Return-to-site of an AUV using Terrain-Relative Navigation: Field Trials

Stephen M. Rock[*†] **Brett Hobson**[†] **Sarah E. Houts**[*]
[*]Stanford University
Stanford, CA 94305, USA
[†]Monterey Bay Aquarium Research Institute
7700 Sandholdt Road, Moss Landing, CA 95039-9644
rock@stanford.edu, hobson@mbari.org, shouts@stanford.edu

Abstract—**This paper presents the results from recent field trials in Monterey Bay, CA which demonstrate the use of Terrain-Relative Navigation (TRN) to perform return-to-site missions using a 53cm diameter Dorado-class Autonomous Underwater Vehicle (AUV). For these tests, a series of targets were identified on a bathymetry map that were recognizable features (e.g. large boulders and outcroppings), and a trajectory was constructed for the AUV to fly over those targets (i.e. the latitudes and longitudes as defined on the map). Multiple missions were flown using a real-time TRN algorithm to identify and correct for georeferencing errors in the bathymetry map. The results were repeatable. They demonstrated that there was a 20m georeferencing error in the map. Without TRN, the targets were missed. With TRN, the AUV flew directly over the targets. Performance was independently validated using acoustic imagery from an on-board mapping multibeam sonar. These field results are presented and described. The TRN algorithm is also discussed.**

I. Introduction

Two return-to-site missions for an Autonomous Underwater Vehicle (AUV) motivate this paper. The first is a mission to discover and then monitor sites in the deep ocean for change. The second is to map and then monitor sites on a free-drifting iceberg in the Antarctic. Developing the technology to complete both of these missions is supported under a NASA ASTEP grant #NNX11AR62G as part of an effort to enable robotic exploration for life in extreme environments.

The focus of the work reported here is on the navigation problem. Specifically, the goal is to develop and demonstrate a technique that does not exploit any infrastructure (such as GPS or long-baseline arrays) and is able to function in environments in which traditional inertial navigation aids may not be sufficient.

The approach developed here is Terrain-Relative Navigation (TRN). In this approach the primary source of position information is a 3-D map of the terrain. Given this map, terrain profiles constructed from range measurements made by the vehicle during flight can be correlated against the map to determine a location. For the AUV missions dealt with in this paper, that map is a pre-existing bathymetric map of the area. For missions such as a return to site on an iceberg, that map must first be created.

Creating the maps for TRN can be done as either an off-line or on-line Simultaneous Localization and Mapping problem (SLAM). In an off-line approach, an AUV flies a pattern over the terrain collecting the multibeam sonar data needed to construct the map (i.e. a 3-D representation). For deep-sea missions this could be a "lawnmower" pattern over an area of the seafloor. For iceberg missions, this could be a series of circumnavigations at multiple depths. Note that control of the AUV during this stage relies on reactive and pre-programmed behaviors. After collecting those data, the map is constructed in an off-line, batch optimization. This is the method typically used to construct current seafloor bathymetry maps using an AUV.

In an on-line SLAM approach, the above two steps are fused into a single process and the AUV control can become a function of its location in the map. Depending on the mission, either approach can be appropriate. For missions where the mapping data can be collected safely using reactive control strategies, the off-line approach is appropriate. When this is not possible, an on-line

978-1-4799-4344-9/14/$31.00 ©2014 IEEE

approach may be more appropriate.

Current work focused on developing maps of free-drifting icebergs using an off-line SLAM method is reported in Kimball [1] and Hammond [2].

II. Terrain-Relative Navigation

Terrain-Relative Navigation (TRN) is a technique that enables an AUV to estimate its position with respect to an existing map of a terrain (e.g. bathymetric map of the seafloor or 3-D reconstruction of an iceberg). The method is based on correlating a set of range measurements obtained as the vehicle flies over the terrain against their expected values as defined by the map. The method can be used to eliminate drift in an on-board inertial navigation solution, offering an alternative to surfacing for a GPS fix. It can also be used to provide a direct measurement of the AUV's position in map coordinates instead of inertial coordinates (latitude and longitude). This latter option is the motivation for this work. It enables return-to-site missions that are robust to map georegistration errors as well as to any drop-outs or drift in the inertial navigation solution.

