Robust Adaptive Terrain-Relative Navigation

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 a technique for localizing a vehicle in GPS-denied environments. TRN augments a dead-reckoned solution with continuous position fixes based on correlations with a pre-stored map. In underwater applications TRN accuracy on the order of 3m has been demonstrated, however convergence to incorrect solutions has been observed when operating for extended periods over featureless terrain. Specifically, the TRN filter can become overconfident in an incorrect position fix.**

Previous work by the authors introduced an adaptive technique for mitigating over-confidence in uninformative terrain. Specifically, the algorithm exponentially down-weights the probabilities with a factor α between zero and one, based on the estimated terrain information. This paper focuses on understanding the source of false fixes in uninformative terrain, and uses this insight to develop the method behind the adaptive technique.

This paper shows that the cause of false fixes using standard TRN weighting in information-poor regions is the assumption that the terrain is uncorrelated. It also introduces a method to analyze the probability of false peaks using the standard TRN measurement weighting. This analysis is an extension of work in the statistics community on robust adjusted likelihood ratios, and is used to bound the probability of large false peaks.

The resulting robust, adaptive technique is capable of real-time operation and its effectiveness is demonstrated in simulation and with 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 drift accrued during dead reckoning by matching range measurements of the terrain to an on-board terrain map. The TRN concept was first deployed as the cruise missile guidance system TERCOM (Terrain Contour Matching), which used a batch correlation of altimeter measurements against an a priori map to generate a position estimate [1].

TRN was adapted to autonomous underwater vehicles (AUVs) starting in the 1990s [2] [3] and TRN packages are now available on commercial vehicles such as the HUGIN AUV [4]. Map-relative position estimates are particularly valuable for AUVs as they reduce the necessity of surfacing for GPS fixes or deploying USBL or LBL positioning systems. Further, many return-to-site underwater missions are to locations identified in a map, and since underwater maps can be inaccurately geo-referenced, a map-relative position is necessary [5].

In informative terrain TRN accuracy is typically on the order of the map resolution. Accuracy of a few meters has been demonstrated in field trials using MBARIs Dorado-class AUVs in Monterey Bay [6], and similar performance was observed in field trials using the HUGIN AUV in Ostfjords [7].

When operating over featureless terrain for an extended period, converging to an overconfident, incorrect solution - a false fix - is a recognized issue with TRN filters [8] [9]. False fixes greatly reduce the value of a TRN estimate. An incorrect TRN estimate can result in failure to accomplish a mission, and in the case of close proximity operations, can lead to trajectories intersecting the terrain.

This has motivated a large amount of effort towards understanding and reducing the incidence of false fixes due to uninformatve terrain. Simulations in [10] showed that false fixes increase with the amount of noise in the map. The information metrics developed in [11] were used to select informative regions of terrain for correlation. The integrity methods in [7] reject TRN position fixes if the variation in the terrain contour is insufficient.

The adaptive method presented in [12] is designed to slow convergence in uninformative terrain to prevent false fixes. This paper focuses on understanding the source of false fixes in uninformative terrain, and uses this to develop the method behind the adaptive technique.

In the standard formulation of the TRN problem, a Gaussian likelihood function is typically used to perform the measurement update. That is:

$$L(x)_{\text{standard}} = p(z, \hat{h}|x)_{\text{standard}} = \eta \exp\left(\frac{-(z-\hat{z})^2}{2(\sigma^2_{\text{map}} + \sigma^2_{\text{sensor}})}\right) \tag{1}$$

where z is the measurement, $\hat{z} = \hat{h}(x)$ is the estimated measurement from the map, and σ^2_{map} and σ^2_{sensor} are the map and sensor noise.

This paper shows that this formulation implicitly assumes that the vehicle is operating over informative terrain, and that this results in false fixes when operating over uninformative terrain. This assumption enters the measurement update as a probabilistic terrain model that characterizes the degree of spatial correlation in the terrain; in standard TRN, the terrain is assumed to be uncorrelated. This assumption is reasonable when the terrain correlation is small with respect to error in the map, but results in false peaks when the terrain correlation

is large with respect to map error.

The relation between the level of terrain correlation and the occurrence of false peaks using standard TRN correlation is demonstrated using simulated informative and uninformative terrains generated from probabilistic terrain models. The probabilistic terrain models employed are Gaussian and are commonly used in the Geographic Information System (GIS) community to characterize spatial correlation. Likelihood surfaces are computed using the standard uncorrelated terrain model and the known Gaussian terrain model for both cases. Both probabilistic terrain models have comparable performance in informative terrain, but in uninformative terrain the correlated Gaussian terrain model eliminates the false peaks that arise using standard TRN correlation. This motivated the development of the adaptive technique presented in [12].

The result of this development is that the likelihood function used in the measurement update becomes

$$L(x)_{\text{adjusted}} = L(x)^{\alpha}_{\text{standard}} = p(z, \hat{h}|x)^{\alpha}_{\text{standard}} \tag{2}$$

Specifically, this adaptive adaptive method exponentially down-weights the probabilities calculated using the standard TRN algorithm to account for the false peaks introduced by the uncorrelated terrain model. This method is based on analyzing the likelihood of false peaks using the standard TRN measurement weighting, and is an extension of work in the statistics community on robust adjusted likelihood ratios [13]. This analysis shows that adjusting the probabilities with a factor α between zero and one, based on the estimated terrain information, can eliminate the large false peaks using the standard TRN weighting. Here, terrain information is the amount of variance in the expected measurements that can not be attributed to map error. This method achieves good performance in simulated runs on both informative and uninformative terrain, at a computational cost comparable to using the standard TRN weighting.

This paper is organized as follows. Section II outlines the basics of the TRN problem. Section III describes how terrain models enter in to the measurement update, and how the standard TRN assumption results in overconfidence. Section IV introduces the method of mitigating false-fixes by matching false peak likelihoods. Section V compares results using the standard and adaptive TRN algorithms in offline runs using field data gathered in flat terrain in Monterey Bay.

II. The Basic TRN Problem

The goal of a TRN estimator is to localize a vehicle with respect to the underlying terrain. This is accomplished using an estimation filter that fuses a vehicle motion model with an observation model.

The filter presented here is characteristic of TRN applications. 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 estimate the vehicle position given range measurements and a terrain.

A. TRN Filter

The TRN filter models the probability distribution of the vehicle position. The fundamental non-linearity of natural terrain introduces multi-modal distributions that require the use of non-parametric filters. A particle filter was selected for use in this work due to its general applicability and ease of expansion to higher state estimates.

