Improving Robustness of Terrain-Relative Navigation for AUVs in Regions with Flat Terrain

Shandor Dektor
Department of Aeronautics and Astronautics
Stanford University
Stanford, CA 94305
Email: sgd@stanford.edu

Stephen Rock
Department of Aeronautics and Astronautics
Stanford University
Stanford, CA 94305
Email: rock@stanford.edu

***Abstract*—Terrain-Relative Navigation (TRN) is an emerging technique for localizing a vehicle in underwater environments. TRN offers a means of augmenting an INS/DVL dead-reckoned solution with continuous position fixes based on correlations with pre-stored bathymetry maps instead of surfacing for periodic GPS fixes. TRN accuracy on the order of 3m has been demonstrated in recent field trials using MBARIs Dorado-class AUVs in Monterey Bay. However, these TRN algorithms have occasionally converged to incorrect solutions when the AUV operates for extended times over featureless terrain. Specifically, the TRN filter can become overconfident in an incorrect position fix.**

This paper demonstrates that the cause of these false fixes in information-poor regions is an incorrect accounting of map uncertainty and sensor noise in standard TRN filter implementations, and offers a modification to the algorithms that can eliminate the false-fixes. Specifically, standard TRN algorithms assume that map noise and vehicle sensor noise can be lumped together when performing measurement updates. In regions where the ratio of terrain information to map error is low, this assumption leads to underestimation of position uncertainty as the filter essentially converges on noise in the map. Adjusting the filter variance to depend on the estimated terrain information in addition to map error and sensor error provides a more robust TRN solution.

An improved algorithm is described which adjusts the filter variance using a technique employed by the robotics and statistics community for reducing the likelihood of overconfidence. The advantage of this adjusted variance technique is that it permits nominal convergence rates of the TRN filter over information-rich terrain while mitigating the risk of false fixes in information-poor terrain.

The effectiveness of the modified TRN algorithm is demonstrated in simulations using field data from MBARI AUV runs over flat terrain in Monterey Bay.

I. Introduction

Terrain Relative Navigation (TRN) is an emerging technique for localizing a vehicle in the underwater environment. TRN offers a means to eliminate the drift accrued from dead reckoning by matching range measurements of the terrain to a stored terrain map. The TRN concept was initially developed for use on aerial platforms and first deployed as the cruise missile guidance system TERCOM (Terrain Contour Matching), which performed a batch correlation of altimeter measurements against an a priori map to generate a position estimate [1].

TRN has since been expanded to many other environments and has found particular value in the underwater environment. Map relative position estimates reduce the necessity of surfacing for GPS fixes or deploying USBL or LBL positioning systems. TRN accuracy on the order of a few meters has been demonstrated in field trials using MBARIs Dorado-class AUVs in Monterey Bay. Similar performace was observed in field trials using the HUGIN AUV in Ostfjords [2].

TRN occasionally converges to overconfident, incorrect solutions when operating for an extended period of time over featureless terrain [3]. These false fixes can occur in many forms - from simple offsets to divergence in the filter estimate, depending on the type of terrain and the route the vehicle is following. False fixes are not limited to underwater vehicles - all TRN filters have proved susceptible in areas with flat terrain [4]. In the Gulf War and in the Iraq War, over 1% of TERCOM guided cruise missiles did not meet accuracy expectations [5] [6], in some instances resulting in missile crashes in friendly nations [7].

False fixes greatly reduce the utility of TRN, as the integrity of a position estimate is paramount. In the case of close proximity operations, incorrect TRN estimates can lead to planned trajectories intersecting the terrain. For extended duration missions, avoiding false fixes requires minimizing the time the vehicle spends flying over flat terrain. This not only requires longer, less direct trajectories but limits operations to high information regions.

This paper demonstrates that the current method of dealing with map error is the cause of false fixes in regions with uninformative terrain. The issue fundamentally lies with the interpretation of map noise as information in the terrain. This overoptimistic expectation of terrain information results in a faster, unmerited reduction of position covariance that can lead to false fixes.

A new method for mitigating false fixes is proposed - directly modifying the correlation step to depend on the amount of map error, sensor error and information in the terrain. This approach is similar in concept to the variance adjustments used in robotics [8] and statistics [9] to reduce the likelihood of overconfidence, but with the important difference that this adjustment is not constant but depends on local terrain.

In order to implement these changes, the TRN filter requires an estimate of the information present in the terrain. This information is viewed as the amount of variance in the expected measurements that can not be attributed to map error, and thus can be used to reliably converge.

This approach is successfully demonstrated in a simulation environment as well as on offline data from a MBARI mapping vehicle run at Portuguese Ledge in Monterey Canyon. In both cases, using a standard TRN filter yields incorrect, overconfident results. The modifications employed result in greater filter reliability as measured by converging to the correct position with a credible estimate of position uncertainty.