Several groups have explored the use of acoustic-based TRN for AUV applications in addition to the work reported here. Most of that work, however, has focused on using TRN as an inertial navigation aid rather than for return-to-site missions. Examples include demonstrations on the Swedish AUV62f [3] [4] and on the HUGIN AUV [5] [6]. The work that is most similar is by Anonsen and Hagen [5] in which sea trials are reported that demonstrated an estimate of inertial position accurate to within 5m over a 5 hour run with the HUGIN AUV.

Basic TRN algorithm:

The form of the TRN estimator developed here is a particle filter. This type of estimator was chosen because of its inherent ability to deal with multimodal likelihood distributions while still providing high resolution in the converged solution. Effectively, it is real-time Monte-Carlo approach that propagates a series of hypotheses (i.e. position estimates) and selects as truth a weighted estimate which best matches the measured sonar ranges over time.

A conceptual block diagram describing its implementation and operation is presented in Figure 1.

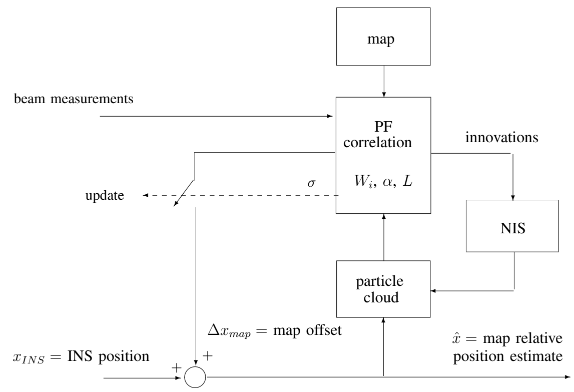

Fig. 1. Block diagram of the TRN algorithm including the NIS innovations check

The core of the algorithm is a collection of particles, each of which represents an hypothesis of the AUV's position (the particle cloud). At each time step a process update and a measurement update are performed for every particle. In the process update, the position estimate associated with each particle is moved by the measured change reported by the Inertial Navigation System (INS) plus noise. That is

$$x_i(n) = x_i(n-1) + \Delta x_{INS} + w \qquad (1)$$

where x_i is the location propagated in the i^{th} particle (Northings, Eastings and depth); x_{INS} is the location reported by the INS; $\Delta x_{INS} = x_{INS}(n) - x_{INS}(n-1)$ and $w = \mathcal{N}(0, \sigma^2_{INS})$.

In the measurement update, the measured terrain profile (e.g. from the DVL or sonar beam lengths), is compared against the profile determined from the map and the location of each particle's position estimate, x_i, is updated. This correlation is done using a likelihood function. The standard function typically used in the literature is

$$L(x)_{typ} = \eta \exp\left(\frac{-(z-\hat{z})^2}{2(\sigma^2_{map} + \sigma^2_{sensor})}\right) \qquad (2)$$

where z are the measured beam ranges and $\hat{z}$ are the ranges predicted at that assumed map location. For this calculation, the map is assumed to have an error that is $\mathcal{N}(0, \sigma^2_{map})$; the range measurements are assumed to have an error that is $\mathcal{N}(0, \sigma^2_{sensor})$; and η is a normalizing factor.

Computing $L(x)$ for every particle results in a likelihood distribution which is used to compute weights on each of the particles, W_i. Weights are increased on particles with high likelihood (i.e. the measured profile agrees with the map at position x_i) and decreased on particles with low likelihood. Given these weights, a best estimate of the AUV's position with respect to the map is computed as a weighted average of the x_i. Also computed is a covariance associated with that position (σ^2).

The particle cloud is initialized to cover a large area of the terrain in order ensure that the true location of the AUV is included. However, as time progresses, resampling of the cloud allows particles associated with rejected hypotheses to be repositioned in regions determined to be more likely locations. This enables the estimator to improve the resolution of its final, converged answer.

At every time step, a map offset is calculated as

$$\Delta x_{map} = \hat{x} - x_{INS}. \tag{3}$$

where $\hat{x}$ is the best estimate of the AUV's position in the map frame.

