A Range-Only Acoustic Localization Technique for Marine Robotics Applications

Jeff Kaeli, Gwyneth Packard, Carl Fiester, Dara Tebo, Mike Purcell
Ocean Systems Laboratory, Applied Ocean Physics and Engineering Department
Woods Hole Oceanographic Institution (WHOI)
Woods Hole, MA, USA
{jkaeli,gepackard,cfiester,dtebo,mpurcell}@whoi.edu

***Abstract*—We present a parametric technique for simultaneous outlier rejection and position estimation from a set of range measurements. This approach utilizes an adapted version of the Random Sample Consensus (RANSAC) algorithm, which has found extensive use in image matching applications. Our implementation has been demonstrated to function in real time on an Autonomous Underwater Vehicle (AUV) for the purposes of localizing a stationary acoustic beacon and aiding navigation.**

***Index Terms*—underwater acoustics, single beacon navigation, autonomous underwater vehicle, parametric modeling, outlier rejection**

I. Introduction

Acoustics plays a critical role in underwater communication, navigation, and localization due to the absence of GPS and the limitations of dead reckoning using inertial sensors. The most ubiquitous use is measuring ranges using acoustic travel time multiplied by the speed of sound. These ranges may be measured between a ship and an autonomous underwater vehicle (AUV) for tracking, between an AUV and a seafloor beacon for use as a navigation reference, or between a ship and a seafloor beacon for surveying the exact location of the beacon. In each case, the range measurements must be made from multiple locations so the solution can be triangulated. However, acoustic range data suffers from errors that are highly non-Gaussian and prone to outliers due to ray bending in the water column and multi-path propagation from surface and seafloor reflections. These outlier ranges are often difficult to reject in the absence of a priori knowledge about platform movement.

This problem has many similarities to the problem of estimating the transformation between two images using matched keypoints. Such an estimation can be used for the purpose of estimating camera motion, uncalibrated stereo image rectification, or constructing a photomosaic. The Random Sample Consensus (RANSAC) algorithm has enjoyed broad success in addressing these problems. In contrast to regressive methods that require outlier rejection as a pre-processing step, RANSAC simultaneously rejects outliers while estimating the parameters of the model. It works by first randomly selecting a number of samples to uniquely constrain a parameterized model. Then, it computes the number of inliers from the entire data set that also fit that model within a specified threshold. Lastly, it iterates upon this process while saving the parameters of the model that maximize the number of inliers.

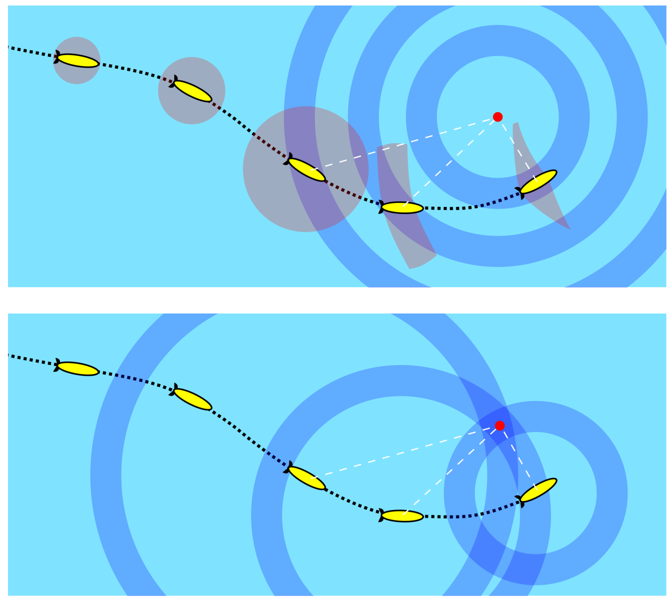

Fig. 1: *Top:* An AUV accumulates uncertainty as it maneuvers in the absence of external navigation fixes. By obtaining acoustic ranges to the beacon, the AUV's uncertainty is constrained by the uncertainty in the range. *Bottom:* Over short distances, the AUV position can be assumed to be perfect for the purpose of localizing the beacon. In both cases, it is readily apparent that orthogonal bearings provide the optimal constraints for single beacon navigation.

We begin with a discussion of prevailing underwater navigation techniques, in particular range-only techniques, and describe the advantages our approach has over existing methods. We then present modifications to the RANSAC algorithm that make it suitable for outlier rejection and position estimation based on acoustic ranges. We also present a refinement step that improves the position estimate using all ranges that are inliers to the estimated model parameters. We then discuss the process of implementation on an AUV for the purpose of estimating the location of and navigating relative to an acoustic beacon on the seafloor.

