science Foundation Industry-University Cooperative Research Center (IUCRC) under NSF award number 2139551, along with significant contributions from CAAMS industry members.

K. Velasco, T. McLain, and J. Mangelson are at Brigham Young University. They can be reached at {kalliyan, mclain, mangelson}@byu.edu.

As motivation for continuing the work on GkCM, we focus on the case of bearing-based measurements. Underwater agents are especially restricted in the types of sensors available for multi-agent localization. GPS and radio signals attenuate rapidly underwater, and visual sensing is severely degraded by poor visibility. Acoustic sensors are one of few options, providing range and bearing information. However, bearing measurements are particularly susceptible to noise and outliers, making robust outlier rejection essential. Beyond underwater robotics, bearing-only measurements are also fundamental in other real-world scenarios, such as vision-based localization and radar-based perception. These applications highlight both the importance of designing a dedicated consistency metric for azimuth–elevation bearings and the need for further development of GkCM.

In this work, we take the natural next step by extending GkCM to edge-weighted *k*-uniform hypergraphs. We demonstrate our method in the context of bearing-based localization using azimuth and elevation measurements using acoustic measurements. The contributions of this work are threefold:

1) We introduce a consistency metric for azimuth-elevation measurements to static landmarks, complementing existing range-based metrics and enabling applications in underwater acoustic positioning.
2) We propose a replicator-dynamics-based approach to consistent set maximization on weighted k-uniform hypergraphs.
3) We present a real-world demonstration of outlier rejection and localization on a WAM-V autonomous surface vessel (ASV) using acoustic beacon landmarks.

II. Related Work

This section reviews the key components underlying our approach: outlier rejection for azimuth–elevation measurements, consistent set maximization, and continuous optimization methods for clique and hypergraph clustering. Together, these elements motivate our proposed framework for weighted group-*k* consistent set maximization.

A. Outlier Rejection for Acoustic Measurements

Ultra-short baseline (USBL) acoustic positioning estimates range from time-of-flight measurements, while a phased hydrophone array computes the azimuth and elevation of incoming signals. Sources of outliers for bearing measurements include multipath propagation, ray bending due to salinity, density, or temperature gradients, and rotational disturbances on the receiving vehicle. In practice, it is common to obtain accurate range data alongside highly corrupted bearing

observations, motivating the need for robust outlier rejection tailored to bearings. As a result, many underwater localization approaches rely primarily on range-only methods [4–6], while more recent bearing-only methods explicitly emphasize the importance of outlier rejection for acoustic bearings [7].

Outlier rejection techniques for acoustic measurements can be broadly divided into two categories: (i) pre-filtering measurements before they enter the state estimator, and (ii) online filtering using outlier-robust estimators [8, 9]. A widely used pre-filtering approach evaluates multiple trilateration hypotheses from range data to obtain a robust initialization, followed by outlier suppression through robust filtering [10]. GkCM provides an alternative by pre-filtering outliers, allowing range and bearing measurements to be filtered independently. This is advantageous in scenarios where ranges are accurate but bearings are unreliable. While GkCM consistency metrics have been proposed for ranges [3], sonar [11], and camera-based mutual observations [12], to our knowledge no such metric exists for azimuth–elevation acoustic bearing measurements.

B. Consistent Set Maximization Approaches

Pairwise consistent set maximization (PCM) [1] is widely used in robotics for pre-processing measurements before adding them to a pose graph. PCM constructs a consistency graph where nodes correspond to measurement associations and edges encode pairwise consistency according to a chosen metric. The maximum clique in this graph is the largest fully connected set of nodes, and identifies inliers to be retained for localization. Compared to RANSAC [13], PCM has the advantage of not requiring a predefined model for how the measurements should relate. Furthermore, compared to robust cost-function approaches, PCM has the ability to handle situations where a full chain of trusted odometry is not available.

Later works extended PCM to weighted consistency graphs, formulating the problem as a maximum edge-weighted clique [14, 15]. CLIPPER [2] further frames the problem as the densest edge-weighted clique (DEWC) and applies a continuous relaxation inspired by [16], later introducing a convex semidefinite relaxation for globally optimal solutions [17]. Weighted graphs allow selecting cliques with strongly consistent edges rather than simply the largest size, improving both runtime and accuracy. They also enable leveraging of non-geometric consistency indicators such as object shape or semantic similarity [18].

To address scenarios where pairwise consistency is insufficient (particularly for low degree-of-freedom measurements), group-k consistent set maximization (GkCM) generalizes PCM to unweighted hypergraphs [3], where an edge connects k nodes instead of a pair. A major computational bottleneck in GkCM is solving the maximum clique problem. Exact algorithms enumerate all cliques to find the largest, while heuristic algorithms iteratively seed and grow cliques based on node degrees and internal connectivity. Although the heuristic approach improves runtime, it remains challenging for real-time deployment [11]. Motivated by the success of weighted graphs in PCM, extending GkCM to edge-weighted hypergraphs promises both improved outlier rejection and better scalability in practice.