II. The Basic TRN Problem

The goal of TRN is to localize a vehicle as best as possible with respect to the underlying terrain. This is accomplished through an estimation filter that fuses a vehicle motion model with an observation model to produce the best estimate of the vehicle position given the inertial and range measurement.

The baseline filter presented here is characteristic of TRN applications, and has had success in many environments. It consists of three components: a particle filter used to model the uncertainty in the vehicle position, a kinematic process model to propagate the vehicle dead reckoning estimate and associated error, and a range measurement model to encapsulate how the vehicle sees the terrain and the error in the a priori map.

A. TRN Filter

The fundamental non-linearity of natural terrain introduces multi-modality that requires the use of non-parametric Bayesian filters. A particle filter was chosen for its general applicability and ease of expansion to higher state estimates.

A particle filter models the belief distribution as the set:

$$\mathcal{X}_k := x_k^{[1]}, x_k^{[2]}, \ldots, x_k^{[M]} \tag{1}$$

where each particle x_k^m $(1 < m < M)$, is a discrete hypothesis with its own likelihood $w_k^{[m]}$.

B. Process Model

The process model is used to describe the growth in uncertainty that is introduce by dead reckoning. This depends on the type and quality of the sensors available on the vehicle. The vehicle used in these tests is a MBARI Dorado-class mapping AUV, which is equipped with a high-grade inertial navigation system with an accuracy of $< .05\%$ DT when aided by DVL velocity measurements. The availability of inertial-grade navigation sensors means that a kinematic process model is merited [10]:

$$x_{k+1} = x_k + \delta x_k^{\text{INS}} + r_k \tag{2}$$

where the state x_k is the vehicle position in northings, eastings and the depth bias relative to the local map

$$x = [x_N \; x_E \; z] \tag{3}$$

and $r_k = \mathcal{N}(0, \Sigma_{\text{INS}})$ models the INS noise.

In cases of poor vehicle odometry, low quality instruments, or frequent DVL outages, the vehicle state and process model can be expanded to accomodate the estimation of additional error states [11]. However, estimators typically perform better if the INS quality justifies the reduction of the number of estimated parameters.

C. Measurement Model

The TRN filter incorporates measurements to bound the growth of error accrued from dead reckoning. The measurement model evaluates the likelihood of observing those ranges at different positions in the map, and is thus able to bound the region of plausible vehicle locations.

For TRN, this is accomplished by correlating the recorded vehicle range measurements with an a priori terrain map. The ranges recorded by the vehicle at time step k are referred to as y_k, where $\boldsymbol{y}_{i,k}$ denotes the i^{th} beam. The a priori map is generated from sonar ranges recorded during mapping runs and is noted as $\hat{h}$. Both the map and the vehicle measurements are treated as random error off the true terrain, h.

The sensor model treats each range measurement $y_{i,k}$ as the true terrain at the projected location, $h(x_k)_i$, with a range dependent error $e_{i,k}$.

$$y_{i,k} = z_k - h(x_k)_i + e_{i,k} \tag{4}$$

$$e_{i,k} = \mathcal{N}(0, \sigma^2_{i,k,\text{sensor}}) \tag{5}$$

The map model assumes the error between the map and the terrain is fixed and independent of position in the map.

$$\hat{h}(x_k)_i = h(x_k)_i + \nu_{\text{map}} \tag{6}$$

$$\nu_{\text{map}} = \mathcal{N}(0, \sigma^2_{\text{map}}) \tag{7}$$

Thus, the expected range measurement, $\hat{y}_{i,k}$, is:

$$\hat{y}_{i,k} = \hat{z}_k - \hat{h}(\hat{x}_k)_i \tag{8}$$

where $\hat{h}(\hat{x}_k)_i$ is the elevation for the a prior map projected from the estimated position, $\hat{x}_k$ for beam i.

If the sonar measurement noise is uncorrelated with the map error and the beam errors are independent, the probability of getting a measurement sequence y_k, given vehicle state x_k is:

$$p(\boldsymbol{y}_k | x_k) = \eta \exp\left(-\frac{1}{2}\sum_i \beta_{i,k}(y_{i,k} - \hat{y}_{i,k}))^2\right) \tag{9}$$

$$\beta_{i,k} = \frac{1}{\sigma^2_{i,e,\text{sensor}} + \sigma^2_{\text{map}}} \tag{10}$$

This likelihood evaluation is the fundamental output of the measurement update, however, it does not take into account the variance in the correlation introduced by map error. This variance grows independent of the information in the terrain, and is the mechanism through which map error can lead to false fixes.

III. False Fixes in TRN

TRN filters have a demonstrated tendency of producing overconfident position estimate [12]. In some cases, this results in an estimate that is both overconfident and wrong - a false fix. This is particularly worrisome in many TRN applications - if the filter estimate is incorporated into vehicle trajectory planning, the planner runs the risk of planning a trajectory intersecting the terrain.