A particle filter models the belief distribution as the set:

$$\mathcal{X}_k := x_k^{[1]}, x_k^{[2]}, \ldots, x_k^{[M]} \tag{3}$$

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

B. Process Model

The process model characterizes the growth in position uncertainty introduced by dead reckoning. The rate of growth of uncertainty 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. Due to the availability of inertial-grade navigation sensors, the state is the vehicle position in northings, eastings and depth:

$$x = [x_{\text{N}}\ x_{\text{E}}\ x_{\text{D}}] \tag{4}$$

and a kinematic process model is used [14]:

$$x_{k+1} = x_k + \delta x_k^{\text{INS}} + r_k \tag{5}$$

The change in position at time step k, δx_k^{INS}, is measured by the INS and $r_k \sim \mathcal{N}(0, \Sigma_{\text{INS}})$ models the INS noise. The inclusion of depth as a state is necessary as local maps often have depth offsets.

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 [15]. Minimizing the number of estimated parameters, however, is valuable for making non-parametric filters computationally tractable.

C. Measurement Model

The TRN measurement update bounds the growth of error from dead reckoning. This model evaluates the likelihood of observing measured ranges at different positions in the map, and is thus able to constrain the region of plausible vehicle locations. The ranges recorded by the vehicle at time step k are referred to as $\boldsymbol{z}_k$, where $z_{i,k}$ 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 $p(z|x, h)$ treats each range measurement $z_{i,k}$ as a noisy measurement of the true terrain at the projected location, $h(x_k)_i$:

$$z_{i,k} = x_{\text{D},k} - h(x_k)_i + \nu_{i,k,\text{sensor}} \tag{6}$$

$$\nu_{i,k,\text{sensor}} \sim \mathcal{N}(0, \sigma^2_{i,k,\text{sensor}}) \tag{7}$$

The map model $p(\hat{h}|h)$ assumes the map, $\hat{h}$, is a noisy model of the terrain, h:

$$\hat{h} = h + \nu_{\text{map}} \tag{8}$$

$$\nu_{\text{map}} \sim \mathcal{N}(0, \Sigma_{\text{map}}) \tag{9}$$

Here $\hat{h}$ covers the entire map, and Σ_{map} is assumed to be diagonal with covariance σ^2_{map}.

The Bayesian update for the particle filter requires calculating the likelihood of the measurements given a position:

$$p(x^m_k|\boldsymbol{z}_k, \hat{h}) = \eta p(\boldsymbol{z}_k, \hat{h}|x^m_k) w^m_k \tag{10}$$

The standard method of computing $p(\boldsymbol{z}_k, \hat{h}|x^m_k)$ is a Gaussian weighting on the squared error:

$$p(\boldsymbol{z}_k, \hat{h}|x^m_k)_{\text{standard}} = \eta \exp\left(-\frac{1}{2}\sum_i \beta_{i,k}(z_{i,k} - \hat{z}^m_{i,k}))^2\right) \tag{11}$$

$$\beta_{i,k} = \frac{1}{\sigma^2_{i,k,\text{sensor}} + \sigma^2_{\text{map}}} \tag{12}$$

where $\hat{z}^m_{i,k}$ is the expected range measurement for each particle:

$$\hat{z}^m_{i,k} = x^m_{\text{D},k} - \hat{h}(x^m_k)_i \tag{13}$$

Here $\hat{h}(x^m_k)_i$ is the elevation at the projected location in the map for particle m at position x^m_k for beam i.

This weighting was demonstrated in [12] to result in false peaks in uninformative terrain, and the author introduced an adaptive parameter α to mitigate the false peaks by down-weighting measurements in uninformative terrain.

III. Probabilistic Terrain Models and TRN

The measurement update in Equation 10 is the foundation of the TRN correlation step, and the challenge is developing the model for $p(\boldsymbol{z}_k, \hat{h}|x^m_k)$. This is not specified by either the sensor model $p(z|x, h)$ or the map model $p(\hat{h}|h)$. Both of these models are relative to the *terrain*; the sensor is not measuring the map.

Bayesian TRN estimators typically use a Gaussian weighting as in Equation 11, however, this weighting does not account for the impact of spatial terrain correlation.

Evaluating the measurement update in Equation 10 with the known sensor and map models requires a probabilistic terrain model to account for correlation in the terrain. This correlation is introduced through $p(h)$ when the terrain h is marginalized:

$$p(\boldsymbol{z}_k, \hat{h}|x^m_k) = \eta\left(\int p(\boldsymbol{z}_k, \hat{h}|x^m_k, h)p(h)dh\right) \tag{14}$$

Incorporating the map and sensor models, this becomes

$$p(\boldsymbol{z}_k, \hat{h}|x^m_k) = \eta\left(\int p(\boldsymbol{z}_k|x^m_k, h)p(\hat{h}|h)p(h)dh\right) \tag{15}$$

The model $p(h)$ represents the prior statistics on the terrain, of which the primary characteristic is the degree of spatial correlation in the terrain. In informative terrain - where the correlation in the terrain is small with respect to map noise - $p(h)$ can be reasonably modeled as uniform:

$$p(h) = 1 \tag{16}$$

Incorporating this into Equation 15 yields the standard TRN weighting - Equation 10. However, this uniform prior assumption is the source of false peaks in uninformative terrain.

In uninformative terrain, $p(h) = 1$ insufficiently models the correlation in the terrain. The remainder of this section will demonstrate that $p(h) = 1$ is reasonable when the terrain is informative, but fails when the terrain is uninformative.

Specifically, simulated informative and uninformative terrain are used to demonstrate the impact of the $p(h) = 1$ assumption, by comparing the results of the standard TRN weighting with the results using a known terrain model.

In order to do this, a terrain model is required.

A. Correlated Terrain Model

The probabilistic terrain model selected for this work is drawn from the GIS community and represents the terrain as a Gaussian distribution with a common mean and the correlation between points in the terrain is only a function of the distance between them [16]:

$$p(h) \sim \mathcal{N}\left(\mu_{\text{terrain}}, \Sigma_{\text{terrain}}\right) \tag{17}$$

where $\mu_i = \mu_j = \mu_{\text{terrain}}$ and

$$\text{var}[h_i - h_j] = E\left[(h_i - h_j)^2\right] \tag{18}$$

$$= 2\gamma(\delta_{i,j}) \tag{19}$$

where $\delta_{i,j}$ is the separation between points h_i and h_j, and $\gamma(\delta)$ is the semi-variance.

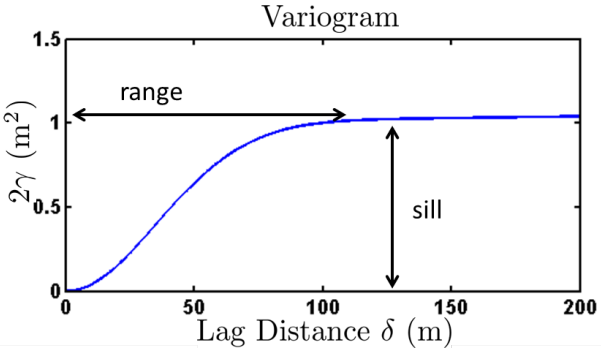

Fig. 1: Example variogram - the range and sill are the characteristic length over which terrain varies and the size of that variation.

