

RESOLVING OBSERVER MOTION BY OBJECT TRACKING.

John Hallam
Department of Artificial Intelligence
Edinburgh University
Edinburgh

ABSTRACT.

A fundamental sub-problem of the task of interpreting the sonar data gathered by a mobile submersible robot vehicle is that of constructing and maintaining a stationary global reference frame in which interpretation can take place. Such a frame may be established by tracking the movements of objects in the local environment of the robot and using the data so obtained to deduce the motion of the observer. In a marine context, where currents make it impossible to hold a vehicle stationary for any extended period, it is important to track the vehicle and object motion continuously in real time. In this paper the problem of establishing a stationary reference frame using noisy, specular, intermittent sonar data is discussed and a computational method for its solution is presented. The method has been tested by computer simulation in a two dimensional context and an evaluation of its performance is given. The solution presented is not restricted to marine vehicles but is especially effective for land-based vehicles, where operating conditions are less severe.

I. INTRODUCTION.

A. Background.

The problem of motion resolution affects all mobile robotic vehicles to some degree. In essence it is the problem of determining a robot's proper motion or position with respect to its 'world'. This task has been well defined and tackled for land-based robots where some simplifying assumptions may be made. The solution is exemplified in Moravec's work with the Stanford CART robot [Mor80] and takes the form of a 'Move, Stop, Think' paradigm. The vehicle estimates its current position, moves a roughly known distance in an approximate direction, and then remains stationary while its position is recalibrated using sensory data (in Moravec's work, nine television images supply this data). When recalibration is complete, it is ready to execute another planned movement. This procedure is a generalised version of the static calibration that fixed robots undergo before beginning work. Indeed, if the planned movements can be made or measured accurately the recalibration with its time-consuming processing need not be done very

often, and the vehicle proceeds by dead-reckoning. Such a paradigm is both simple and effective provided that it is possible approximately to predict or estimate the effects of the deliberate movements and, more importantly, that it is possible to remain stationary while the position recalibration is done.

These assumptions are reasonable on land, but immediately break down in a marine environment. It is impossible to keep a vehicle stationary for any extended period since water currents will push it about, and without expensive navigational equipment it is impossible to keep track of the vehicle position accurately. If motion resolution is to be done in this context it must be performed continuously, in real time, rather than using a discrete paradigm like the land-based one.

B. System Input.

This constraint on a motion resolution scheme allows some of the potential difficulties of marine acoustic sensing to be turned to advantage. Acoustic data, unlike optical data, is strongly specular and it is difficult to recover anything resembling an optical image in bulk spatial resolution without extensive processing and using special techniques to overcome the specularity. Sonar images naturally contain relatively few data points, but it is easy to obtain accurate range and bearing (in two or three dimensions) directly from the data and most sonar equipment is capable of measuring Doppler shift velocity for each data point. This highly concentrated geometric information is ideal for driving a motion resolution system.

With these facts in mind, therefore, the input to the motion resolution system is assumed to comprise time series of vectors of measurements describing the data points (targets) in a sequence of sonar images. Each vector contains the polar coordinates of its target with respect to the sonar transducer, and may carry a Doppler velocity if one is available and reliable. The vector also carries a flag that informs the system whether the target has a non-zero absolute velocity -- this can be computed from Doppler data even when the Doppler velocity cannot be determined exactly.

The number of targets seen in a given sonar scan is variable and because of specularity and the relative movement of observer and target a given target will in general be seen only intermittently. The geometric measurements obtained for each target, in polar coordinates, are subject to errors of constant variance. However, when transformed into Cartesian coordinates the errors are range dependent -- the volume of uncertainty in which the target lies is a sector with width dependent on the sonar beam-width and thickness dependent on the sonar range resolution (see Figure 1). The errors are independent in polar coordinates but are coupled, depending on the bearing, in Cartesian coordinates.