Finally, if the covariance associated with the estimate is low, the estimator is assumed to be converged and the current map offset is added to the INS position to yield the best estimate of the AUV's position. If the covariance is high, the calculated map offset is ignored and the last good value is used for the update.

The filter implemented for the tests described here used 10,000 particles and updated the estimate of map offset every 3 seconds. The accuracy achieved with this method has been approximately 2 to 3m in previous field trials when using 1-m level bathymetry maps. Further details on the basic TRN estimator can be found in [7] and [8].

TRN algorithm enhancements:

Two augmentations have been made to the TRN algorithm described above in order to improve its robustness. Both have been implemented in response to behaviors observed in previous field trials.

The first augmentation was the addition of a reasonableness check based on the innovations in the filter. This test was added because the filter could occasionally be in a converged state at an incorrect location yet have a high confidence (i.e. small σ) and tight particle cloud. For example, this could occur when crossing discontinuities in the map (e.g. artifacts associated with creating a large map from multiple submaps). If undetected, this situation could cause the filter to become "lost" and unable to re-converge.

To detect and correct for this situation, a Normalized Innovation Squared (NIS) test has been implemented. The test monitors the innovations sequence and declares a fault if the covariance of the normalized innovations exceeds a threshold. When this occurs, the filter is reinitialized. Details are provided in [9].

The second augmentation was a modification to the likelihood function used in the correlation step, $L(x)_{typ}$. This modification was made in response to the observation that TRN can occasionally converge to an incorrect position estimate and assign a high probability of confidence to that answer when operating over benign terrain. The reason for this incorrect convergence is a result of the way in which TRN estimators are implemented using particle filters. Several ad hoc solutions are available in the literature, but [10] presents a first-principles approach to a solution. Simply stated, the standard likelihood function, $L(x)_{typ}$, implicitly assumes that the noise/uncertainty in the map is small with respect to the amount of variability in the terrain. However, this condition is not satisfied when the terrain is benign (e.g. flat). Hence, when flying over benign terrain, a different likelihood function needs to be used.

The new likelihood function is

$$L(x)_{mod} = L(x)^{\alpha}_{typ} \tag{4}$$

where the paramenter, α, is calculated based on the statistical properties of the local terrain map. Details are presented in [10].

III. Vehicle Description

The AUV used for testing and development is a variant of MBARI's standard mapping vehicle. It is a Dorado/Bluefin 21" (54 cm) diameter type,

Fig. 2. The MBARI modified Mapping AUV used for the return-to-site tests

which has a cylindrical midsection and a truncated hemispherical nose.

For these tests, a mapping sonar (Reson 7125) was mounted in a down-look configuration and perpendicular to the direction of travel. This sonar was included to enable an independent verification of the performance of the navigation solution. Two Imagenex 837A multibeam sonars were also included, but only used for obstacle avoidance. One was mounted in a down-look configuration and parallel to the direction of travel. The other was mounted in a forward-look configuration and parallel to the direction of travel.

The beam ranges used in the TRN estimator were taken only from the DVL for these tests. Any or all of the multibeam sonars could also be included, but the goal was to demonstrate a successful mission using a minimal sensor suite.

IV. Field Trials

In May 2014, a demonstration of a return-to-site AUV mission was performed at Portuguese Ledge in Monterey Bay, CA (Figure 3). This site is approximately 800m in diameter and contains a series of rock formations rising 30m above the seafloor which are surrounded by a large area of featureless terrain. This mix of informative and uninformative terrain was ideal for testing the TRN estimation logic under varying conditions. Specifically, it made it possible to observe the behavior of the logic both in its converged and un-converged states as well as during the transitions between the two.