C. Continuous Optimization for Clique and Hypergraph Clustering

Identifying dense subgraphs in edge-weighted graphs is a classic NP-hard problem that has been widely studied. One field where this concept has been well addressed is dominant set clustering. While dominant sets do not require full connectivity like maximum cliques, related methods have been shown to be effective heuristics for both maximum clique and densest k-subgraph problems [19, 20]. Two prominent optimization strategies for dominant set clustering are replicator dynamics and Frank-Wolfe methods.

Frank-Wolfe methods are gradient-ascent based: at each iteration, a linearized version of the problem is solved to obtain a search direction, followed by step-size selection (e.g., via line search) and an update along that direction [21, 22]. When the initial point lies within the feasible set, no additional projection is required.

Replicator dynamics, originating from evolutionary game theory, models a population of strategies evolving according to a fitness function. Strategies with higher payoffs are adopted by a larger fraction of the population in subsequent generations. This framework has been applied to hypergraph clustering, where nodes correspond to strategies and fitness encodes the importance of each node [23, 24]. Extensions such as path-following replicator dynamics introduce a path parameter and truncated simplex projection, enabling multiple clusters to be revealed [20]. Additional improvements, including line maximization for step size selection, have further enhanced convergence and performance [25]. Variants like infection-immunization dynamics have also been applied in robotics contexts, including geolocalization [26, 27].

Building on these insights, we use replicator-dynamics-based clustering as a stepping stone towards consistent set maximization on edge-weighted k-uniform hypergraphs. This extension enables the implementation of weighted group-k consistency metrics for outlier rejection in scenarios with low degree-of-freedom measurements.

III. Edge-Weighted Group-k Consistent Set Maximization

Consistent set maximization aims to identify the largest set of mutually consistent associations among sensor measurements. This is achieved by defining a consistency metric that depends on the sensor measurement type, and then building a consistency graph or hypergraph. In a standard graph, each node represents a single measurement association and edges represent pairwise consistency. In a k-uniform hypergraph, each hyperedge connects k nodes.

In unweighted GkCM (uGkCM) [3], a set of measurements $\tilde{\mathbf{Z}}$ is group-k internally consistent with respect to a consistency metric C and threshold γ if

$$C(\{\mathbf{z}_0, \ldots, \mathbf{z}_k\}) \leq \gamma, , \quad \forall \quad \{\mathbf{z_0}, \ldots, \mathbf{z_k}\} \in \mathcal{P}_k(\tilde{\mathbf{Z}}), \quad (1)$$

where C is some defined consistency function for the set of measurements $\{\mathbf{z}_0, \ldots, \mathbf{z}_k\}$, γ is a threshold chosen based on the measurement distribution, and $\mathcal{P}_k(\tilde{\mathbf{Z}})$ is the set of all permutations of $\tilde{\mathbf{Z}}$ with cardinality k. For each consistent group of k nodes, a hyperedge is created. The largest fully connected subgraph (maximum clique) in the hypergraph corresponds to the set of inlier associations that can safely be added to the pose graph.

Weighted GkCM extends this formulation by assigning a positive weight to each hyperedge, reflecting the level of consistency. The consistency hypergraph can formally be defined as $H = (V, E, \omega)$, where $V = \{1, \ldots, n\}$ is the set of vertices (associations), E is the set of hyperedges consisting of k-tuples of vertices, and $\omega : E \to \mathbb{R}_+$ assigns a positive weight to each hyperedge using the consistency definition. The goal becomes identifying the densest edge-weighted subset of measurements. In this work, we will refer to unweighted GkCM as uGkCM, and the replicator-dynamics-based approach on weighted hypergraphs as rGkCM.

IV. Consistency Metric for Azimuth–Elevation

A consistency metric quantifies the extent to which a set of measurements agree with one another, and the order k is determined by the number of observations required to evaluate the metric.

A. Computing the Consistency Metric

To complement the $k = 4$ range-based consistency check proposed in [3], we introduce a $k = 3$ consistency metric tailored to azimuth–elevation bearing measurements.

$$C(\beta_a^i, \beta_b^i, \beta_c^i) = \left\| h(\mathbf{X}_{abc}, \mathbf{Z}_{ab}^i) - \beta_c^i \right\|_\Sigma \leq \gamma, \qquad (2)$$

where $\beta = [\alpha, \epsilon]$ denotes the azimuth α and elevation ϵ of beacon i observed from a given pose a, b, or c, $\mathbf{X}_{abc}$ is the triplet of poses $\{a, b, c\}$, and $\mathbf{Z}_{ab}^i$ is the pair of bearing measurements from poses a and b to beacon i. The term $\|\cdot\|_\Sigma$ denotes the Mahalanobis distance [28] between the predicted and observed bearings, and γ is a consistency threshold chosen from the chi-squared distribution with two degrees of freedom.