II. Related Work

Once an AUV dives below the surface, it loses its connection with GPS satellites and must estimate its position using a combination of on-board inertial sensors and external acoustic references. In typical cases, an acoustic beacon replies to an interrogation ping from the AUV, and the range can be inferred from the two-way travel time based on the local speed of sound. When multiple beacons respond to the same ping, termed long-baseline (LBL) navigation, the AUV is able to to triangulate its position based on a single interrogation [1]. However, based on the added time and cost to deploy and survey in each acoustic beacon, there is interest in ways that a single beacon can be utilized to aid navigation.

A. Single Beacon Navigation

Ultra-short baseline (USBL) arrays can be used to provide bearing information in addition to range from a single ping. These have proven effective for AUV docking [2] as well as tracking applications [3]. A more exotic beacon has been demonstrated that generates a spiral wave front such that bearing can be determined based on frequency [4]. A fair amount of work has also been done to improve upon the interrogate-reply paradigm used in LBL navigation. Chip-scale atomic clocks (CSAC) can be utilized to provide one-way-travel time from synchronized pings. These have been demonstrated both where the USBL array is on the vehicle [5] as well as where the USBL array is on the ship and the position solution can be transmitted to the AUV [6]. This latter approach has also been implemented on autonomous surface platforms [7].

B. Range-Only Navigation

In many cases, range information alone is ample to aid navigation. Multiple pings made relative to a moving platform create a synthetic array over time, enabling triangulation similar to LBL navigation. However, consideration of geometry is critical. Erroneous mirror solutions can arise when the AUV moves straight without turning. Furthermore, the beacon only constrains the error along the bearing to the beacon, as shown in Figure 1, so orthogonal bearings provide the optimal constraints on the uncertainty. As a result, much of the work in range-only navigation involves pre-defined maneuvers for areal coverage [8] or for docking applications [9]. These approaches have utilized both two-way [10] and one-way travel times [11].

C. Outlier Rejection

An important consideration for position estimates using acoustic ranges is rejecting bad ranges, or outliers. Acoustic ranges are subject to noise that is highly non-Gaussian, so linear methods such as regression or simple averaging over time will not improve results. These bad ranges sometimes arise where the ping bounces off the surface of the seafloor. When the AUV detects this multi-path ping rather than the direct path, the range can be overestimated. Ray bending in stratified ocean layers exacerbates this. Bad ranges can also arise from false detections on background noise in the ocean. Lastly, sometimes there is no detection, creating gaps in the time series. Methods of outlier rejection have included manual rejection by trained operator, [11], particle filters [5], [6], and single cluster graph partitioning [12].

Another method that simultaneously estimates parameters while rejecting outliers is the Random Sample Consensus (RANSAC) algorithm [13]. It has found extensive use in automated image matching because it is robust to the large amounts of non-Gaussian noise inherent in computing image transforms using matched features. [14]. In the context of underwater acoustics, it has been used as a pre-processing step for bearing-only estimation [15], [16].

III. Methods

A. Random Sample Consensus

The goal of the RANSAC algorithm is to find a solution S given a set of N records $R = \{R_1; R_2; :::; R_N\}$ that maximizes the number of records that contribute to the solution (the inliers) while rejecting the records that do not contribute to the solution (the outliers). A threshold h is specified, within which a record will be considered an inlier to the solution. First, a subset of K records is chosen at random, where K is the minimum number of records required to uniquely constrain the solution. Then, the solution to this subset is found, and each record is tested to see if it is an inlier to that solution. If the current solution has a higher fraction of inliers f than the existing solution, then the current solution is saved as the best solution thus far. This process continues for Q iterations, determined from the probability p that a solution with inlier fraction f can be found by randomly sampling K records Q times. [17].

$$Q \geq \frac{\log(1-p)}{\log(1-f^K)} \tag{1}$$

Each record $R_i = [x_i; y_i; r_i]$ is a measurement of AUV position, computed in a local latitude/longitude reference frame, and the corresponding range. Two records produce 0, 1, or 2 solutions $S = [x_s; y_s]$ based on the intersections between the range circles centered at the AUV positions. A third record is therefore required to disambiguate the solution, so $K = 3$. Computing the intersections between two circles is a straightforward mathematical operation. We maintain a circular buffer of the N most recent records to avoid degraded performance from old records with higher uncertainty relative to the current record.

B. Refinement

At this point, the adapted RANSAC algorithm has provided a solution that fits 2 records exactly, with the remaining inlier records lying within a threshold distance. For certain applications, it may be desirable to refine the solution to minimize the error within the entire set of inlier records. For each inlier record, we can compute the distance between the record and the solution.