The primary input to this model is the semi-variogram, $\gamma(\delta)$, which describes the spatial characteristics of the terrain. Two main characteristics are the range and sill, shown in Figure 1, where the range is the length scale over which the terrain varies, and the sill is the maximum amount that the terrain varies. In uninformative terrain the variance is small and grows

much more slowly as a function of distance. Subsitituting the Gaussian terrain model $p(h)$ into Equation 15 results in

$$p(\boldsymbol{z}_k, \hat{h}|x_k^m)_{\text{Corr. Terrain}} = \eta \frac{1}{\det(A_{x_k^m})} \exp\left(\frac{1}{2} b_{x_k^m}^T A_{x_k^m}^{-1} b_{x_k^m}\right) \tag{20}$$

where

$$A_{x_k^m} = H_{x_k^m}^T \Sigma_{k,\text{sensor}}^{-1} H_{x_k^m} + \Sigma_{\text{map}}^{-1} + \Sigma_{\text{terrain}}^{-1} \tag{21}$$

$$b_{x_k^m} = -H_{x_k^m}^T \Sigma_{k,\text{sensor}}^{-1} \boldsymbol{z}_k - \Sigma_{\text{map}}^{-1} \hat{h} - \Sigma_{\text{terrain}}^{-1} \mu_{\text{terrain}} \tag{22}$$

In Equation 21 $H_{x_k^m}$ is the matrix for the observation of the terrain for the estimate x_k^m, Σ_{map} is the map error (diagonal with value σ_{map}^2). As before, $\hat{h}$ is the map, μ_{terrain} and Σ_{terrain} are the mean and variance the terrain, $\boldsymbol{z}_k$ and Σ_{sensor} are the measurement and measurement noise, respectively. Both Σ_{map} and Σ_{terrain} are of dimension n by n and $H_{x_k^m}$ is of dimension n by l, where n is the total number of points in the map and l is the number of measurements. The matrix Σ_{sensor} is of dimension l by l.

B. TRN in Simulated Terrain

TRN runs are simulated using a short trajectory across simulated terrain. The terrain patches 50m by 50m are generated with variograms using a Gaussian variogram model,

$$\gamma(\delta) = \beta_0 + \beta_{\text{sill}} \left(1 - \exp\left(\frac{-\delta^2}{\beta_{\text{range}}^2}\right)\right) \tag{23}$$

where the parameters β_0, β_{sill} and β_{range} control the amount of variation in the terrain and the range over which it occurs. For both cases β_{range} is set to 20m and β_0 is set to $(.01\text{m})^2$, for informative terrain β_{sill} is set to $(5\text{m})^2$, for uninformative terrain β_{sill} is set to zero. The simulated informative and uninformative terrains are shown in Figures 2 and 6.

Maps of the terrain are made by adding uncorrelated Gaussian noise of variance σ_{map}^2 to all points in the terrain. Comparing the map for informative terrain in Figure 3 and uninformative terrain in Figure 7, the map noise is much more evident when there is little terrain information.

Measured terrain profiles are simulated by adding uncorrelated Gaussian measurement noise of variance σ_{sensor}^2 to the actual terrain profiles.

Likelihood surfaces are evaluated over a 20m by 20m area centered on the correct location at (0,0). Surfaces are computed using the standard weighting with $p(h) = 1$ from in Equation 11 and Gaussian weighting using Equation 20 with the known correlated Gaussian $p(h)$ used to generate the terrain.

In informative terrain $p(h) = 1$ is a reasonable approximation for the true $p(h)$, and both weightings produce similar results - a peak at the correct location. Results are shown in Figures 4 and 5 for the standard and Gaussian weightings, respectively.

In uninformative terrain, there is a stark difference between $p(h) = 1$ and using the correct $p(h)$. Using $p(h) = 1$ results in a number of false peaks, shown in Figure 8, whereas the correlated Gaussian $p(h)$ produces a low, broad likelihood surface, shown in Figure 9

This motivates the development of a method for reducing flat peaks in uninformative terrain.

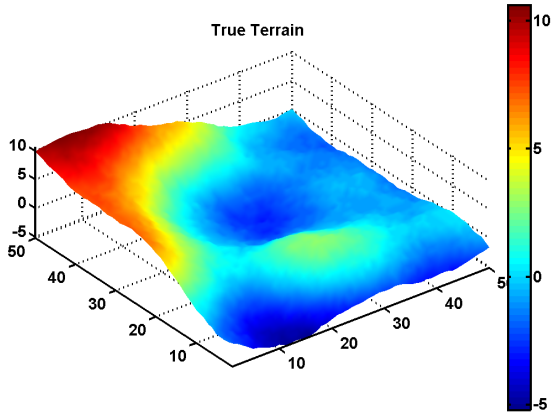

Fig. 2: Simulated informative terrain

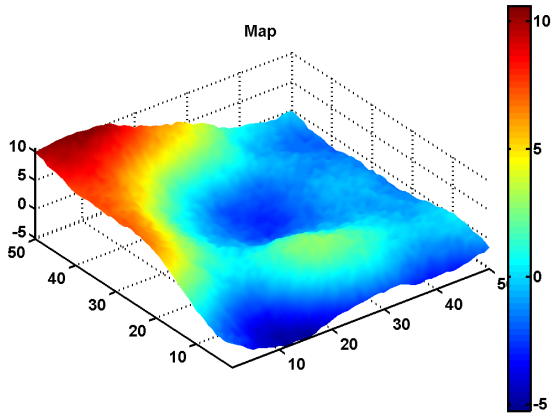

Fig. 3: Map of simulated informative terrain

IV. False Peak Bounds and TRN

The previous section illustrated how the terrain model used for standard TRN can cause issues in uninformative terrain. While one method of approaching this problem is to pursue improved terrain models, this section introduces a different method to achieve this based on bounding the size of false peaks.

This analysis is related to the false fix analysis by Nygren in [11]. Nygren analyzed likelihood ratios to determine the confidence in a PDF peak for use with a maximum-likelihood estimator, however, in Nygrens analysis the difference between the map at the two locations was assumed to be the true difference.