a) Measurement model: Much like in camera geometry, an azimuth–elevation measurement can be modeled as a line of sight from the robot to a landmark. The landmark position can be estimated from the intersection of two lines of sight. The consistency metric evaluates the measured bearings against the predicted bearing measurements from a third line of sight to the triangulated landmark. The expected bearing is given by

$$h(\mathbf{X}_{abc}, \mathbf{Z}_{ab}^i) = \mathrm{bearing}\big(\mathrm{tri}(\mathbf{X}_{ab}, \mathbf{Z}_{ab}^i), \mathbf{X}_c^i\big), \qquad (3)$$

where $\mathrm{tri}(\cdot)$ denotes the triangulation function that computes the estimated 3D landmark position as observed from poses a and b, and $\mathrm{bearing}(\cdot)$ computes the relative azimuth–elevation measurement from Cartesian coordinates according to the sensor convention. For triangulation, we adopt the midpoint method [13], which estimates the 3D point that minimizes the sum of squared distances between the two lines of sight. Although it is less robust than iterative optimization methods, its computational efficiency makes it well suited for the large number of consistency checks required during hypergraph construction.

b) Uncertainty propagation: To compute Σ, we consider both pose and measurement uncertainties. Let Σ_j denote the joint covariance of poses a, b, and c, which can be extracted from the factor graph using [29], and let Σ_β denote the diagonal covariance matrix of the three corresponding azimuth–elevation measurements. The block covariance is $\Sigma_T = \mathrm{blockdiag}(\Sigma_j, \Sigma_\beta)$. Uncertainty is propagated to the azimuth–elevation space via $\Sigma = H\,\Sigma_T\,H^\top$, where $H = \frac{\partial h}{\partial(\mathbf{X}_{abc}, \beta_{abc})}$ is the Jacobian of the measurement model.

c) Consistency check: The metric (3) evaluates whether the predicted bearing from the triangulated point to pose c is consistent with the measured bearing at c. Each bearing in the triplet is tested against the other two, and the set is considered consistent only if all three checks pass.

d) Weighting: To obtain the edge weight w in the hypergraph, the Mahalanobis distance is converted into a consistency score following [18] and expressed as

$$s(\beta_a^i, \beta_b^i, \beta_c^i) \triangleq \exp\left(-\tfrac{1}{2}\tfrac{C(\beta_a^i, \beta_b^i, \beta_c^i)}{\sigma^2}\right), \qquad (4)$$

where σ is a tunable parameter. Larger Mahalanobis distances yield smaller scores, reflecting weaker consistency among the triplet.

B. Degenerate Configurations

The midpoint triangulation can fail or become ill-conditioned when the two lines of sight are parallel (or nearly parallel) or when the two poses are co-located, since no unique intersection exists in these cases. Near-parallel lines can arise either from a small baseline-to-distance ratio (i.e., the observing poses are close relative to the distance to the landmark) or from a small angle between the two lines of sight (i.e., the landmark is nearly collinear with the poses). These unstable configurations can amplify bearing noise into large position errors. Uncertainty propagation may not fully capture these effects, which can lead to wrongly rejected edges in the hypergraph.

In the acoustic domain, however, the low rate of transmissions partially mitigates these issues. Because pings are typically received sequentially from different beacons, an agent usually moves a sufficient distance between observations to avoid the most severe degeneracies. Nevertheless, in high-rate scenarios with azimuth–elevation updates, many consecutive observations may be taken from nearly the same location, reducing the effective baseline and increasing the likelihood of midpoint failure. In such cases, failed triangulations cause consistency checks to fail, preventing related nodes from participating in the maximum clique and thereby reducing the true positive rate.

V. Finding the Densest Edge-Weighted Clique

The proposed approach to finding the densest edge-weighted clique first solves for the densest cluster in the

hypergraph using replicator dynamics, then peels or adds nodes to the cluster using classic graph-search methods to ensure the clique property is maintained. The scoring of nodes using replicator dynamics assists speed of checking if nodes can be added or removed from a graph, as early termination can happen once the clique property is broken.

A. Densest Cluster Replicator Dynamics

A game in game theory consists of a set of players, strategies, and a payoff function. Inspired by Darwin's principle of survival of the fittest, evolutionary game theory was originally developed to model the adaptation of animal populations. In this setting, a population of infinite size evolves by inheriting strategies from previous generations. Strategies with higher payoffs are adopted more frequently, gradually dominating the population, while strategies with poor fitness eventually go extinct. A Nash equilibrium is reached when no individual agent can achieve a higher payoff than the population average, and an *evolutionarily stable strategy* (ESS) is one that cannot be invaded by a mutant strategy with greater expected payoff.

This concept has been directly mapped to clustering problems for graphs and hypergraphs, and [24] demonstrates how clusters in an edge-weighted k-uniform hypergraph can be represented as an ESS in an evolutionary game. The nodes of the hypergraph correspond to strategies in the game, and the payoff is a function of the edge weights.

For finding a dense cluster in the hypergraph, the expected payoff over the entire population is defined as

$$f(\mathbf{x}) := \sum_{e \in E} \omega(e) \prod_{j \in e} x_j, \tag{5}$$

where $\mathbf{x}$ is the fitness vector (with x_i corresponding to node i in the consistency graph) and $\omega(e)$ is the weight, or consistency, of edge e.