Figure 1: Measurement Uncertainty Region

II. A METHOD FOR MOTION RESOLUTION.

An alternative and more uniform definition of motion resolution may be formulated in terms of target tracking. If the proper motions of all targets, including the sonar transducer as a special case target that is never seen, are tracked a global frame may be established using the tracking data. Implicitly it is the frame with respect to which the transducer is tracked; explicitly, it is that frame in which stationary targets are fixed. In fact, although it is known whether a target is absolutely stationary, only the relative motion of the target with respect to the observer may be measured; and further, a stationary target does not imply a fixed absolute position (consider the stationary line reflector in Figure 2, which mirrors observer motion parallel to the line). This motivates the fundamental insight of the method: if the relative motion of targets is tracked the observer movement will appear as a common component of those relative motions, that may be extracted and tracked to derive the observer motion.

The basic block of the tracking system is the Kalman filter. Devised by R. E. Kalman [Kal60], this filter is a general and powerful recursive data processor. It is able to deal with input data vectors with variable error covariance and variable inter-sample times. For data with a unimodal error

distribution with zero mean the Kalman filter is optimal in the sense that its estimates and predictions of tracked quantities have minimal error variance, and the filter computes error covariance matrices for its outputs in the course of its operation. The filter has been used extensively in radar tracking applications (for example, Morley and Wilsdon's multiradar tracking system [Mor77]) though in these cases there is no observer motion to extract. It is computationally demanding in general, but can be simplified considerably in particular cases.




Figure 2: Behaviour of a Line Reflector.

The motion resolution algorithm may be divided into two parts: a set of per-target channels (Figure 3); and a kernel (Figure 4). The kernel tracks the observer, which has three degrees of freedom -- two translational and one rotational. Translational tracking maintains the global frame origin position while rotational tracking compensates for fluctuations in the vehicle's heading which shift the sonar transducer's polar origin. The kernel receives inputs from all active channels and may incorporate navigational data from the vehicle's control system if available.

Each channel tracks a given target and is activated when that target is seen. Targets, for the purpose of tracking, are considered to be points and so have two degrees of freedom. The channel contains both relative and absolute motion data for its target and mechanisms to pass certain data to the kernel. With this organisation new targets are accommodated by creating a new channel, while defunct targets are eliminated by destroying the associated channel.

Tracking is done in Cartesian coordinates. The kernel tracking is done with respect to the global reference frame; channels track their targets using coordinates in the observer relative frame. This latter frame is defined as the Cartesian reference frame with origin at the current position of the sonar transducer and with x axis aligned with the azimuth axis of the transducer.

The algorithm is cyclic in its operation; each cycle is initiated by target observation vectors (event vectors) arriving at their appropriate channels (the implied segmentation problem is discussed later). With reference to Figure 3, an arriving event vector is transformed into Cartesian coordinates in the observer relative frame (ORF), an error covariance matrix is computed for it, and it is used as input to update the target relative state (TRS) vector using the Kalman filter K_r . This state is a four component vector of position and velocity in the ORF. The TRS is rotated so that the ORF and global frame polar axes are aligned, and three quantities computed from this transformed TRS (TTRS) provide kernel inputs.

i). the target absolute position less the TTRS position is used as an observer position estimate;
 ii). the TTRS velocity, adjusted for the current ORF angular velocity estimate (and for target absolute motion), is used as an observer linear velocity estimate.

iii). the TTRS velocity is corrected using the current observer absolute velocity, and for target absolute motion if necessary, and used to compute an estimate of the observer (or ORF) angular velocity. These three estimates drive the kernel translational and rotational tracking.





Figure 4: Kernel Processing.

The next phase of activity takes place in the kernel (Figure 4). Inputs from all active channels are collated, and combined using a non-recursive minimum variance averaging filter. This filter is designed to have unit common mode gain while rejecting uncorrelated differential inputs as much as possible, and is defined by the equations below:

$$\bar{P} = \{\sum_i P_i^{-1}\}^{-1}; \quad \bar{x} = \bar{P} \sum_i P_i^{-1} x_i$$

where P_i is the variance of x_i

Figure 3 (below): Per-Target Channel Organisation.
 The dashed link is used only for moving targets.