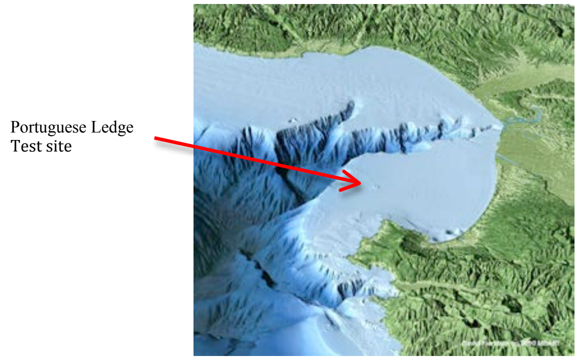

Fig. 3. Portuguese Ledge Test Site

Four runs comprised the test. Each was defined by a waypoint at the end of three path segments, and each waypoint was defined as a latitude-longitude in an existing bathymetric map of the site. Each run began at the Western most waypoint and proceeded East, South and then North. A nominal trackline is presented in Figure 4. Note that each commanded path was similar, but not identical.

Importantly, the map used to define the waypoints for this mission had a known georeferencing error of approximately 20m.

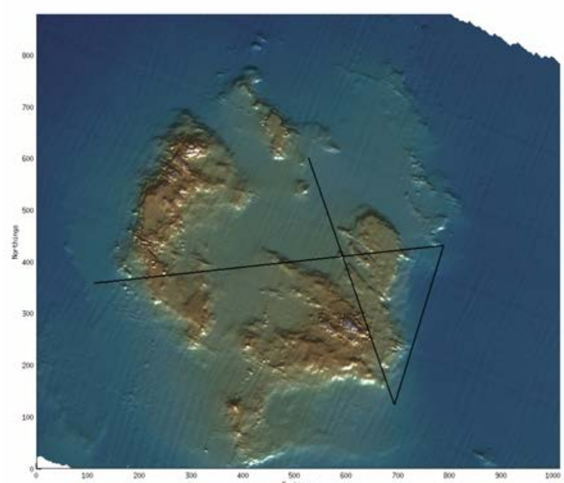

Fig. 4. The commanded trackline specified on the original map

TRN performance:

TRN was active during all of these runs. That is, a map offset was calculated in real time and

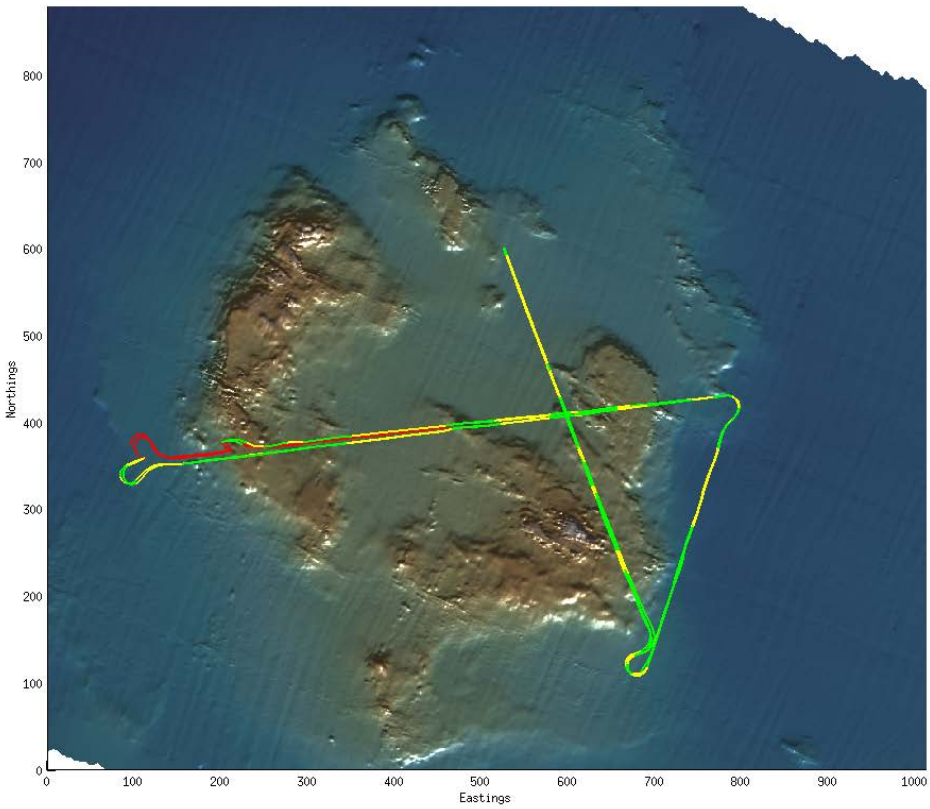

Fig. 5. Corrected tracklines for the four tests. For each trackline, green indicates that the TRN estimator is converged and the map offsets are being updated; red indicates that the estimator is not converged; and yellow indicates that estimator has modified the likelihood function in response to reduced information in the terrain