The expected payoff of an individual node j is given by

$$[g(\mathbf{x})]_j := \sum_{e \ni j} \omega(e) \prod_{\ell \in e \setminus \{j\}} x_\ell, \tag{6}$$

where $g(\mathbf{x}) = \nabla f(\mathbf{x})$ is the gradient of the global payoff.

Finding the equilibria of replicator dynamics is equivalent to locally optimizing a homogeneous polynomial over the standard simplex [24]. Each iteration in the optimization problem corresponds to a new generation in the population. The optimization problem can then be written as

$$\hat{\mathbf{x}} = \arg\max_{\mathbf{x} \in \Delta} f(\mathbf{x}), \tag{7}$$

where $\Delta = \{\mathbf{x} \mid \sum_i x_i = 1,\ x_i \in [0,1]\}$ is the standard simplex. Intuitively, this maximizes the overall payoff of the population. In the context of consistent set maximization, nodes gain importance if they are involved in many highly consistent edges, similar to how strategies become more prevalent in a population when they yield higher fitness and are adapted by more members.

The optimization problem is a homogeneous polynomial with nonnegative coefficients, and thus the Baum–Eagon theorem applies [30]. The Baum–Eagon inequality guarantees that the proposed update rule yields monotonic improvement of the objective function. Specifically, the replicator update for node j is

$$x_j(t+1) = x_j(t)\, \frac{[g(\mathbf{x}(t))]_j}{f(\mathbf{x}(t))}, \tag{8}$$

and its vector form is

$$\mathbf{x}(t+1) = \mathrm{proj}_\Delta(\mathbf{x}(t) \odot g(\mathbf{x}(t)))\,, \tag{9}$$

where proj_Δ projects onto the simplex Δ. Projection can be performed via normalization or, when prior information about clique size or number of clusters is available, via truncated simplex projection [20].

To solve the densest edge-weighted clique problem, $\mathbf{x}$ is initialized uniformly. At each iteration, the gradient is computed from the current hypergraph, followed by the replicator update (9). Iterations proceed until convergence, defined when the change in $\mathbf{x}$ falls below a threshold δ. The choice of δ controls a trade-off: larger thresholds yield faster convergence and higher recall but may admit more outliers, while smaller thresholds improve precision at the cost of runtime.

The support of $\mathbf{x}$ reflects the fitness of nodes. After convergence, elements below a cutoff threshold τ are discarded, and remaining nodes represent the densest edge-weighted cluster.

B. Densest Edge-Weighted Clique Refinement

The replicator dynamics optimization does not explicitly enforce the clique property. As a result, the solution may include nodes that are not part of the true maximum clique, or, less frequently, exclude nodes that should belong due to conservative thresholding. To address these issues, we perform post-processing on the replicator solution to ensure a valid clique for consistent set maximization. This involves both *node removal* and *node recovery*, which verify whether adding a node to a candidate clique preserves the clique property. For a k-uniform hypergraph, this requires that all k-tuples formed with the new node and the existing clique members are present in the current hyperedge set [3].

We leverage the fitness scores from replicator dynamics as a heuristic to prioritize node processing, avoiding the need to examine all nodes as exhaustively as done in the exact and heurstic maximum clique solvers presented in [3]. Starting from the replicator solution $\hat{\mathbf{x}}$, we sort nodes by descending fitness and iteratively process nodes.

1) Node Removal: Node removal begins with the lowest-scoring node within the accepted threshold τ. For each node, we check if adding it to the current candidate clique (composed of higher-fitness nodes) preserves the clique property. Nodes that violate the property are removed. Once a node is found to successfully maintain the clique property, or the candidate clique becomes empty, the process stops, and the predicted final clique is returned. If a node is removed, nodes with fitness lower than a removed node are assumed unlikely to belong to the maximum clique, enabling early termination.

Algorithm 1: Densest Edge-Weighted Clique (rGkCM)

Input: Weighted k-uniform hypergraph $H = (V, E, \omega)$, convergence threshold δ, fitness threshold τ
Output: Maximum clique C

```
// Replicator dynamics cluster
```
Initialize $\mathbf{x}^{(0)} \leftarrow \frac{1}{|V|}\mathbf{1}$, $t \leftarrow 0$;
while $\|\mathbf{x}^{(t+1)} - \mathbf{x}^{(t)}\| \geq \delta$ **do**
  $\mathbf{x}^{(t+1)} \leftarrow \text{proj}_{\Delta}(\mathbf{x}^{(t)} \odot g(\mathbf{x}^{(t)}))$;
  $t \leftarrow t + 1$;
end
```
// Clique refinement
```
Sort $\hat{\mathbf{x}}$ in descending order and form x_pairs;
Initialize clique: $C \leftarrow \{i \in V \mid \hat{\mathbf{x}}[i] > \tau\}$;
```
// Node removal
```
$trimmed \leftarrow$ false;
for $i = |C| - 1$ ***down to*** 0 **do**
  $node \leftarrow C[i]$;
  Remove $node$ from C;
  if *not canAddNode*$(C, node)$ **then**
    $trimmed \leftarrow$ true;
  end
  else
    Add $node$ back to C;
    break;
  end
end
```
// Node recovery
```
if $\neg trimmed$ **then**
  for $i = |C|$ ***to*** $|\hat{\mathbf{x}}| - 1$ **do**
    $node \leftarrow x\_pairs[i].first$;
    if *canAddNode*$(C, node)$ **then**
      Add $node$ to C;
    end
    else
      break;
    end
  end
end
return C;

2) Node Recovery: If no nodes were removed from the original clique, we can attempt to recover additional nodes to maximize the clique size. We check if the highest scoring node that did not meet the threshold can be added to the clique, and iteratively process the remaining nodes by descending fitness. Once the clique property is violated, the recovery process terminates, yielding the final maximum clique.