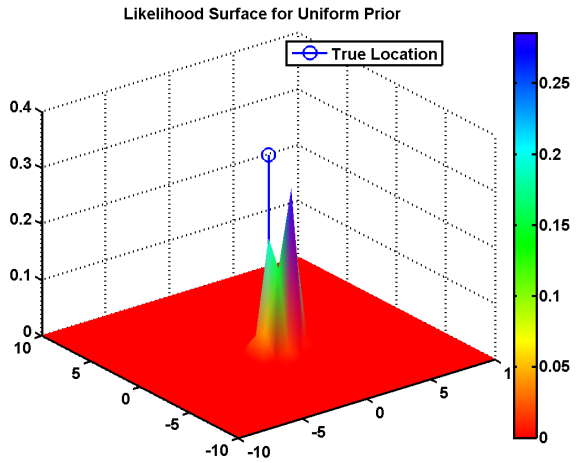

Fig. 4: Likelihood surface using standard TRN weighting on informative terrain

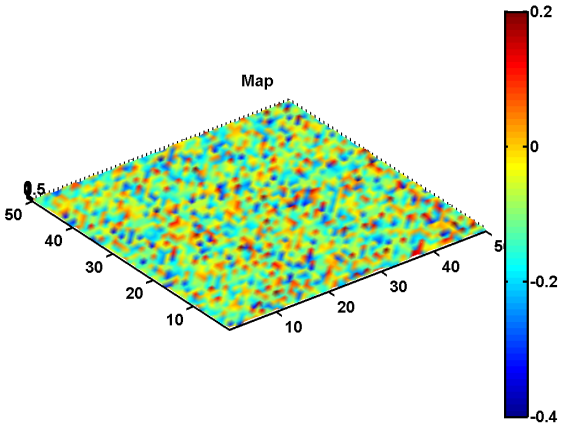

Fig. 7: Map of simulated flat terrain

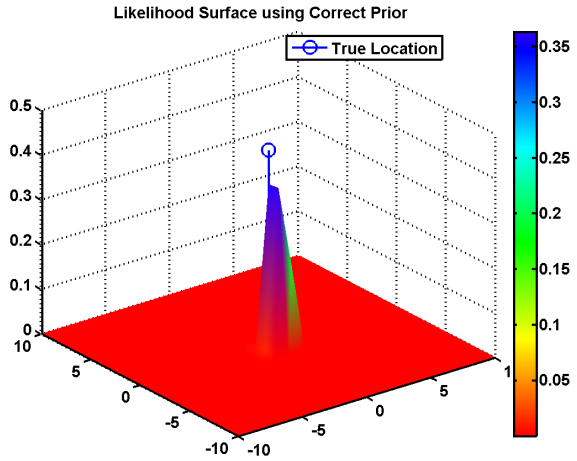

Fig. 5: Likelihood surface using known Gaussian terrain model on informative terrain

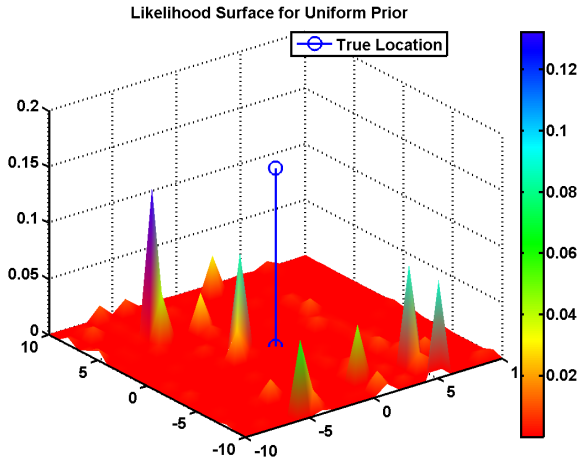

Fig. 8: Likelihood surface using standard TRN weighting on uninformative terrain

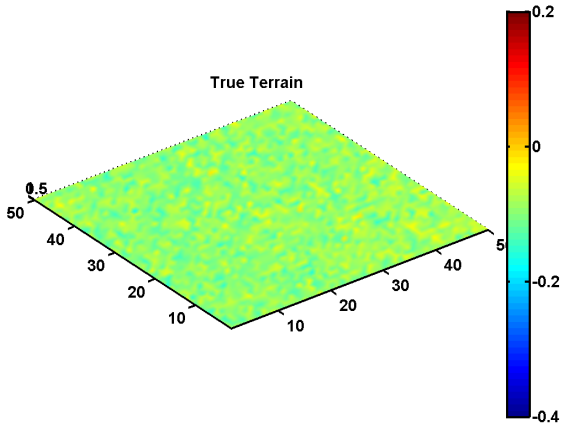

Fig. 6: Simulated flat terrain

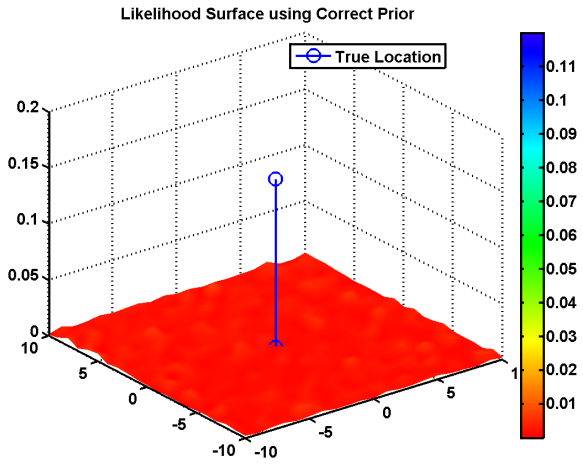

Fig. 9: Likelihood surface using known Gaussian terrain model on uninformative terrain

A. False Peak Bounds

A false peak occurs when the data appears to provide strong support for the wrong hypothesis. That is, when comparing for a correct hypothesis θ_{true} to a false hypothesis θ_{false}:

$$\frac{L(\theta_{\text{false}})}{L(\theta_{\text{true}})} = \frac{p(z|\theta_{\text{false}})}{p(z|\theta_{\text{true}})} >> 1 \tag{24}$$

For the TRN problem, the likelihood $L(\theta)$ and $p(z|\theta)$ are the likelihood $L(x)$ and probability $p(z, \hat{h}|x)$, respectively. While occasional support for the false hypothesis is a statistical certainty, large support is unlikely when the models underlying $p(z|\theta)$ are correct. Royall demonstrated in [17] that Equation 24 can be analyzed to determine the bound on the likelihood of false peaks, by considering

$$\max_N \left(\frac{L(\theta_{\text{false}})}{L(\theta_{\text{true}})} > k \right) \tag{25}$$

Where k is the size of the false peak, and N is the number of measurements, and $L(\theta_{\text{true}})$ is shorthand for $p(z|\theta)$. Maximizing this as a function of N returns a bound on the probability of vs the size of a false peak. In [17] it was shown that when using correct, Gaussian models this bound is

$$\max_N \left(\frac{L(\theta_{\text{false}})}{L(\theta_{\text{true}})} > k \right) = \Phi\left(-\sqrt{2\ln(k)} \right) \tag{26}$$

where Φ is the cdf of a normal distribution. Royall showed that using inaccurate models to calculate $p(z|\theta)$ can result in a high probability of large false peaks, where this bound is broken. Royall adopted the bound in Equation 26 as a *robust bound* in [13]. These bounds are shown in Figure 10. The robust bound is shown in solid blue, and an example of a bound calculated for an imperfect model is shown in dashed red.