used to correct the estimated location of the AUV with respect to the map as described above. This map-relative position was the signal used by the tracking control system to control the AUV's location. Consequently, if the offset calculated by the TRN logic was correct, the paths actually flown would match the commanded paths when viewed in the frame of the map.

Results for the four test runs are shown in Figures 5, 6 and 7. Figure 5 presents the tracklines displayed on the map after adjusting for the georeferencing errors. These paths represent the best estimate of the trajectories flown over the true terrain. Figure 6 presents a zoomed in view of a section of one of the paths. The dashed line is the location of the commanded path without the TRN adjustment applied. The solid line is the path after the correction. The 20m offset between the two sets of plots is clearly visible. (See Figure 11 for additional views.)

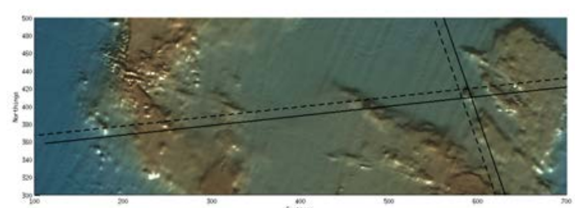

Fig. 6. Zoomed in view showing the offset between the commanded path before and after adjustment for Run #3

Figure 7 displays the map offset calculated by the TRN logic for each of the four runs. Also plotted is the 90% confidence band reported by the logic for the Run #3. These results show both the consistency of the estimated offset as well as the accuracy of the confidence interval (i.e. σ).

As an independent check, data from the mapping sonar collected during this run was used to generate a map of the terrain using the tools

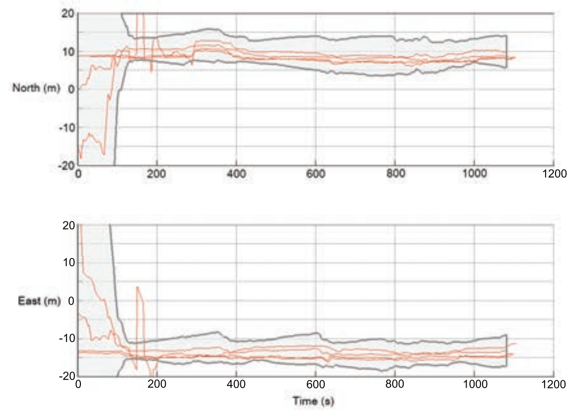

Fig. 7. Map offset calculated by the TRN algorithm in four tests. The grey area is the 90% confidence bound calculated during Run #3

available in the MBSystem software package. This is an offline process that creates a map as well as a best estimate of the AUV's trajectories with respect to this map based on measured multibeam sonar and INS data The result of this process is presented in Figure 8. A comparison of Figures 5 and 8 indicates qualitative agreement between the real-time TRN results and those obtained using the mapping sonar data. Figures 9 and 10 provide a zoomed in view of the regions where the trajectories cross.

Additional insight into the operation of the TRN estimator is also provided in Figure 5. The solid line in the figure indicates the path computed by the on-line estimator. When the line is green, the filter is converged and the map offset calculated by the estimator is updated continuously. When the line is red, the TRN filter is not converged and the last good map offset is used. This occurs for the two trajectories that approach the first waypoint over benign terrain from the South. (The trajectories that approach from the North arrive at the first waypoint converged since they have