False fixes occur most often in flat, uninformative terrain. This has motivated a large amount of effort towards choosing informative terrain for TRN operation, as well as identifying and reducing the incidences of false fixes [3]. Simulations in [13] showed that false fixes increase with the amount of noise in the map, and [14] used information metrics to choose good regions of terrain for correlation. The correlator integrity methods in [3] use heuristics considering the amount of variation in the the terrain contour measured by the vehicle and it's correlation with the map, but do not separate the impact of map error from sensor error.

Separating map error from sensor error shows that the amount of information conveyed by the map is limited by the map error. In the case of uninformative terrain, the correlation variance introduced by map error may exceed the benefit gained by incorporating the measurement information. This motivates the introduction of a modified algorithm that reduces the imporantance of uninformative measurements but uses informative measurements when they are available.

A. Algorithm Modification

There is a strong history in the robotics and statistics communities of modifying correlation properties to improve estimator robustness. In robotics, this is frequently done to mitigate the impact of unmodeled or correlated errors.

One common method in robotics of dealing with overconfidence is reducing the information conveyed in a measurement by exponentiating the probabiltity [8]:

$$p(\boldsymbol{y}_k|x_k, m) = p(\boldsymbol{y}_k|x_k, m)^\alpha \tag{11}$$

where $\alpha < 1$. This type of modification is also used in statistics on a similar problem - correctly assessing the relative likelihood of multiple hypotheses. Royall et al [9] use this method for estimating model parameters given imperfect models. Royall shows that adjusting the correlator variance to depend on modeling parameters reduces the likelihood of producing strong, misleading evidence - an analog to the false fix problem.

Following this methodology, the parameter α is chosen to depend on the information in the terrain in addition to the map and sensor error. The choice of α can be interpreted as scaling the observed correlation by the information in the expected measurement.

$$\alpha_{i,k} = \frac{(\sigma^2_{i,e,\text{sensor}} + \sigma^2_{\text{map}})\hat{\delta}^2_{\text{info}}}{(\sigma^2_{i,e,\text{sensor}} + \sigma^2_{\text{map}})(\hat{\delta}^2_{\text{info}} + \sigma^2_{\text{map}}) + \sigma^2_{i,e,\text{sensor}}\sigma^2_{\text{map}}} \tag{12}$$

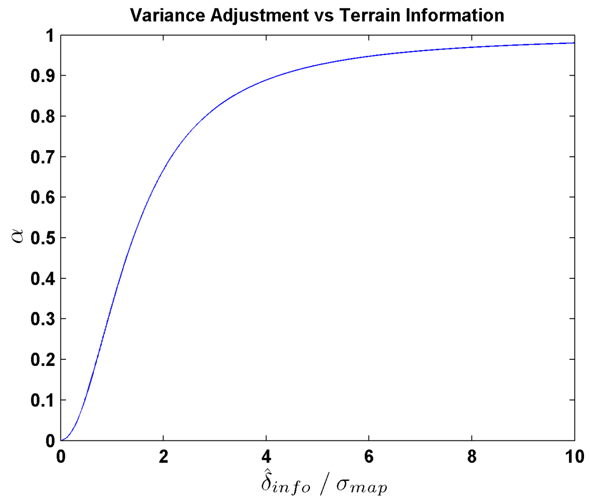

Fig. 1. Modification α vs $\frac{\hat{\delta}}{\sigma_{\text{map}}}$, for the case of $\sigma_{\text{map}} = \sigma_{\text{sensor}}$. Note that the change in α slows notable for values greater than 3.

where $\hat{\delta}^2_{\text{info}}$ is the estimate of the information contained in the true terrain. The behavior of α vs terrain information is shown in Figure 1 for the case of equal map and sensor. This shows that α has several appealing properties. In the case of uninformative terrain where $\hat{\delta}^2_{\text{info}}$ is small, $\alpha \to 0$. This essentially disregards any information in the terrain as it likely to be overshadowed by noise. In information-rich terrain, $\hat{\delta}^2_{\text{info}}$ is large and $\alpha \to 1$, so the estimator retains its nominal convergence rate.