The final algorithm is summarized in Algorithm 1. The resulting procedure identifies a dense edge-weighted clique in a weighted k-uniform consistency hypergraph, combining replicator dynamics with clique-aware post-processing. However, the proposed algorithm is a heuristic that does not guarantee a globally optimal solution.

VI. Experiments

We evaluate the proposed azimuth–elevation consistency check and replicator-based DEWC algorithm (rGkCM) on both simulated and real-world localization scenarios. We compare rGkCM to the heuristic unweighted GkCM (uGkCM), which is currently the fastest method for consistent set maximization on k-uniform hypergraphs. For all experiments, we use the parameters $\delta =$ 1e-3 and $\tau =$ 1e-12.

A. Localization with Simulated Bearings

For simulations, we generate random trajectories of 100 steps. Two types of trajectories are considered: (1) an ideal arc, which minimizes degenerate triangulation configurations, and (2) a randomized trajectory, where the robot selects a random yaw at each step, increasing the likelihood of degenerate configurations (such as poses occupying the same space). Examples of the trajectories can be found in Fig. 1. Beacons are placed randomly according to the trajectory type. At each step, the robot receives observations from five beacons, and we test outlier fractions ranging from 50% to 90%. For each setting, 100 simulations are performed, with 500 hypergraphs evaluated per run across the five beacons. To simplify analysis, the same set of outlier indices is applied to all beacons.

The unweighted hypergraph is evaluated using uGkCM, while the weighted hypergraph uses rGkCM with replicator dynamics. Both methods were run with four threads, with only the gradient computation in rGkCM parallelized. Factor graphs are built and evaluated using GTSAM [31].

Results indicate that uGkCM outperforms rGkCM in extreme outlier scenarios (90% or higher). This is expected, as the replicator dynamics payoff function prioritizes node groups with many internal edges, making it more susceptible to local minima when the true clique is relatively sparse compared to spurious structures. Overall, uGkCM achieves slightly higher true positive rates with comparable false positive rates. However, rGkCM offers a substantial speed advantage, consistently solving problems in under one second, whereas uGkCM may require up to 50 seconds (Figure 2). This runtime improvement makes rGkCM more practical for real-world applications.

We note the sensitivity of the azimuth–elevation consistency metric. Comparing Tables I and II, both uGkCM and rGkCM exhibit a marked decrease in true positive rate and an increase in false positive rate, likely due to degenerate configurations that cause the consistency checks to fail. The relative geometry of the robot trajectory and landmarks, as well as the sensitivity of the triangulation, should therefore be carefully considered before adopting the azimuth–elevation consistency metric for a given localization scenario.

B. Localization with Real-World Bearings

We validated our approach on an acoustic dataset collected by a WAM-V 8 autonomous surface vehicle (ASV) operating in a marina environment with water depths ranging from 2–2.5 m. The ASV was equipped with a SeaTrac X150 USBL acoustic modem, an inertial measurement unit (IMU), and GNSS for ground truth. Four additional SeaTrac X150

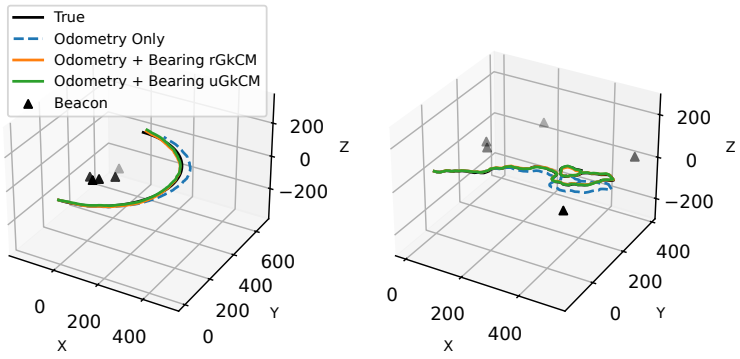

Fig. 1. Left: Arc trajectory with ideal triangulation geometry. Right: Random trajectory with degenerate triangulation geometry.

TABLE I

PERFORMANCE COMPARISON OF RGKCM AND UGKCM UNDER VARYING OUTLIER RATIOS (ARC TRAJECTORY).