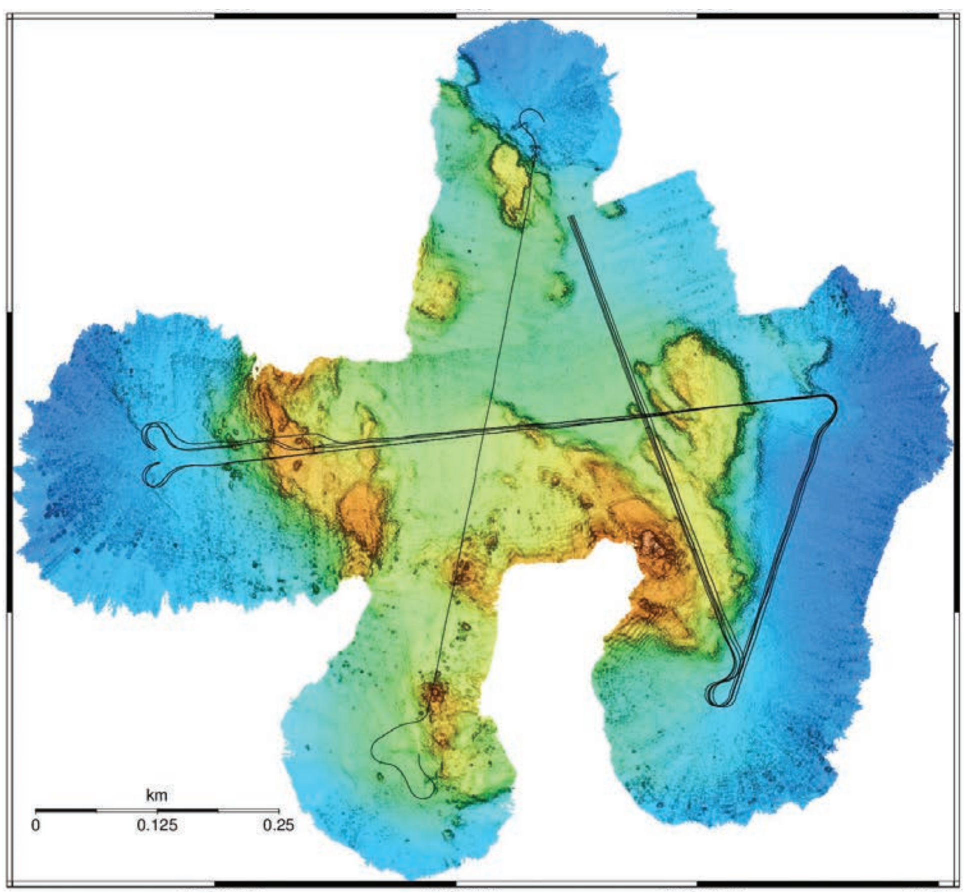

Fig. 8. Map and trajectories calculated from multibeam sonar data using MB-System

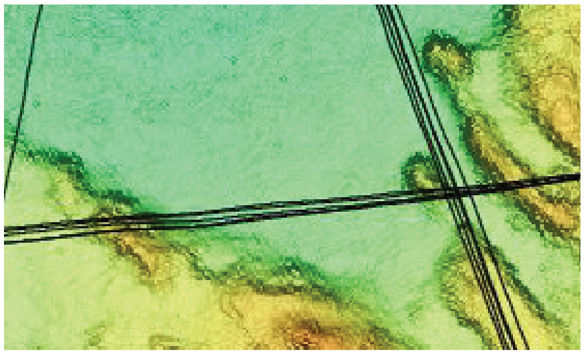

Fig. 9. Zoomed in view of the trajectories calculated from the multibeam sonar data

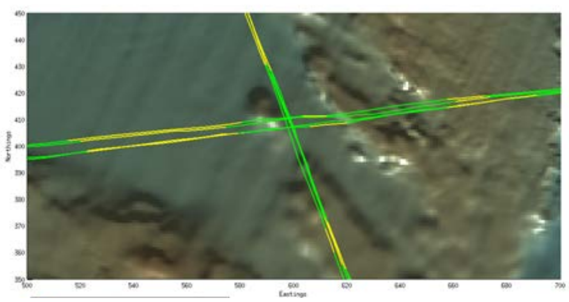

Fig. 10. Zoomed in view of the trajectories calculated by the TRN estimator

travelled over feature rich terrain.)