This modification differs from prior estimators as it depends on amount of information in the true terrain - a quantity which is unknown. Therefore the amount of information in the terrain must be estimated prior to calculating α. In this implementation, the information is derived from the spread of the distribution of expected measurements, where $\hat{y}_{i,k}$ is the expected measurement of beam i at time step k. For a particle filter, each particle has its own estimated measurement $\hat{y}^{[m]}_{i,k}$ that comes from the stored map:

$$\hat{y}^{[m]}_{i,k} = z^{[m]}_k - \hat{h}(x^{[m]}_k)_i \tag{13}$$

This may also be rewritten as error off the true terrain

$$\hat{y}^{[m]}_{i,k} = z^{[m]}_k - h(x^{[m]}_k)_i - \nu_{\text{map}} \tag{14}$$

where the quantity $z^{[m]}_k - h(x^{[m]}_k)_i$, noted as $y^{[m]}_{i,k}$ is what the particle $x^{[m]}_k$ would estimate as its measurement for beam i were it to know the true terrain. The information used to converge thus comes from the difference between the true terrain at the hypothesized location, $x^{[m]}_k$, and the true measurement $y^{\text{TRUE}}_{i,k}$ at the true location, x^{TRUE}_k.

$$\delta^{[m]}_{i,k} = y^{[m]}_{i,k} - y^{\text{TRUE}}_{i,k} \tag{15}$$

where the true terrain information driving filter convergence is

$$\delta^2_{\text{info}} = \sum_m w^{[m]}_k (\delta^{[m]}_{i,k})^2 \tag{16}$$

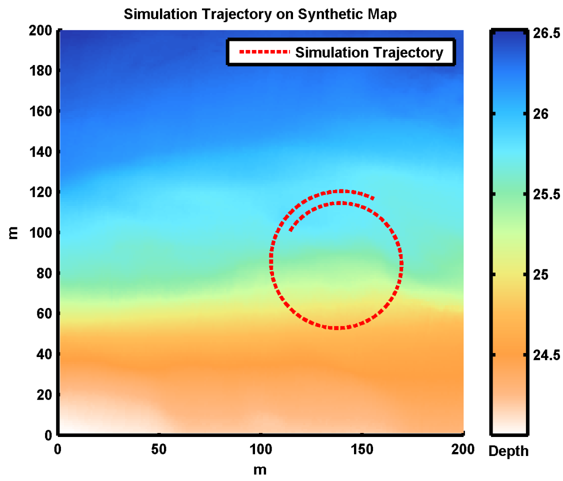

Fig. 2. Simulated trajectory on synthetic map. The slight north-south slope and east-west curvature provide terrain information for correlation.

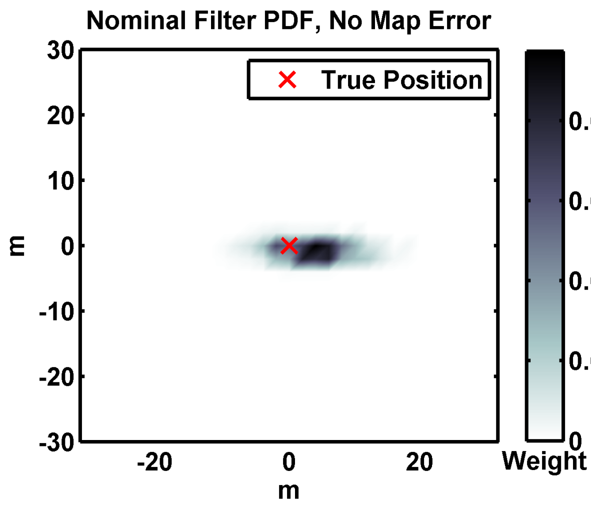

Fig. 3. Typical simulation results for nominal TRN filter in the case of no map error. Particle weights are summed in 2m x 2m bins to illustrate the weight distribution. This run resulted in a belief distribution centered off the true location but with uncertainty broad enough to encompass truth.

It is important to note that the neither the true location, the true terrain nor the true expected measurement are in fact known. Given this constraint, the goal is to come up with a conservative estimate for the terrain information. This estimate is derived from comparing the expected variance with the measured variance in $\hat{y}_{i,k}$.

$$E\left[Var(\hat{y}_{i,k})\right] = \sum_m w_k^{[m]}(\delta_{i,k}^{[m]})^2 - \left(\sum_m w_k^{[m]}\delta_{i,k}^{[m]}\right)^2 + \left(1 - \sum_m (w_k^{[m]})^2\right)\sigma_{\text{map}}^2 \quad (17)$$

In terms of the expected information.

$$E\left[Var(\hat{y}_{i,k})\right] = \delta_{\text{info}}^2 - \bar{\delta}_{i,k}^2 + \left(1 - \sum_m (w_k^{[m]})^2\right)\sigma_{\text{map}}^2 \quad (18)$$

From this it is clear that the minimum expected δ_{info}^2 is for $y_{i,k}^{\text{TRUE}} = \bar{y}_{i,k}$. Using the actual variance on the expected measurement, the terrain information can be estimated.

$$\hat{\delta}_{info}^2 = Var(\hat{y}_{i,k}) - \sigma_{\text{map}}^2 \quad (19)$$

This information estimate is inserted into equation 12 and used in the reweighting step, now posed as:

$$p(y_k|x_k) = \eta \exp\left(-\frac{1}{2}\sum_i \alpha_{i,k}\beta_{i,k}(y_{i,k} - \hat{y}_i(x_k))^2\right) \quad (20)$$

B. False Fixes in Simulation

The utility of the information adjusted variance filter modifications is readily demonstrated in simultion. The simulation environment takes the models described in section II and applies them to an AUV operating over synthetic terrain. The synthetic terrain shown in Figure 2 is derived from stored meter-resolution MBARI bathymetry, scaled down to reduce the impact of the original map noise. The synthetic map is simply this terrain with with gaussian noise added at each DEM gridpoint, as per the map model.