Restoring this robust bound is the basis for making likelihood calculations robust to model error. The method introduced in [13] adjusts the probabilities with an exponential factor α

$$p(z|\theta)_{\text{adjusted}} = p(z|\theta)^{\alpha} \tag{27}$$

The ability to restore the false peak bound through this exponential adjustment is the primary driver behind the approach presented in this paper.

This adjustment is calculated from the data and model used in $p(z|\theta)$, and the output of this is termed the robust adjusted likelihood. Although the parametric models this was developed for are not directly applicable, the mindset of restoring bounds on false peaks is well suited to the problem of false peaks in uninformative terrain.

B. Applying False Peak Bounds to TRN

Section III demonstrated that the probabilistic terrain model used in standard TRN is the cause of false peaks in uninformative terrain. The remainder of this section is focused on adapting the robust adjusted likelihood methodology to the TRN measurement update. The goal is to estimate the parameter α, such that

$$p(z, \hat{h}|x)_{\text{adjusted}} = p(z, \hat{h}|x)^{\alpha}_{\text{standard}} \tag{28}$$

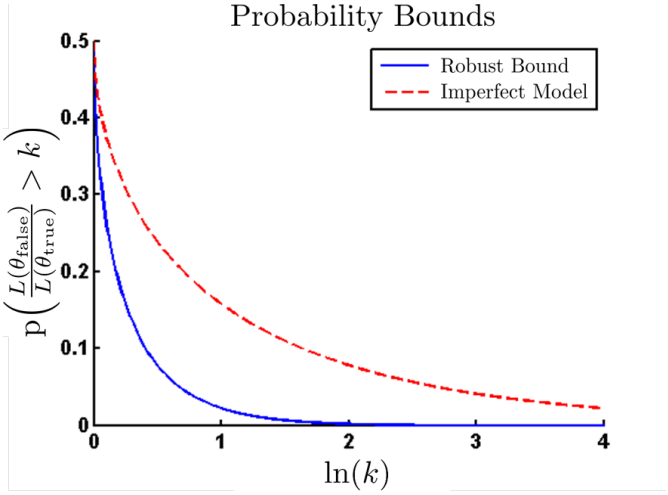

Fig. 10: Bounds on the probability of false peaks. The size of the false peaks increases to the right, and the bound is on the probability of a false peak of that magnitude occurring.

Applying likelihood ratio analysis to TRN, the likelihood for a single measurement update is

$$L(x_k^m) = p(\boldsymbol{z}_k, \hat{h}|x_k^m) \tag{29}$$

Adapting the methods from [13] requires modeling the probability distribution on this likelihood as the filter continues to integrate measurements. Therefore rather than a single measurement, the terrain profiles measured at the true and false trajectories are considered:

$$h_{\text{true}} = h(x_{\text{true}}) \tag{30}$$

$$h_{\text{false}} = h(x_{\text{false}}) \tag{31}$$

and strong support for the wrong position is when

$$\frac{p(\boldsymbol{z}, \hat{h}|x_{\text{false}})}{p(\boldsymbol{z}, \hat{h}|x_{\text{true}})} >> 1 \tag{32}$$

where x_{true} denotes the true vehicle position, x_{false} is a wrong position, and $\boldsymbol{z}$ are the vehicle measurements. At this point likelihood notation will be adopted, and applied to the standard Gaussian weighting on the terrain profile this means

$$\frac{L(x_{\text{false}})_{\text{standard}}}{L(x_{\text{true}})_{\text{standard}}} = \frac{p(\boldsymbol{z}, \hat{h}|x_{\text{false}})_{\text{standard}}}{p(\boldsymbol{z}, \hat{h}|x_{\text{true}})_{\text{standard}}} \tag{33}$$

The likelihood of large false peaks occurring as a function of the terrain information can then analyzed, with the simplifying assumption that σ^2_{sensor} and σ^2_{map} are constant. Incorporating these assumptions into Equation 11 yields

$$\frac{L(x_{\text{false}})_{\text{standard}}}{L(x_{\text{true}})_{\text{standard}}} = \exp\left(\frac{-e^T_{\text{false}} e_{\text{false}} + e^T_{\text{true}} e_{\text{true}}}{2\left(\sigma^2_{\text{sensor}} + \sigma^2_{\text{map}}\right)} \right) \tag{34}$$

where

$$e_{\text{true}} = z - \hat{z}_{\text{true}} \tag{35}$$
$$= (x_{\text{D,true}} - h_{\text{true}} + \nu_{\text{sensor}}) - (x_{\text{D,true}} - \hat{h}_{\text{true}}) \tag{36}$$
$$= (x_{\text{D,true}} - h_{\text{true}} + \nu_{\text{sensor}}) - (x_{\text{D,true}} - h_{\text{true}} - \nu_{\text{map,true}}) \tag{37}$$
$$= \nu_{\text{sensor}} - \nu_{\text{map,true}} \tag{38}$$

At the false location,

$$e_{\text{false}} = z - \hat{z}_{\text{false}} \tag{39}$$
$$= (x_{\text{D,true}} - h_{\text{true}} + \nu_{\text{sensor}}) - (x_{\text{D,false}} - \hat{h}_{\text{false}}) \tag{40}$$
$$= (x_{\text{D,true}} - h_{\text{true}} + \nu_{\text{sensor}}) - (x_{\text{D,false}} - h_{\text{false}} - \nu_{\text{map,false}}) \tag{41}$$
$$= \nu_{\text{sensor}} + \delta_{\text{t-f}} - \nu_{\text{map,false}} \tag{42}$$

where the difference between the expected measurements of the true and false trajectories is $\delta_{\text{t-f}}$

$$\delta_{\text{t-f}} = (x_{\text{D,true}} - x_{\text{D,false}}) + (h_{\text{true}} - h_{\text{false}}) \tag{43}$$

Equations 35 and 39 can be interpreted as follows: for the true trajectory the expected difference between the measurements and the map is the sum of the noise in the measurements and the map, whereas at the false location there is also a contribution from the difference between the terrain profiles and error in the depth estimate. In the case of known depth, $\delta_{\text{t-f}}$ is the difference between terrain profiles, and as such is referred to as the terrain information. With these models, the false peak likelihood using Equation 34 can be evaluated as a function of the map noise, sensor noise, and the difference $\delta_{\text{t-f}}$.

$$p\left(\frac{L(x_{\text{false}})_{\text{standard}}}{L(x_{\text{true}})_{\text{standard}}} > k\right) = f\left(\delta_{\text{t-f}}, \sigma^2_{\text{map}}, \sigma^2_{\text{terrain}}, k\right) \tag{44}$$