The position and velocity data are tracked by the Kalman filter, K_0 , which maintains the four component observer absolute state (OAS) vector. The angular velocity average is used directly as the new ORF angular velocity and the polar axis alignment is computed by integrating the angular velocity estimate.

In the third phase of each cycle the kernel passes the new estimate of the OAS to each active channel which adds the OAS to the TRRS to compute a new estimate of its absolute state vector. Each estimated quantity passed through the system is accompanied by an appropriate variance or covariance.

III. INPUT SEGMENTATION.

As noted above, there is an implicit segmentation problem in matching target observation events to the appropriate channels. This is dealt with by using the prediction ability of the Kalman filters in the channels. For each arriving event the set of potentially matching channels is enumerated and a predicted TRS is computed for each channel in the set. The potential match set is then pruned as follows. First, estimates of the radial and angular distances between the event and TRS are made and compared, allowing for an unknown observer rotation, with estimates of the corresponding variances. Matches with distances that are too large are rejected at once. Secondly, the filter innovation distance (the distance between the TRS and event using the filter covariance matrix as a metric) is computed for each match. The matches are classified into 'good' if the distance is less than 3.5 standard deviations, or 'poor' if less than nine. Matches missing the 'poor' threshold are discarded. Thirdly, poor members of multiple match classes that contain a good match are eliminated. Lastly, the remaining matches are implemented: uniquely matched events are passed to their channels; unmatched events initialise new channels; multiple good matches of events are dealt with by a multiple target correlation scheme (developed by Fortmann et al. [For80], and to be implemented shortly); multiple poor matches are deferred pending further data.

A problem that particularly affects segmentation arises from allowing the ORF to rotate. The rotation of the ORF appears as common angular velocity in the channel relative tracking filters. However, because these filters have a linear prediction model, an error arises in the TRS due solely to polar origin shift. This is illustrated in figure 5. If the filter state is not estimated for a time t , the state will be predicted at point P , $r(1, wt)$ where w is the angular velocity. However, the actual position of the target is point T , $r(\cos wt, \sin wt)$. The prediction error is roughly proportional to $r(wt)^2$, and must be corrected if this is large. This error also affects the angular velocity estimates computed by the target channels, since the position estimate in each TRS will be displaced radially outwards relative to the correct position. The error is proportional to the relative distance between observer and target, and

correction is made for it in the angular velocity estimator.




Figure 5: Rotation-induced Prediction Error.

IV. RESULTS.

The method has been tested using simulated events generated by a modeller program. The modeller reads initial values for target and observer position and velocity, and predicts their positions at future times. Each simulated sighting of a target causes an event vector to be generated. The modeller can move targets and observer deterministically, randomly, or in a combination of both modes of movement, and adds measurement error into the computed vectors if desired. The random movement and polar measurement errors have Gaussian distribution.

Two sets of tests were run. In the first set the angular compensation part of the motion resolution system was turned off and the observer restricted to translational motion only. The test world contained four objects positioned at the corners of a 70m sided square, centred on the origin, and the sonar scan repetition rate was 2.5Hz, corresponding to a range limit of 300m. Simulations ran from one hundred to one thousand seconds of simulated time. Extensive error data were logged during the runs and the results of four of the tests are reproduced in tables one and two. In the second set, the angular compensation system was turned on, and tests run for about thirty seconds simulated time, using the same world model as the first set of tests. The sonar scan rate remained 2.5Hz and the angular measurement standard deviation was 0.01 radians (about 0.57 degrees). In each test the radial measurement error standard deviation was 0.01 metres (corresponding to a sonar range resolution of about 40cm).