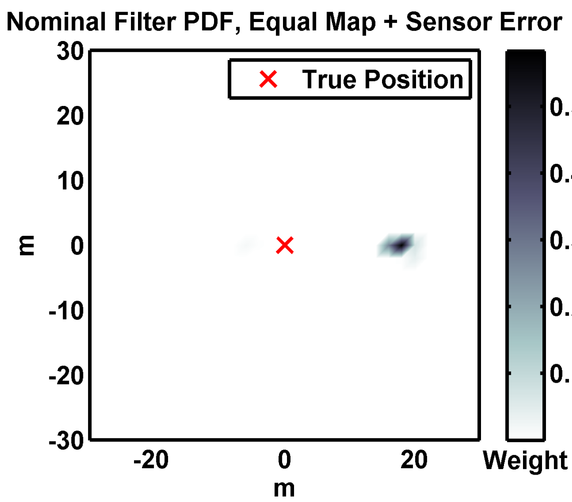

Fig. 4. Typical simulation results for nominal TRN filter with equal map and sensor error. Particle weights are summed in 2m x 2m bins to illustrate the weight distribution. This run resulted in a belief distribution that is far from truth and very overconfident.

The trajectory is chosen to be unique and non-repeating, in this case a simple spiral. The measurement model is simplified to an altimeter measuring the terrain profile along the trajectory. Depth bias is eliminated to ensure that only the map and sensor error drove the filter convergence. The simulations use a filter similar to the one described in section II, reduced to estimating only the east and north position. The particle filter is initialized with 10000 particles randomly spread over an area 40 meters by 40 meters, and operates at an update frequency of 2Hz (measurements taken at $\sim$1m intervals).

The important feature of these simulations is that the distribution of map and sensor error is exactly controlled. This allows examination of the validity of the common TRN assumption of summing map and sensor error in filter design. The total error, is set to an arbitrary value σ_0^2. The map and sensor error are then set such that $\sigma_{\text{map}}^2 + \sigma_{\text{sensor}}^2 = \sigma_0^2$. In the test case with no map error, this means that $\sigma_{\text{sensor}}^2 = \sigma_0^2$. In the scenario with map error, the error is divided evenly between the map and the sensor so that $\sigma_{\text{sensor}}^2 = \sigma_{\text{map}}^2 = 0.5\ \sigma_0^2$. This distribution is chosen as the sonar measurements used to build the map are likely to be of similar quality to those measured by the AUV along the trajectory.

The impact of map error versus sensor error on standard TRN algorithms can be seen in Figures 3 and 4. These figures show the particle weights after 250 measurements, grouped into 2m x 2m bins. This offers a good visualization of the distribution of the estimate uncertainty: broader spread and lower peak indicates greater uncertainty, whereas a tighter spread and higher peak indicates less uncertainty.

Figure 3 shows the distribution in the case of no map error. When no map error is present, the standard TRN filter behaves just as expected. Although the terrain is deliberately chosen to be of low information, the filter produces an estimate that is close to truth and does not appear to be overly confident.

Figure 4 shows what happens when there is an even distribution of map and sensor error. Although both simulations have the same amount of total error, the standard TRN filter performs markedly worse in the presence of map error. The standard TRN filter produces an erroneous, overconfident estimate. The scale in Figure 4 shows that the peak of this distribution is almost 5 times higher than in Figure 3, and has no weight at the true location. This result is clearly erroneous - introducing map error does not increase the information in the underlying terrain and should not result in lower covariance.

These simulations illustrate that map error has a profoundly different effect on filter convergence than sensor error, and must therefore be treated differently. The modified filter outlined in section II is designed to mitigate exactly this kind of problem - by slowing the correlation weight in less informative areas, the filter should produce an estimate that with a lower likelihood of false fixes.

Figure 5 shows a typical result when running the modified algorithm under the same equal error conditions that results in false fixes when using a standard TRN filter. The modified algorithm avoids converging on the error in the map and produces an estimate of the position that does not discard the true location. Although the filter estimate is more uncertain than the no map error case, this is to be expected as it must pay the penalty of a broader covariance for avoiding a false fix.