Algorithm 1 RANSAC

input:
  set of N records $R_1; R_2; :::; R_N$
  inlier threshold h
output:
  best solution B
initialize:
  iteration count $i \leftarrow 0$
  required iterations $Q \leftarrow \infty$
  inlier fraction $f \leftarrow 0$
while $i < Q$ **do**
  $i \leftarrow i + 1$
  randomly select K records $R_{n_1}; R_{n_2}; :::; R_{n_K}$
  compute solution S using K records
  inlier count $c \leftarrow 0$
  for $j = 1$ to N **do**
    compute error e between R_j and S
    if $e < h$ **then**
      $c \leftarrow c + 1$
  if $\frac{c}{N} < f$ **then**
    $B \leftarrow S$
    $f \leftarrow \frac{c}{N}$
    $Q \leftarrow$ updated required iterations
return B

$$d_i = \sqrt{(x_i - x_s)^2 + (y_i - y_s)^2} \quad (2)$$

The difference between this distance and the range is the error between the record and the solution.

$$e_i = d_i - r_i \quad (3)$$

By scaling the difference between the record and the solution along each axis by the ratio of the error to the distance, we can compute a vector pointing from the solution to the edge of the range circle. This error vector represents the amount the solution would have to move to reduce the error to zero for that record.

$$x_{e;i} = (x_i - x_s)\frac{e_i}{d_i} \quad (4)$$

$$y_{e;i} = (y_i - y_s)\frac{e_i}{d_i} \quad (5)$$

We can compute the mean of these error vectors over the set of inliers $i \in I$.

$$x_e = \sum_{i \in \boldsymbol{I}} w_i e_x \quad (6)$$

$$y_e = \sum_{i \in \boldsymbol{I}} w_i e_y \quad (7)$$

The weight w_i can be used to favor certain records over others, for instance, to trust more recent records where the AUV's position estimation is more reliable. The weights must sum to unity. For the results presented here we have used uniform weighting.

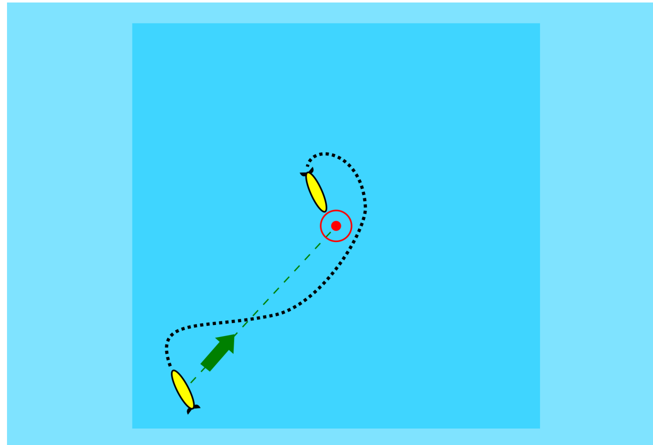

Fig. 2: Diagram of the simulation setup. The direct path is illustrated by the green dotted line.

The solution can then be updated by augmenting the current solution by the mean error vector.

$$x_{s;\text{updated}} = x_s + x_e \quad (8)$$

$$y_{s;\text{updated}} = y_s + y_e \quad (9)$$

This process can be iterated until the magnitude of the mean error vector drops below some threshold.

IV. Results

A. Simulation

We developed a simulation environment in MATLAB to test both the estimation algorithm as well as the behaviors acting on these estimates. Since these estimates subsequently affect the next record that is observed, such simulations help us to analyze many different behaviors before testing in the water. To realistically model the acoustic ranges between the AUV and the beacon, several kinds of noise are added to the actual values. First, Gaussian noise is added to each computed range. Next, there is a 10% likelihood that uniformly distributed noise is added to the range, simulating a multi-path return. Then, there is a 10% likelihood that a completely random range is returned, simulating a detection off of ambient noise in the water column. Lastly, there is a 10% likelihood that no range is returned, simulating no detection of the reply ping. These values have been tuned heuristically to provide acoustic range variations similar to those encountered locally.

The simulation setup is illustrated in Figure 2. We started with the simplest possible behavior where the vehicle always bears towards the estimated beacon location at some predefined speed. To model realistic vehicle dynamics, the maximum heading rate of the vehicle was limited to 10 degrees per second. The AUV began the simulation at a random position and with a random heading somewhere in a box of 1000m radius around the beacon. As it estimated the beacon's