Although the analysis can be done with any $\delta_{\text{t-f}}$ profile difference, it is the sum of the squared difference, $\delta^T_{\text{t-f}}\delta_{\text{t-f}}$ that enters into the equation. This term is therefore modeled as

$$\delta^T_{\text{t-f}}\delta_{\text{t-f}} \sim N\delta^2_{\text{rms}} \tag{45}$$

Where δ_{rms} is used to characterize the RMS terrain profile difference. The pdf of the likelihoood ratio can thus be written as a log-normal distribution,

$$p\left(\frac{L(x_{\text{false}})_{\text{standard}}}{L(x_{\text{true}})_{\text{standard}}}\right) \sim \ln\mathcal{N}\left(\mu_{\text{LR}}, \sigma^2_{\text{LR}}\right) \tag{46}$$

where the terms μ_{LR} and σ^2_{LR} are a function of δ_{rms}, σ^2_{map},$\sigma^2_{\text{terrain}}$, and the profile length N.

$$\mu_{\text{LR}} = \frac{-N\delta^2_{\text{rms}}}{2\left(\sigma^2_{\text{sensor}} + \sigma^2_{\text{map}}\right)} \tag{47}$$
$$\sigma^2_{\text{LR}} = \frac{N\left(\left(\sigma^2_{\text{sensor}} + \sigma^2_{\text{map}}\right)\left(\delta^2_{\text{rms}} + \sigma^2_{\text{map}}\right) + \sigma^2_{\text{sensor}}\sigma^2_{\text{map}}\right)}{\left(\sigma^2_{\text{sensor}} + \sigma^2_{\text{map}}\right)^2} \tag{48}$$

With this log-normal distribution, a more negative mean indicates a lower likelihood of a false peak, but a larger standard deviation results in a larger spread in possible weightings and higher likelihood of false peaks. A simple example is the case for flat terrain - as there is no information in the terrain, the mean stays at zero (eg, the true location is indistinguishable from the false location), however because of the growing variance there is will be a high likelihood of strong evidence in favor of either location. Thus, the likelihood of false peaks is determined by the growth of the mean relative to the standard deviation.

Calculating the likelihood of a mis-weighting using the standard TRN formula is therefore done with the cumulative distribution of a normal variable. For a mis-weighting greater than k, that returns:

$$p\left(\frac{L(x_{\text{false}})_{\text{standard}}}{L(x_{\text{true}})_{\text{standard}}} > k\right) = \Phi\left(\frac{\mu_{\text{LR}} - \ln(k)}{\sigma_{\text{LR}}}\right) \tag{49}$$

The bound on the likelihood of a false peak of size k is calculated by maximizing this probability. As δ_{rms}, σ^2_{map} and $\sigma^2_{\text{terrain}}$ are fixed, this is with respect to N, and is

$$\max_N p\left(\frac{L(x_{\text{false}})_{\text{standard}}}{L(x_{\text{true}})_{\text{standard}}} > k\right) = \Phi\left(-\sqrt{\frac{a_{\text{LR}}}{b_{\text{LR}}}}\sqrt{2\ln(k)}\right) \tag{50}$$

where

$$a_{\text{LR}} = \delta^2_{\text{rms}} \tag{51}$$
$$b_{\text{LR}} = \frac{\left(\sigma^2_{\text{sensor}} + \sigma^2_{\text{map}}\right)\left(\delta^2_{\text{rms}} + \sigma^2_{\text{map}}\right) + \sigma^2_{\text{sensor}}\sigma^2_{\text{map}}}{\sigma^2_{\text{sensor}} + \sigma^2_{\text{map}}} \tag{52}$$

and a_{LR} and b_{LR} are dependent only on the terrain information, map and sensor error.

Examining this equation shows two important features. The first is that, in the case where the map error is small with respect to the terrain information, σ^2_{map} is small with respect to δ^2_{rms}, $a_{\text{LR}} \approx b_{\text{LR}}$ and the bound on false peaks matches the robust bound. The second is that the robust bound will be significantly broken in flat terrain, when σ^2_{map} is large with respect to δ^2_{rms}. This makes a_{LR} small compared to b_{LR}, and consequently the likelihood of false peaks will be much higher. The broken bound shown in Figure 10 is for the case where σ^2_m,σ^2_m are 1 and δ^2_{rms} is 0.5.

The general conclusion is that the lower the amount of terrain information relative to map error, the further to the right the false peak bound is pushed. That is, the likelihood of large false peaks is increased. This supports the simulation results from Section III-B, where there were many large peaks in uninformative terrain.

The robust bound on false peaks can be restored in uninformative terrain through the parameter α, which changes the mean and variance of the likelihood ratio. Noting the adjusted likelihood ratio as

$$L(x)_{\text{adjusted}} = L(x)^{\alpha}_{\text{standard}} = p(\boldsymbol{z}, \hat{h}|x)^{\alpha}_{\text{standard}} \tag{53}$$

the mean and variance of the log-normal likelihood ratio become

$$\mu_{\text{LR, adjusted}} = \alpha\mu_{\text{LR}} \tag{54}$$
$$\sigma^2_{\text{LR, adjusted}} = \alpha^2\sigma^2_{\text{LR}} \tag{55}$$

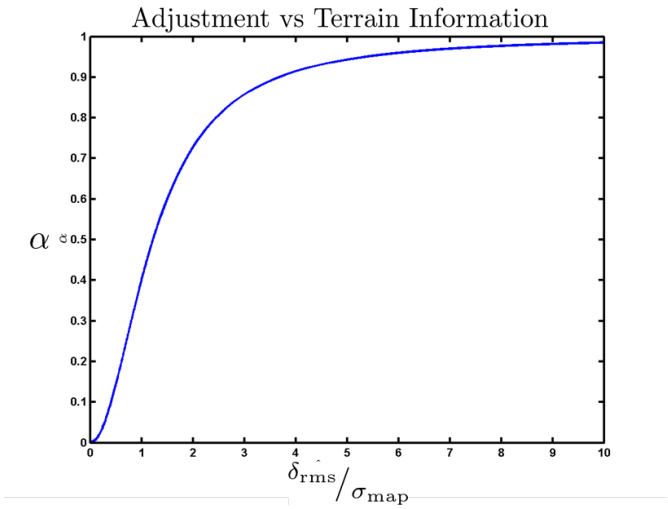

Fig. 11: Adjustment α vs $\frac{\hat{\delta}_{\text{rms}}}{\sigma_{\text{map}}}$, for the case of $\sigma_{\text{map}} = \sigma_{\text{sensor}}$.