IV. Offline Demonstration on MBARI Dorado Class AUV

The performance of the modified and standard TRN algorithms are compared using data taken by the MBARI Dorado-Class mapping AUV (Figure 6) during a mapping run at

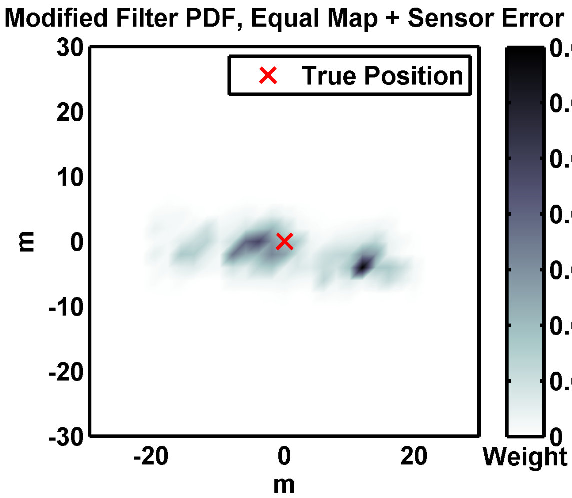

Fig. 5. Typical simulation results for modified TRN Filter with equal map and sensor error. Particle weights have been summed in 2m x 2m bins to illustrate the weight distribution. Although the simulation results shown in Figure 4 have the same total error, using the modified filter results in a more reasonable estimate of position uncertainty.

Fig. 6. MBARI Dorado-class mapping vehicle.

Portuguese Ledge in Monterey Bay. This run gathered the multibeam range data used to produce the map shown in Figure 7. The map was generated by projecting the multibeam data taken along the vehicle trajectory, so the position of the vehicle in the map is precisely known. This allows the performance of the standard and modified TRN algorithms to be compared relative to a known true trajectoy.

The offline simulations use a particle filter with 10000 particles initialized over a 150m x 150m region at the start of the dashed trajectory shown in Figure 7. Figure 8 and Figure 10 show typical results of using a standard TRN algorithm and the modified TRN algorithm, respectively. These plots are close-ups in the lower right corner of the map, to demonstrate the filter performance when the vehicle reaches the informative rocky outcropping after flying over flat terrain for 2000 meters. The plots show the evolution of the mean in dashed red and uncertainty in red ellipses relative to the true trajectory.

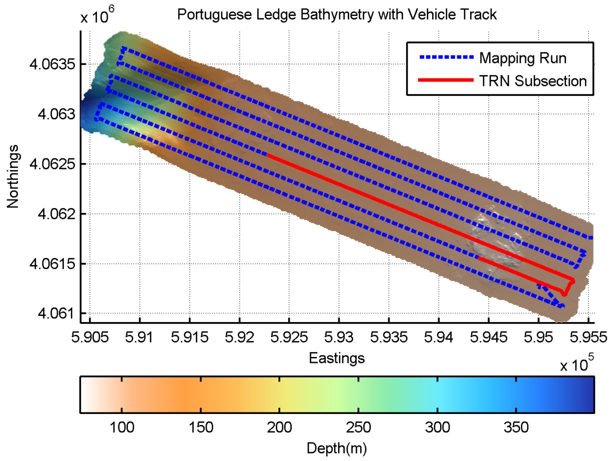

Fig. 7. MBARI Dorado Mapping Run at Portuguese Ledge 04/24/08. UTM North is plotted on the y-axis, and UTM East is plotted on the x-axis, both in meters. The Kearfott INS-predicted trajectory is overlaid on the map in black, and the red, dashed section is used for testing the TRN algorithms. This section was chosen as it flies over a very gentle slope ($< 0.5°$) for 2km before reaching a rocky outcropping of 20m prominence

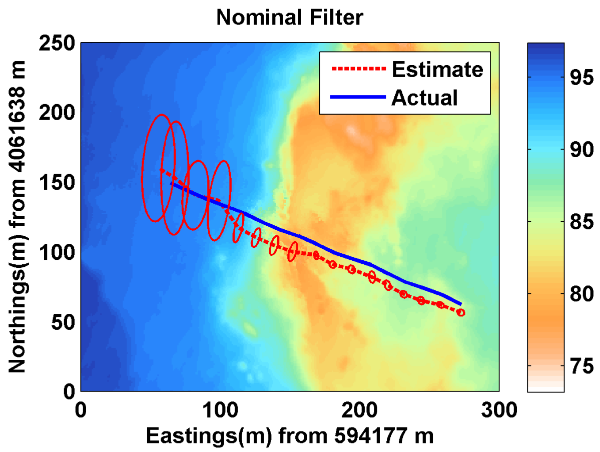

Fig. 8. Offline Results for Nominal TRN Filter. The red dashed line and ellipses show the evolution of the mean and 90 percent confidence interval for the nominal TRN filter. In this instance, the nominal filter converges to an estimate 8m south of the true position.