Table 1 gives the error variance ranges for the four runs in series one. These figures were computed by comparing the filter states with the modeller states after each cycle of the algorithm, accumulating the errors in the input measurements, the TRS position, the target global state for each

TABLE 1: Variances for Tests 1 (Maximum / Minimum).

TEST	CONDITIONS.	Measurement Variance	TRS Position Variance	Global Position Variance	Global Velocity Variance	OAS Position Variance	OAS Velocity Variance
1a	Observer Velocity (0.06, 0.08)	12	1.3	1.8	0.2	0.45	0.11
	Objects Stationary. Time : 100 seconds. θ standard deviation 0.1rad	10	0.7	0.5	0.04	0.4	0.08
1b	Observer Velocity (0.06, 0.08)	0.17	0.019	0.128	7.3e-3	0.110	7.1e-3
	Two Objects moving: (0.35, 0.2) (0, 0.6) Time : 100 seconds θ standard deviation 0.01 rad	0.020	0.003	0.040	1.8e-3	0.040	1.8e-3
1c	Observer Velocity (0.06, 0.08) + random	0.450	0.040	0.120	3.5e-5	3	1.6e-4
	Objects stationary Time: 200 seconds θ standard deviation 0.01 rad	0.002	0.004	0.012	1.0e-5	1.8	1.3e-4
1d	As 1c, with Time 1000 seconds	0.180	0.047	0.156	7.3e-4	0.600	3.2e-5
	Random position force 0.04m in each dimension for tests 1c, 1d	0.140	0.010	0.054	3.2e-6	0.210	2.7e-5

TABLE 2: Position Offsets for Tests 1.

TEST	Observer		Target 1		Target 2		Target 3		Target 4	
	x	y	x	y	x	y	x	y	x	y
1a	3.3	-0.6	2.8	-0.1	4.3	-0.2	2.7	-1.1	4.1	-0.3
1b	1.2	1.9	1.1	1.9	1.3	1.9	1.3	2.0	1.2	2.0
1c	30.0	39.8	29.9	39.9	30.0	39.8	30.0	39.7	30.0	39.9
1d	30.6	40.2	30.5	40.1	30.6	39.9	30.3	39.9	30.6	40.0

target and the observer global state. The variances presented are the maximum and minimum in each category. Note that they are not the data variances computed by the Kalman filters -- the filters were uniformly more pessimistic about the validity of their estimates.

Table 2 shows the coordinate origin offset vectors for the tests. These values are difference between the target and observer absolute positions and the corresponding data in the modeller states. The sizes of these numbers depends upon the difference between the modeller's initial position for the observer and the motion resolution system's starting observer position of (0,0), and on the amount of drift of the global frame origin which occurs during the initialisation period of the filter system. The stability of the global frame of reference in each test is indicated by the similarity between corresponding displacements for different targets and observer. Note that test 1d is a lengthier run of test 1c, so that comparison between these rows of Table 2 shows the time stability of the global frame; comparison within the rows shows the spatial stability of the frame.

Table 3 shows the observer absolute state variances and the biases in the angular state components for each of the tests in series two. These figures are the means and variances of the differences between system computed quantities and modeller generated quantities, as in Table 1. Table 4 shows the coordinate origin offset vectors for tests 2a and 2b. Results of this nature for tests 2c and 2d are not meaningful because of the large values of the polar axis bias in those tests. The values in Tables 3 and 4 are inherently less reliable than those in Tables 1 and 2 because of the smaller number of samples involved (some seventy as against two hundred and forty).

V CONCLUSIONS.

1). The system performs well when presented with data containing only translational motion. It rejects object motion satisfactorily and maintains a stable frame of reference under extreme input error conditions -- the standard deviation of 0.1 radians in test 1a corresponds to a sonar beam width of more than ten degrees. The reference frame is stable over long periods of time, its origin drifting about 60cm in 800 seconds, and the movement of the observer with respect to the reference frame is tracked correctly. Thus under these conditions the system fulfils its goal of implicitly defining a stable stationary frame of reference.