and the probability of false peaks is

$$p\left(\frac{L(x_{\text{false}})_{\text{adjusted}}}{L(x_{\text{true}})_{\text{adjusted}}} > k\right) = \Phi\left(\frac{\alpha\mu_{\text{LR}} - \ln(k)}{\alpha\sigma_{\text{LR}}}\right) \tag{56}$$

Selecting α is based on restoring the robust bound, and results in

$$\alpha = \frac{a_{\text{LR}}}{b_{\text{LR}}} \tag{57}$$

which can be written as

$$\alpha = \frac{\delta^2_{\text{rms}}\left(\sigma_{\text{s}}^2 + \sigma_{\text{m}}^2\right)}{\left(\sigma_s^2 + \sigma_m^2\right)\left(\delta^2_{\text{rms}} + \sigma_m^2\right) + \sigma_s^2\sigma_m^2} \tag{58}$$

The parameter α is a function of the terrain information, map and sensor noise and is lower bounded by 0 and upper bounded by 1. In the case of perfect maps, $\sigma_m^2 = 0$ and α is 1; no adjustment is necessary. If there is map error, the amount of information does matter, and the scaling of α is largely a function of the ratio of terrain information to map error. The behavior of α vs terrain information is shown in Figure 11 for the case of equal map and sensor. For uninformative terrain, the ratio of terrain information to map error tends to zero, and α must tend to zero to make the approximation robust.

C. Real-time TRN Filter Implementation

The robust, adaptive technique was developed for use in real-time TRN filters in [12]. This method calculates $\alpha_{i,k}$ - the adjustment to apply to each beam at each timestep based on the expected information, $\hat{\delta}^2_{i,k,\text{rms}}$ for each beam. The term $\beta_{i,k}$ in Equation 11 becomes

$$\beta_{i,k} = \frac{\alpha_{i,k}}{\sigma^2_{i,k,\text{sensor}} + \sigma^2_{\text{map}}} \tag{59}$$

where the adjustment depends is taken from Equation 58

$$\alpha_{i,k} = \frac{(\sigma^2_{i,k,\text{sensor}} + \sigma^2_{\text{map}})\hat{\delta}^2_{i,k,\text{rms}}}{(\sigma^2_{i,k,\text{sensor}} + \sigma^2_{\text{map}})(\hat{\delta}^2_{i,k,\text{rms}} + \sigma^2_{\text{map}}) + \sigma^2_{i,k,\text{sensor}}\sigma^2_{\text{map}}} \tag{60}$$

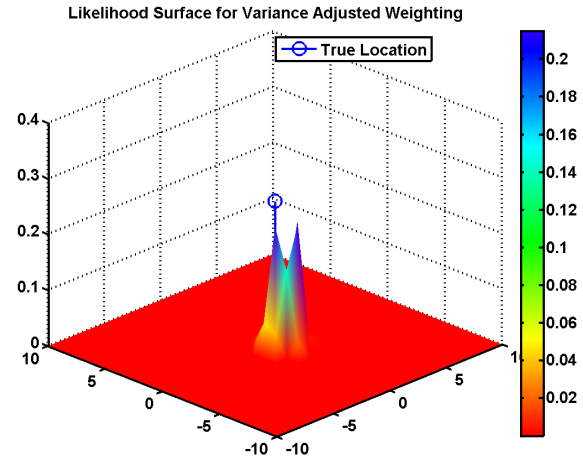

Fig. 12: Likelihood surface using adaptive weighting on informative terrain

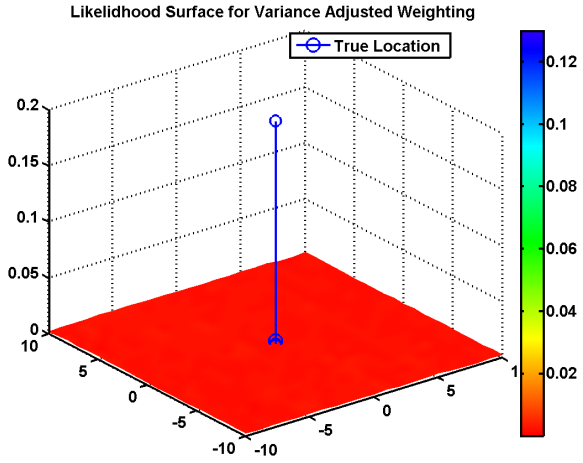

Fig. 13: Likelihood surface using adaptive weighting on uninformative terrain

and $\hat{\delta}^2_{i,k,\text{rms}}$ is the estimate of the terrain information

$$\hat{\delta}^2_{i,k,\text{rms}} = \text{Var}(\hat{z}_{i,k}) - \sigma^2_{\text{map}} \tag{61}$$

D. Adaptive TRN in Simulated Terrain

The efficacy of the adaptive algorithm is demonstrated using the simulated trajectories and maps as in Section III-B.

The result of using the adaptive filter in informative terrain is shown in Figure 12. In informative terrain, the filter produces a result similar to the standard filter - a peak at the correct location, as the higher terrain information pushes α towards 1.

The result of using the filter in uninformative terrain is shown in Figure 13. Here, the adaptive filter eliminates the false peaks that arise using standard TRN weighting, and yields a low, broad distribution similar to the Gaussian filter

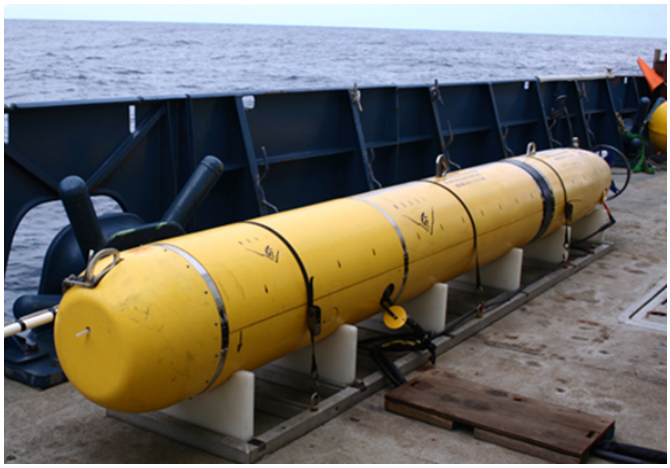

Fig. 14: MBARI Dorado-class Mapping AUV

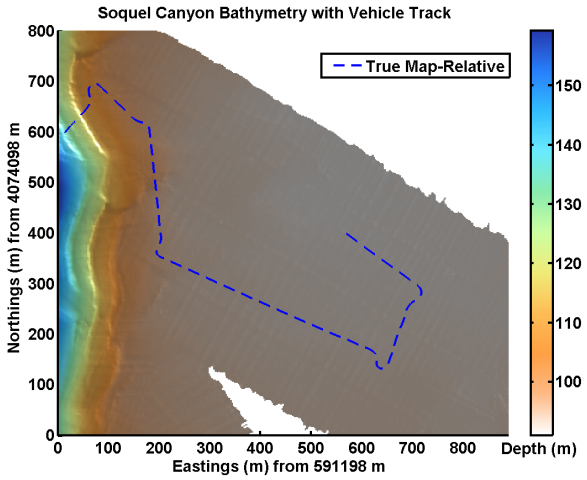

Fig. 15: AUV trajectory plotted over Soquel Canyon map. The AUV starts at the right and drives over flat terrain for ∼1km before entering the canyon.

using the known terrain model, as the low terrain information pushes α towards zero.