Although the simulations use a particle filter, ellipses are chosen to represent the position uncertainty for the sake of clarity.

The nominal filter results, overlaid on the map in Figure 8, converge too much before reaching informative terrain. Though the mean of this distribution is correct, the filter has already discarded the estimate in the true location. After reaching the information rich terrain, it converges to an incorrect location - a false fix. When looking at the 90pct confidence interval in east and north, the filter is shown to converge to offset of ~8m.

The modified filter results, shown in Figure 10, look con-

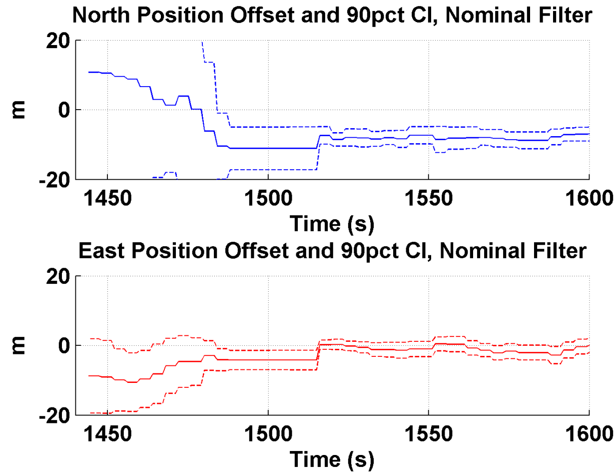

Fig. 9. Offline Results for Nominal Filter, 90 percent confidence interval for north and south position estimates. The convergence to a biased estimate is most evident in the 8m offset and low uncertainty in the north position.

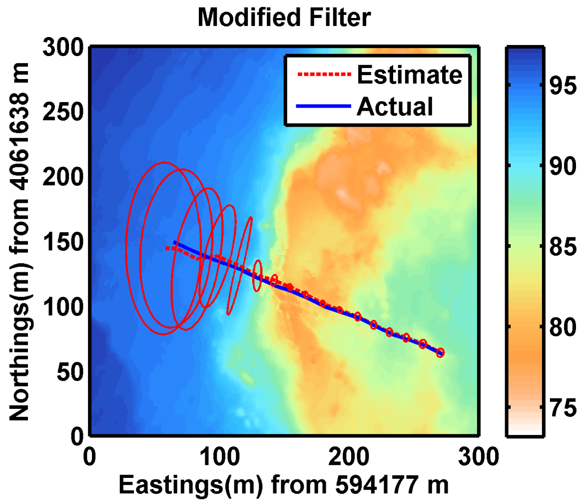

Fig. 10. Offline Results for Modified Filter. The red dashed line and ellipses show the evolution of the mean and 90 percent confidence interval for the nominal TRN filter. Although the initial position uncertainy entering the rocky area is high, the filter quickly converges to an estimate close to the true position.

siderably different. The filter has disregarded most of the measurements in the flat terrain and has an uncertainty roughly twice that of the standard filter. Upon entering informative terrain, however, the modified filter converges rapidly to the correct location. By avoiding growing overconfident in the flat terrain it does not discard the true location. The filter converges to an estimate with an uncertainty on the order of a few meters in east and north (see Figure 11, characteristic of TRN positioning accuracy.

The beam weighting parameters α calculated by the modified filter are in Figure 12. The terrain information in the rocky outcropping results in high α's as the vehicle approaches. At 1475s, the outcropping first appears in the expected measurements of the front facing DVL beams 1 and 3.

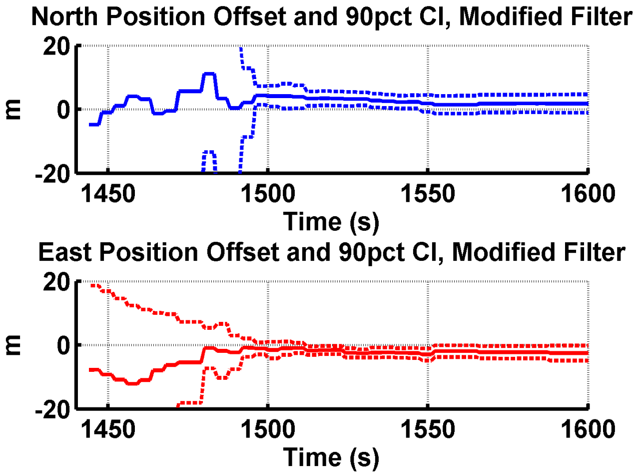

Fig. 11. Offline results for modified filter, 90 percent confidence interval for north and south position estimates.