Num Outliers	**Time** μ, σ **(s)**		**TPR**		**FPR**		**Translational RMSE** μ, σ **(m)**		**Rotational RMSE** μ, σ **(rad)**	
	rGkCM	uGkCM	rGkCM	uGkCM	rGkCM	uGkCM	rGkCM	uGkCM	rGkCM	uGkCM
50	**0.43, 0.11**	39.80, 15.51	**0.7431**	0.6684	0.0000	0.0000	5.26, 6.69	5.31, 4.11	0.03, 0.04	0.03, 0.02
60	**0.23, 0.09**	11.20, 4.16	0.6893	**0.7049**	0.0000	0.0000	5.14, 3.64	5.76, 4.37	0.03, 0.02	0.03, 0.02
70	**0.10, 0.04**	2.30, 0.87	0.7483	**0.7619**	0.0000	0.0000	5.28, 3.84	5.70, 4.21	0.03, 0.02	0.03, 0.02
80	**0.05, 0.02**	0.27, 0.09	0.8206	**0.8271**	0.0000	0.0000	7.41, 16.34	6.28, 5.29	0.03, 0.04	0.03, 0.02
90	0.03, 0.02	**0.02, 0.01**	0.7774	**0.9192**	0.0077	**0.0001**	17.40, 22.10	8.25, 8.30	0.07, 0.09	0.03, 0.04

TABLE II

PERFORMANCE COMPARISON OF RGKCM AND UGKCM UNDER VARYING OUTLIER RATIOS (RANDOM TRAJECTORY).

Num Outliers	**Time** μ, σ **(s)**		**TPR**		**FPR**		**Translational RMSE** μ, σ **(m)**		**Rotational RMSE** μ, σ **(rad)**	
	rGkCM	uGkCM	rGkCM	uGkCM	rGkCM	uGkCM	rGkCM	uGkCM	rGkCM	uGkCM
50	0.47, 0.17	13.57, 9.68	**0.5310**	0.5081	0.0001	**0.0000**	4.09, 2.99	4.00, 2.75	0.03, 0.02	0.03, 0.02
60	**0.27, 0.13**	4.37, 2.76	0.5406	**0.5540**	0.0003	0.0003	4.19, 3.61	3.86, 3.16	0.03, 0.02	0.03, 0.02
70	**0.13, 0.07**	1.02, 0.72	0.5899	**0.6029**	0.0002	**0.0001**	4.45, 4.84	3.95, 2.85	0.03, 0.04	0.03, 0.02
80	**0.06, 0.04**	0.18, 0.13	0.6715	**0.6785**	0.0007	0.0007	5.80, 6.35	5.32, 4.50	0.03, 0.04	0.03, 0.03
90	0.03, 0.02	**0.02, 0.02**	0.6412	**0.7780**	0.0109	**0.0032**	12.61, 9.40	11.58, 12.70	0.08, 0.07	0.06, 0.07

modems were suspended from surrounding docks to serve as beacons, approximately one meter above the seabed. Because the beacon depths were not measured, and given that depth is small relative to the vehicle's trajectory, we evaluate only azimuth angles for outlier rejection in this experiment.

The ASV sequentially queried the four beacons and received messages every 3 seconds. It received a total of 1021 measurements: 94, 326, 289, and 312 from each beacon, respectively. Even in this shallow environment, where multipath is severe, the modems provide relatively clean range estimates by using the time-of-flight of the first arrival path. Figure 3 shows the accuracy of range measurements without any postprocessing or rejection, confirming their reliability. In contrast, bearing estimates contained many significant outliers, especially at the start of the mission. These may be attributed to vehicle instability or corrupted signals caused by dock proximity. These measurement results highlight the importance of independent outlier rejection for range and bearing measurements when using acoustic sensors.

Because IMU alone is not sufficient to provide the odometry backbone needed for consistent set maximization, we made psuedo-odometry measurements by computing the relative pose between GPS measurements at each received acoustic signal, and composing the poses with randomly generated Gaussian noise. In an actual application, an additional sensor such as a Doppler Velocity Log (DVL) should be used to supplement the IMU and provide the between factors for the initial pose graph.

We construct a SLAM problem using a factor graph built in GTSAM [31]. The initial graph was constructed from a prior factor located at the origin and between factors from pseudo-odometry. For each beacon, we then constructed a consistency hypergraph using the proposed azimuth-elevation metric. The densest edge weighted clique was extracted using rGkCM, and the time elapsed to solve the hypergraphs for each of the four modems was 2.83, 238.21, 130.62, and 233.051 seconds. Due to its high computational cost, uGkCM could not be applied to this dataset. Because the location of the landmarks was unknown, two inlier bearing measurements from rGkCM and the associated poses were used to obtain an initial guess of each beacon location using triangulation. Finally, the range and inlier bearing measurements were incorporated into the factor graph for localization.

Figure 3 illustrates the effectiveness of our approach

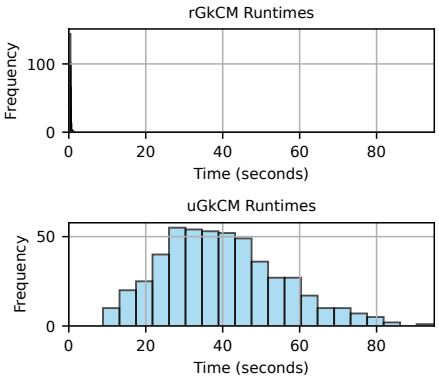

Fig. 2. Distribution of time elapsed to evaluate one hypergraph for an arc trajectory of 100 steps and 50 outliers. rGkCM consistently identifies a clique in under one second when using four threads to compute the gradient, while uGkCM takes up to 95 seconds when using four threads.