When the line is yellow, the estimator has determined that there is reduced information in the measurements (i.e. operating over benign terrain). During this time, the filter propagates normally, but uses the modified correlation function, $L(x)_{mod}$, in the correlation.

Return-to-site performance:

The above results also demonstrate a successful return-to-site. Specifically, the waypoints that define the desired trajectory for Run #3 over the map were chosen to cause the AUV to fly over the lower corner of a boulder. This is indicated in the upper part of Figure 11. The path actually flown is shown in the lower part of the figure.

In the figure, the dashed black line indicates where the requested trackline appears in the map without accounting for georeferencing errors. This is the path that would have been flown without TRN. The 20m offset is clearly visible and, if flown, the boulder would have been missed. The solid line is the path flown using the TRN estimate of map offset.

Note that the solid line appears to start by tracking the desired path. However, the TRN estimator has not yet converged (red) in this region and the AUV is unaware that the map is not correctly georeferenced. When the estimator converges, the estimate of the AUV's position jumps North and West to the true position on map. The control system then drives this error to zero causing the AUV to track the true commanded path.

V. Conclusion

The field tests presented here demonstrate the feasibility of using Terrain-Relative Navigation as an effective augmentation to estimating the position of an AUV. The motivation for this work was enabling return-to-site missions for an AUV, for example, to monitor a site for change over time. These missions require that an AUV visit targets identified in the map's frame rather than in an inertial frame. That is, the navigation system must be able to deal with (potentially large) map georeferencing errors (or even a moving target such as a free-drifting icebergs) as well as offsets that may occur in the INS estimate of position (e.g. drift). This capability could also be used as an alternative to surfacing for GPS navigation fixes to correct for the drift that naturally occurs in any dead-reckoned navigation technique assuming the map is adequately georeferenced.

In addition, these field tests demonstrate robust performance of the TRN system when transitioning between flight over information-rich and benign (flat) terrains. Specifically, a modification to the likelihood function used in the measurement update improves the prediction of the uncertainty in the position estimate while enabling smooth transitions between the regions.

VI. Acknowledgements

The work presented above was funded by NASA ASTEP Grant #NNX11AR62G with additional support form the Monterey Bay Aquarium Research Institute. Thanks go to Hans Thomas and his team plus the crew of the *R/V Rachel Carson* for supporting AUV operations, to Dave Caress for analysis of the mapping sonar data, to Rob McEwen for help with the AUV control