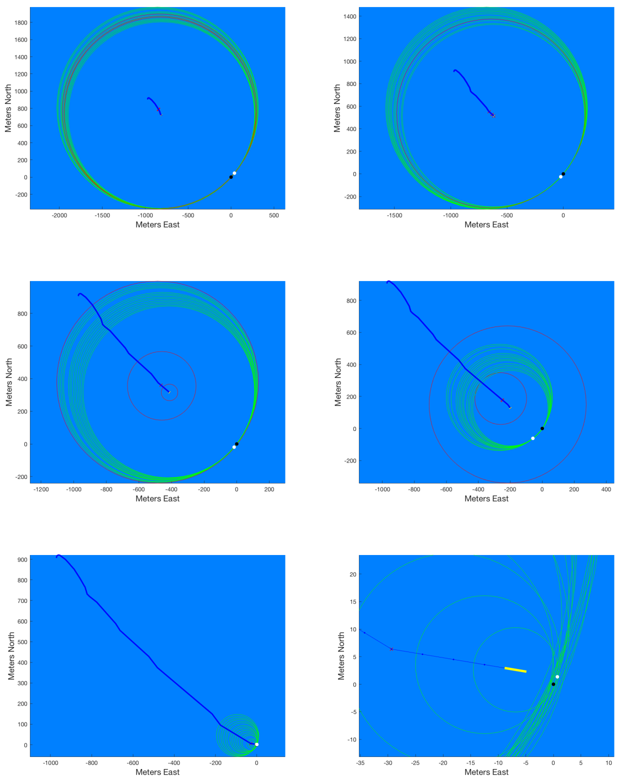

Fig. 3: *From top to bottom, left to right:* snapshots from one simulation run. Outlier ranges are shown as red circles, while inlier ranges that contribute to the location estimate, shown as a white dot, are green. The actual beacon location is shown as a black dot.

location, it altered its bearing directly towards the estimate. Its speed was a constant 1:9 meters per second, a typical value at the higher end of AUV maneuverability. When the AUV came within a 10 meter threshold of the beacon, the current iteration was ended and the next iteration began.

Figures 3 illustrates several snapshots from one simulation run. To evaluate the performance of the algorithm and the behavior, the length of the actual path taken by the vehicle is normalized by the length of the direct path. This is the shortest path in a straight line that the AUV would have to travel to reach the beacon, less the 10 meter threshold. A Monte Carlo approach of 5,000 simulation runs was computed, and the normalized distance calculated for each one. A histogram of these distances is shown in Figure 4. Over 90% of the runs resulted in vehicle paths that were less than 10% greater than the direct path.

B. Relative Loitering

An adaptive vehicle behavior was developed as a REMUS objective that enables the AUV to loiter near seafloor node equipped with an acoustic beacon [18]. The motivation here is twofold. First, the AUV is used to offload data from an ocean bottom seismometer (OBS) using the high throughput of an

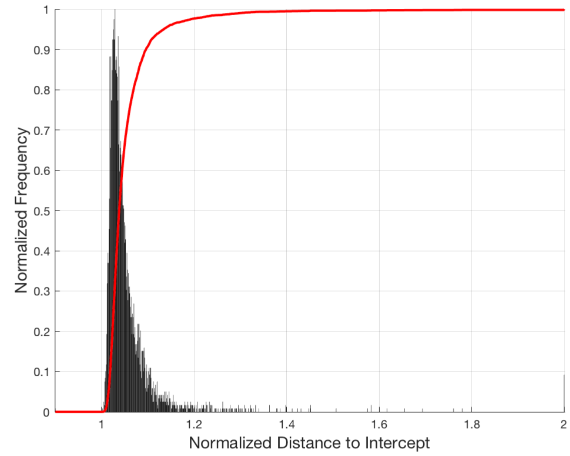

Fig. 4: Histogram (*black*) and cumulative distribution (*red line*) of distance ratios accumulated over 5,000 iterations of a Monte Carlo simulation.

optical modem at short ranges [19]. Second, the optical modem also permits the AUV to transmit an accurate time, enabling the OBS to reset its time base which can drift nonlinearly over long deployments [20]. The AUV continuously estimates the relative position of the beacon after each ping while attempting to circle at a specified radius. This enables it to account both for the initial error in the OBS location as well as its own accumulated uncertainty as it maneuvers. Furthermore, this circling behavior keeps the AUV and the node within each others' optical beam pattern, ensuring the the link remains continuously active throughout the loiter.

Figure 5 shows vehicle interface program (VIP) playback from one of several test missions in Buzzards Bay, MA. The AUV begins to circle the location where it has been told the beacon is prior to its mission. As more records are collected, the position estimate shifts to the actual location. Circling a beacon is the optimal geometry for single transponder, range only navigation because it continuously provides orthogonal bearings at the fastest possible update rate. For this experiment, the inlier threshold was set to $h = 2$ meters and the confidence was set to $p = 0{:}9999$. Through experimentation we have found that using $N = 20$ records provides a good balance between reactiveness and a stable solution. The AUV was able to circle at a radius of 10 meters and maintain the optical link while continuing to update the position estimate.

V. Discussion

We have successfully demonstrated the ability of an AUV to estimate the location of a stationary seafloor beacon and use this estimate to inform adaptive behaviors. This is shown in simulation for transiting towards a beacon as well as in practice for circling a beacon where the application is establishing and maintaining an optical link with a seafloor node. Since geometry is vitally important to single beacon, range only position estimation, the circling behavior is ideal