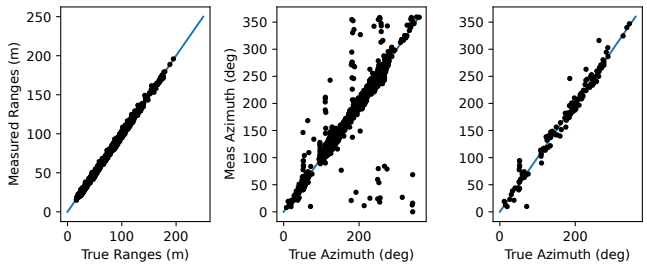

Fig. 3. Left: Range measurements have few outliers despite the shallow marine environment. Middle: Azimuth measurements are highly noisy with outliers. Right: Remaining azimuth measurements after performing outlier rejection. The blue represents the ideal *true = measured* line.

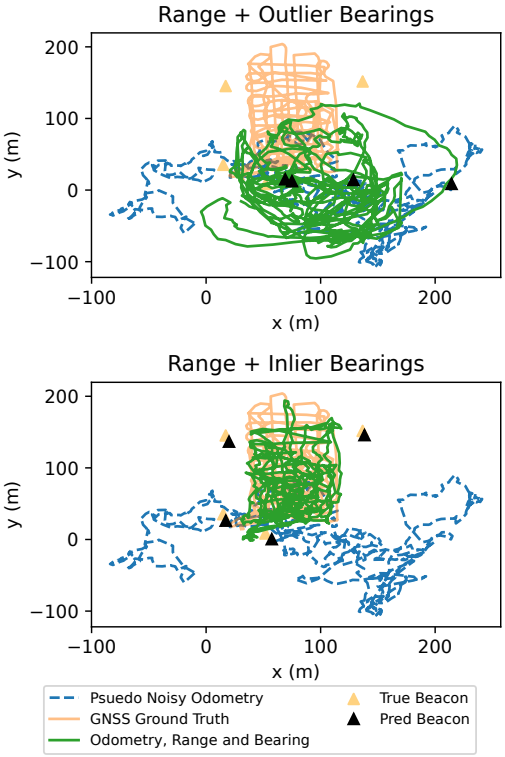

Fig. 4. Localization results using pseudo-noisy odometry from GNSS measurements, real-world acoustic bearing measurements, and triangulation from inlier bearings to initialize beacons.

in removing outlier bearings. As is common with GkCM methods, a portion of valid inliers were also discarded, reflecting the conservative nature of clique-based selection. Nevertheless, the remaining inliers were sufficient to provide an initial estimate of beacon locations, and discarded measurements could potentially be re-evaluated against these estimates, as suggested in [12]. Figure 4 shows the corrected localization estimates after incorporating only the inlier bearing measurements into the pose graph. The root mean squared error (RMSE) for the robot's x and y positions was 6.97 m and 12.83 m, respectively, while the RMSE for the beacons was 3.39 m and 7.94 m. These results demonstrate that in real-world conditions, our method improves robustness to bearing outliers while maintaining tractable runtimes, compared to the heuristic version of uGkCM.

VII. Conclusion

We have presented rGkCM, a replicator-dynamics-based method for outlier rejection using the weighted group-k consistent set maximization approach. By introducing a $k=3$ azimuth-elevation consistency metric, our approach extends consistent set maximization to low-degree-of-freedom bearing measurements, complementing existing range-based methods. Future work includes comparing replicator-dynamics with alternative continuous optimization approaches, such as Frank-Wolfe and infection-immunization dynamics, to further pursue real-time solutions for weighted GkCM. Additionally, hypergraph construction was outside the scope of this work but remains a significant computational bottleneck. Investigating more efficient techniques for building consistency hypergraphs could enable additional real-time GkCM applications.

Acknowledgment

ChatGPT models were used as a collaborative tool for code development of experiments and figures. All generated code was thoroughly edited by the authors.

References

[1] J. Mangelson, D. Dominic, R. Eustice, and R. Vasudevan, "Pairwise consistent measurement set maximization for robust multi-robot map merging," in *Proceedings of IEEE International Conference on Robotics and Automation*, May 2018.

[2] P. C. Lusk, K. Fathian, and J. P. How, "CLIPPER: A graph-theoretic framework for robust data association," in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

[3] B. Forsgren, R. Vasudevan, M. Kaess, T. W. McLain, and J. G. Mangelson, "Group-k consistent measurement set maximization for robust outlier detection," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.

[4] W. Gao, Y. Liu, B. Xu, and Y. Che, "An improved cooperative localization method for multiple autonomous underwater vehicles based on acoustic round-trip ranging," in *Proceedings of the IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 2014.

[5] M. V. Jakuba, J. C. Kinsey, J. W. Partan, and S. E. Webster, "Feasibility of low-power one-way travel-time inverted ultra-short baseline navigation," in *Proceedings of the MTS/IEEE OES OCEANS Conference*, 2015.