V. Offline Demonstration on MBARI Dorado Class AUV

The performance of the adaptive and standard TRN algorithms are compared using data taken by the MBARI Dorado-Class mapping AUV shown in Figure 14 at Soquel Canyon in Monterey Bay. The geo-referencing error of this map is known, and the 'true map-relative' trajectories incorporate this offset.

The test trajectory is shown in Figure 15. The AUV descends from the surface to 70m depth, then makes a 1km traverse across flat terrain of 100m average depth before dropping into Soquel Canyon.

The offline simulations use a particle filter with 4000 particles initialized over a 200m by 200m region at the start of the trajectory. Figures 16 and 17 show typical results of using a standard TRN algorithm and the modified TRN algorithm,

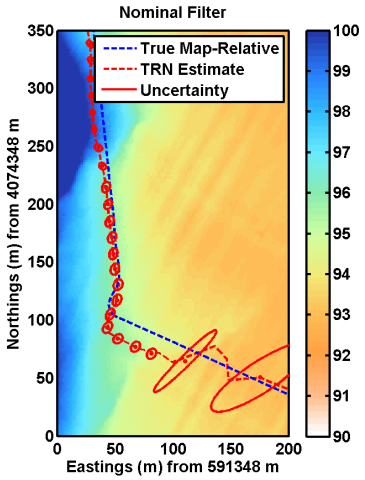

Fig. 16: Offline results for TRN filter using standard weighting. The filter converges to a tight peak with a ∼20m offset before the vehicle reaches informative terrain.

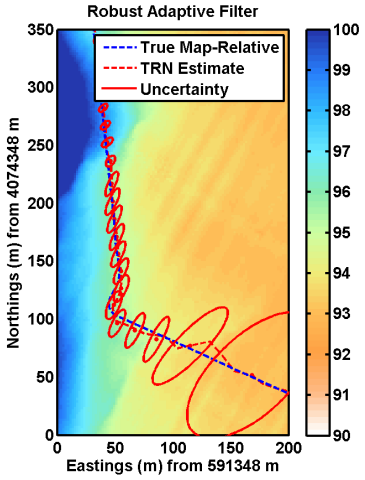

Fig. 17: Offline results for TRN filter using adaptive weighting. The adaptive filter maintains a broader distribution but converges tightly to the true position when the vehicle reaches informative terrain.

respectively. These plots are close-ups on the left side of the map, to demonstrate the filter performance when the vehicle reaches the dropoff into Soquel Canyon after flying over flat terrain. The plots show the evolution of the mean in dashed red and uncertainty in red ellipses relative to the true trajectory. Although the simulations use a particle filter, ellipses are chosen to represent the position uncertainty for the sake of clarity.

TRN results using the standard weighting, shown in Figure 16 , demonstrate false convergence over uninformative terrain. By the time the vehicle reaches informative terrain the filter has converged to an incorrect location with ∼20m offset.

TRN results using the adaptive weighting, shown in Figure 17, show slower convergence in flat terrain. Upon entering informative terrain the adaptive filter converges rapidly to the correct location with an accuracy of a few meters.

VI. Conclusion

The presented analysis demonstrates the standard TRN estimators implicitly assume that the vehicle is operating in informative terrain. This assumption can cause the filter to converge to the wrong location when the vehicle is operating in uninformative terrain.

Adjusting the filter measurement weighting to meet a robust bound on the likelihood of false peaks presents an effective means of preventing false peaks in flat terrain.

Acknowledgment

The authors would like to thank the Monterey Bay Aquarium Research Institute, the Stanford Graduate Fellowship, the National Defense Science and Engineering Graduate Fellowship, and the NASA ASTEP Grant NNX11AR62G for support of this work. Researchers at MBARI who were particularly invaluable in acquiring the presented data include Brett Hobson, Hans Thomas, Erik Martin, Rob McEwen and Dave Caress.

References

[1] J. P. Golden, "Terrain contour matching (TERCOM): A cruise missile guidance aid," *Proc. SPIE*, vol. 0238, pp. 10–18, 1980.

[2] O. Bergem, "Bathymetric navigation of autonomous underwater vehicles using a multibeam sonar and a Kalman filter with relative measurement covariance matrices," 1993.

[3] D. E. Di Massa, *Terrain-relative navigation for autonomous underwater vehicles*. PhD thesis, Massachusetts Institute of Technology, 1997.

[4] K. Anonsen, O. Hagen, O. Hegrenaes, and P. Hagen, "The HUGIN AUV terrain navigation module," in *Oceans - San Diego, 2013*, pp. 1–8, Sept 2013.

[5] S. E. Houts, S. G. Dektor, and S. M. Rock, "A robust framework for failure detection and recovery for terrain-relative navigation," in *Unmanned Untethered Submersible Technology 2013*, (Portsmouth, NH), 09/2013 2013.

[6] D. Meduna, S. M. Rock, and R. McEwen, "AUV terrain relative navigation using coarse maps," in *Proceedings of the 2009 Unmanned Untethered Submersible Technology Conference* (R. Blidberg, ed.), (Durham, NH), August 2009.

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

[8] 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.

[9] N. Bergman, "A Bayesian approach to terrain-aided navigation," 1997.

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

[11] I. Nygren, *Terrain Navigation for Underwater Vehicles*. PhD thesis, Royal Institute of Technology (KTH), 2005.

[12] S. Dektor and S. Rock, "Improving robustness of terrain-relative navigation for AUVs in regions with flat terrain," in *Autonomous Underwater Vehicles (AUV), 2012 IEEE/OES*, pp. 1 –7, September 2012.

[13] 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.

[14] 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.

[15] D. Meduna, S. M. Rock, and R. McEwen, "Closed-loop terrain relative navigation for AUVs with non-inertial grade navigation sensors," in *IEEE-OES Autonomous Underwater Vehicles Conference (AUV)*, (Monterey, CA), September 2010.

[16] N. A. Cressie, *Statistics for spatial data*, vol. 900. J. Wiley, 1993.

[17] R. Royall, "On the probability of observing misleading statistical evidence," *Journal of the American Statistical Association*, vol. 95, no. 451, pp. pp. 760–768, 2000.