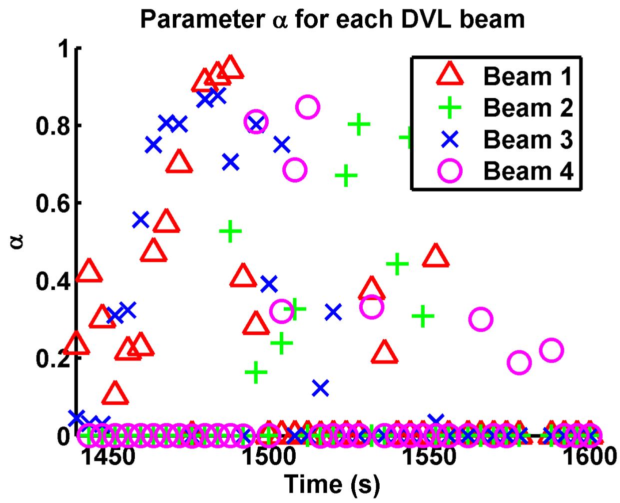

Fig. 12. Parameter α used in Modified Filter weighting for each beam. DVL beams 1 and 3 are the first to see the outcropping are weighted as very informative once this terrain comes into view. After this information is incorporated into the estimate, their weighting drops as the expected measurement variance is small.

These are therefore given a high α due to the reliability of their information. Although this same transition occurs for rear facing DVL beams 2 and 4 at 1525s, they are not weighted as highly. The estimate has already converged, so the amount of expected measurement variation is much lower. The conservative nature of the information estimate is also evident, as measurements are frequently completely ignored ($\alpha = 0$).

These results support the use of the modified TRN filter. The modified TRN filter is capable of the same level of accuracy when operating in informative terrain, while mitigating false fixes due to map noise.

V. Conclusion

The presented simulations and offline results demonstrate that map error is a source of false fixes in flat terrain. The simulations show that these errors can be mitigated by adjusting the filter measurement weighting to depend on the relative amounts of map error, sensor error, and terrain information. The modified algorithm introduced in this paper is adaptable to all Bayesian TRN filters operating in all types of terrain.

The offline demonstration shows that slowing convergence of the filter in uninformative terrain can lead to higher integrity estimates when entering informative terrain. Although the modfied filter produced a much higher uncertainty in the flatter portion of the map, it reliably converged to the true path. The positioning accuracy achieved using the modified filter also matches typical TRN accuracy when operating in informative terrain.

Acknowledgment

The authors would like to thank MBARI for support and the Stanford Graduate Fellowship and the National Defense Science and Engineering Graduate Fellowship for partial funding of this work. Researchers at MBARI who were particularly invaluable in acquiring the presented data include Dave Caress and Hans Thomas.

References

[1] J. Golden, "Terrain contour matching (TERCOM): a cruise missile guidance aid," vol. 238, pp. 10–18, SPIE, 1980.

[2] K. Anonsen and O. Hagen, "An analysis of real-time terrain aided navigation results from a HUGIN AUV," in *OCEANS 2010*, pp. 1–9, Sept. 2010.

[3] O. Hagen, K. Anonsen, and M. Mandt, "The HUGIN real-time terrain navigation system," in *OCEANS 2010*, pp. 1–7, Sept. 2010.

[4] N. Bergman, "A bayesian approach to terrain-aided navigation," 1997.

[5] U. S. G. A. Office, "OPERATION DESERT STORM, evaluation of the air campaign," June 1997.

[6] J. Lewis, "Why the Navy should retire TLAM-N," December 2009.

[7] A. Purvis, "In search of a lost cruise missile," *Time*, April 2003.

[8] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Intelligent Robotics and Autonomous Agents, MIT Press, 2005.

[9] R. Royall and T.-S. Tsou, "Interpreting statistical evidence by using imperfect models: robust adjusted likelihood functions," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 65, no. 2, pp. 391–404, 2003.

[10] G. Donovan, "Position error correction for an autonomous underwater vehicle inertial navigation system (INS) using a particle filter," *Oceanic Engineering, IEEE Journal of*, vol. 37, pp. 431–445, July 2012.

[11] D. Meduna, S. Rock, and R. McEwen, "Closed-loop terrain relative navigation for AUVs with non-inertial grade navigation sensors," in *Autonomous Underwater Vehicles (AUV), 2010 IEEE/OES*, pp. 1–8, Sept. 2010.

[12] K. Anonsen and O. Hallingstad, "Terrain aided underwater navigation using point mass and particle filters," in *Position, Location, And Navigation Symposium, 2006 IEEE/ION*, pp. 1027–1035, 2006.

[13] W. Tang and R. L. McClintock, "Terrain correlation suitability," vol. 2220, pp. 50–58, SPIE, 1994.

[14] I. Nygren and M. Jansson, "Terrain navigation for underwater vehicles using the correlator method," *Oceanic Engineering, IEEE Journal of*, vol. 29, pp. 906–915, July 2004.