[6] N. R. Rypkema, E. M. Fischell, and H. Schmidt, "One-way travel-time inverted ultra-short baseline localization for low-cost autonomous underwater vehicles," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

[7] Y. Sekimori, H. Horimoto, Y. Noguchi, T. Matsuda, and T. Maki, "Scalable real-time global self-localization of multiple AUV system using azimuth, elevation, and depth difference acoustic positioning," in *Proceedings of the MTS/IEEE OES OCEANS Conference*, 2021.

[8] K. M. Fauske and O. Hallingstad, "A comparison of outlier detection algorithms for hydro-acoustic positioning," in *Proceedings of the MTS/IEEE OES OCEANS Conference*, 2006.

[9] D. Fenucci and A. Munafo, "Experimental evaluation of outliers filtering techniques in networked acoustic localisation systems," *IFAC-PapersOnLine*, vol. 53, pp. 14 582–14 588, 2020.

[10] J. Vaganay, J. Leonard, and J. Bellingham, "Outlier rejection for autonomous acoustic navigation," in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 1996.

[11] J. Su, S. Zou, J. Qian, Y. Wei, F. Qu, and L. Yang, "Rejecting outliers in 2D-3D point correspondences from 2D forward-looking sonar observations," *arXiv:2503.16066*, 2025.

[12] B. Forsgren, M. Kaess, R. Vasudevan, T. McLain, and J. Mangelson, "Group-k consistent measurement set maximization via maximum clique over k-uniform hypergraphs for robust multi-robot map merging," *International Journal of Robotics Research*, vol. 43, 07 2024.

[13] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.

[14] H. Do, S. Hong, and J. Kim, "Robust loop closure method for multi-robot map fusion by integration of consistency and data similarity," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5701–5708, 2020.

[15] Z. Chen, J. Zhao, T. Feng, C. Ye, and L. Xiong, "Robust loop closure selection based on inter-robot and intra-robot consistency for multi-robot map fusion," *Remote Sensing*, vol. 15, no. 11, 2023.

[16] M. Leordeanu and M. Hebert, "A spectral technique for correspondence problems using pairwise constraints," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, vol. 2, 2005.

[17] P. Lusk and J. How, "CLIPPER: Robust data association without an initial guess," *IEEE Robotics and Automation Letters*, pp. 1–8, 2024.

[18] M. B. Peterson, Y. X. Jia, Y. Tian, A. Thomas, and J. P. How, "ROMAN: Open-set object map alignment for robust view-invariant global localization," in *Robotics: Science and Systems (RSS)*, 2025.

[19] M. Pelillo and A. Torsello, "Payoff-monotonic game dynamics and the maximum clique problem," *Neural Computation*, vol. 18, no. 5, pp. 1215–1258, May 2006.

[20] H. Liu, L. Latecki, and S. Yan, "Revealing cluster structure of graph by path following replicator dynamic," March 2013.

[21] M. Frank and P. Wolfe, "An algorithm for quadratic programming," *Naval Research Logistics Quarterly*, vol. 3, pp. 95–110, 1956.

[22] M. Leordeanu and C. Sminchisescu, "Efficient hypergraph clustering," in *Proceedings of the International Conference on Artificial Intelligence and Statistics*, N. D. Lawrence and M. Girolami, Eds., vol. 22, 2012, pp. 676–684.

[23] S. Bulò and M. Pelillo, "A game-theoretic approach to hypergraph clustering," in *Advances in Neural Information Processing Systems*, Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta, Eds., vol. 22. Curran Associates, Inc., 2009.

[24] S. Rota Bulò and M. Pelillo, "A game-theoretic approach to hypergraph clustering," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 6, pp. 1312–1327, 2013.

[25] F. Ahmed and G. Still, "Two methods for the maximization of homogeneous polynomials over the simplex," *Computational Optimization and Applications*, vol. 80, Nov. 2021.

[26] S. Rota Bulò, M. Pelillo, and I. M. Bomze, "Graph-based quadratic optimization: A fast evolutionary approach," *Computer Vision and Image Understanding*, vol. 115, no. 7, pp. 984–995, 2011, Special issue on Graph-Based Representations in Computer Vision.

[27] E. Zemene, Y. T. Tesfaye, H. Idrees, A. Prati, M. Pelillo, and M. Shah, "Large-scale image geo-localization using dominant sets," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 1, p. 148–161, Jan. 2019.

[28] Y. bar shalom, X. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Theory, Algorithms and Software*, Jan. 2004.

[29] M. Kaess and F. Dellaert, "Covariance recovery from a square root information matrix for data association," *Robotics and Autonomous Systems*, vol. 57, pp. 1198–1210, 2009.

[30] L. E. Baum and J. A. Eagon, "An inequality with applications to statistical estimation for probabilistic functions of markov processes and to a model for ecology," *Bulletin of the American Mathematical Society*, vol. 73, pp. 360–363, 1967.

[31] F. Dellaert and Contributors, "GTSAM," May 2022.