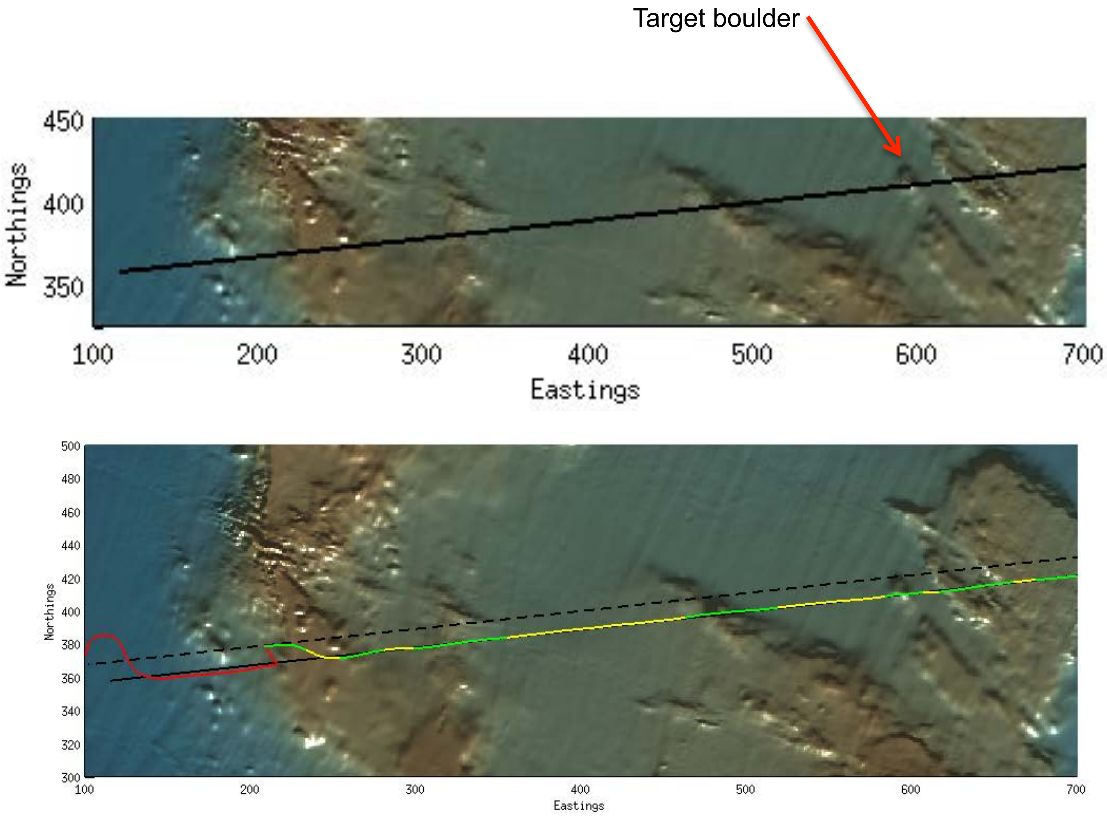

Fig. 11. Upper view: The commanded path on the map. The target (lower edge of a boulder) is indicated by the red arrow in the upper view. Lower View: Zoomed in view of the trajectory flown by the AUV. The dashed line is the location of the commanded path before correction.

system, to Rich Henthorn for software support and to Sarah Houts, Steve Krukowski, Shandor Dektor and David Stonestrom for data analysis.

REFERENCES

[1] P. Kimball and S. M. Rock, “Mapping of translating, rotating icebergs with an autonomous underwater vehicle,” *IEEE Journal of Ocean Engineering*, vol. PP, no. 99, pp. 1–13, February 2014.

[2] M. H. Hammond and S. M. Rock, “Enabling AUV mapping of free-drifting icebergs without external navigation aids,” in *Proceeding of the MTS/EDDD OCEANS Conference*, St. Johns, Newfoundland, September 2014.

[3] J. Carlstrom and I. Nygren, “Terrain navigation of the Swedish AUF62f vehicle,” in *UUST*, Durham, NH, August 2005.

[4] J. Carlstrom, “Results from sea trials of the Swedish AUV62f’s terrain navigation system,” in *Proceedings of the Unmanned Untethered Submersible Technology Conference*, Durham, NH, August 2007.

[5] K. Anonsen and O. Hagen, “An analysis of real-time terrain aided navigation results from a HUGIN AUV,” in *Proceeding of the MTS/EDDD OCEANS Conference*, September 2010, pp. 1–9.

[6] K. Anonsen, O. Hagen, O. Hegrenaes, and P. Hagen, “The HUGIN AUV terrain navigation module,” in *Proceeding of the MTS/EDDD OCEANS Conference*, San Diego, September 2013, pp. 1–8.

[7] D. Meduna, S. M. Rock, and R. McEwen, “AUV terrain relative navigation using coarse maps,” in *Proceedings of the Unmanned Untethered Submersible Technology Conference*, Durham, NH, August 2009.

[8] D. Meduna, S. Rock, and R. McEwen, “Closed-loop terrain relative navigation for AUVs with non-inertial grade navigation sensors,” in *Proceedings of the IEEE AUV Conference*, Monterey, CA, September 2010.

[9] S. E. Houts, S. G. Dektor, and S. M. Rock, “A robust framework for failure detection and recovery for terrain-relative navigation,” in *Proceedings of the Unmanned Untethered Submersible Technology Conference*, Portsmouth, NH, September 2013.

[10] S. G. Dektor and S. M. Rock, “Robust adaptive terrain relative navigation,” in *Proceeding of the MTS/EDDD OCEANS Conference*, St. Johns, Newfoundland, September 2014.