2). When presented with data containing rotational error the system performs poorly. The cause of failure is the consistent low estimate of the ORF angular velocity (low by 3.8 milliradians per second in tests 2c,d) which integrates over the test period to account for half of the bias in the polar axis offset (filter initialisation time accounts for the other half). This gradual drift in the polar axis offset causes the global frame to rotate and interferes severely with absolute

position and velocity extraction (thus there are no offset results for these tests). The error appears to be independent of the observer linear motion since it is seen identically in the results of tests 2c and 2d. Further testing will show whether the angular velocity deficit is due to a biased estimator or is a further rotational non-linearity error of the type already described.

3). It requires between five and ten activations of a given channel to initialise the relative tracking filter successfully. Since the global reference frame can be established using only four targets the system will cope with intermittent sightings of targets easily. For a sonar scan rate of 2.5Hz it may take up to about thirty seconds to collect the necessary observations (assuming a very pessimistic rate of appearance for the targets); with more sophisticated sonar systems it will take only a few seconds.

4). While the system is currently inadequate for marine vehicles, where angular effects cannot be eliminated easily, it is suitable for use in land based vehicles. The structure of the system allows it to incorporate position or velocity data from other sources into the kernel averaging process. Thus mechanically obtained estimates of the motion of a robot vehicle could be used to improve the settling time of the global frame and increase the stability of the observer tracking further. The long term and spatial stability of the reference frame make it useful for navigation. The modifications necessary to enable the system to cope with marine conditions will be done in further work when the cause of the low angular velocity estimate has been fully investigated.

REFERENCES.

- For80 Fortmann TE, Bar-Shalom Y, Scheffe M, 1980, 'Multitarget tracking using joint probabilistic data association'; IEEE Decision & Control 1980, 2, p807.
- Kal60 Kalman RE, 1960, 'A new approach to linear filtering and prediction problems'; Trans ASME Journal of Basic Engineering March 1960 pp35-45.
- Mor80 Moravec HP, 1980, 'Obstacle avoidance and navigation in the real world by a seeing robot rover'; Stanford Artificial Intelligence Laboratory Memo AIM-340.
- Mor77 Morley AR & Wilsdon AS, 1977, 'Multiradar tracking in a multisite environment'; RADAR-77 IEE International Conference, p66.

TABLE 3: Biasses and Variances for Tests 2.

TEST	CONDITIONS	Polar Axis Offset Bias	Bias Variance	ORF angular velocity bias	Bias Variance	Observer Position Var.	Observer Velocity Var.
2a	Observer Stationary Objects Stationary Θ standard deviation 0.01 rad Time : 30 seconds	1.36e-7	1.35e-6	5.87e-5	2.27e-7	5.00e-5	1.10e-5
2b	Observer Moving (0.08, 0.06) Objects Stationary Θ standard deviation 0.01 rad Time : 30 seconds	-2.06e-5	1.39e-6	5.75e-5	2.27e-7	4.24e-5	9.91e-6
2c	Observer rotating 0.05 rad/s Objects stationary Θ standard deviation 0.01 rad Time : 30 seconds	-0.200	7.09e-4	-3.80e-3	9.11e-6	1.18e-3	8.74e-5
2d	Observer Moving (0.08, 0.06) and rotating 0.05 rad/s Objects stationary. Θ standard deviation 0.01 rad Time : 30 seconds	-0.200	7.25e-4	-3.84e-3	8.85e-6	1.22	3.97e-3

TABLE 4: Linear Offsets for Tests 2a, 2b. (Unit is 10^{-2} metres)

TEST	Observer		Target 1		Target 2		Target 3		Target 4	
	x	y	x	y	x	y	x	y	x	y
2a	104	-44.2	55.4	41.9	15.1	-11.2	-82.5	-102	-30.9	-24.0
2b	-198	-156	-136	-102	-176	-159	-274	-245	-162	-171