

Terrain-Aided Navigation for Long-Endurance and Deep-Rated Autonomous Underwater Vehicles

Georgios Salavasidis

Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
geosal@noc.ac.uk

Andrea Munafò

Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
andmun@noc.ac.uk

Catherine A. Harris

Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
cathh@noc.ac.uk

Thomas Prampart

Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
thoram@noc.ac.uk

Robert Templeton

Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
robert.templeton@noc.ac.uk

Micheal Smart

Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
micsmart@noc.ac.uk

Daniel T. Roper

Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
danrope@noc.ac.uk

Miles Pebody

Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
miles.pebody@noc.ac.uk

Stephen D. McPhail

Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
sdm@noc.ac.uk

Eric Rogers

Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, UK
etar@ecs.soton.ac.uk

Alexander B. Phillips

Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
abp@noc.ac.uk

Abstract

Terrain-Aided Navigation (TAN) is a localisation method which uses bathymetric measurements for bounding the growth in inertial navigation error. The minimisation of navigation errors is of particular importance for long-endurance Autonomous Underwater Vehicles (AUVs). This type of AUV requires simple and effective on-board navigation solutions to

undertake long-range missions, operating for months rather than hours or days, without reliance on external support systems. Consequently, a suitable navigation solution has to fulfil two main requirements: a) bounding the navigation error, and b) conforming to energy constraints and conserving on-board power. This research proposes a low-complexity particle-filter based TAN algorithm for Autosub Long Range, a long-endurance deep-rated AUV. This is a light and tractable filter that can be implemented on-board in real-time. The potential of the algorithm is investigated by evaluating its performance using field-data from three deep (up to 3700 m) and long-range (up to 195 km in 77 hours) missions performed in the Southern Ocean during April 2017. The results obtained using TAN are compared to on-board estimates, computed via dead-reckoning, and USBL measurements, treated as baseline locations, sporadically recorded by a support ship. Results obtained through post-processing demonstrate that TAN has the potential to prolong underwater missions to a range of hundreds of kilometres without the need for intermittent surfacing to obtain GPS fixes. During each of the missions, the system performed 20 Monte Carlo runs. Throughout each run, the algorithm maintained convergence and bounded error, with high estimation repeatability achieved between all runs, despite the limited suite of localisation sensors.

List of Abbreviations

AABW	Antarctic Bottom Water	KF(s)	Kalman Filter(s)
ADCP	Acoustic Doppler Current Profiler	LBL	Long Baseline
ALR	Autosub Long Range	LL	Latitude/Longitude
ALR1500	Autosub Long Range 1500m-rated variant	MMSE	Minimum Mean Square Error
AUV(s)	Autonomous Underwater Vehicle(s)	MPF	Marginalized Particle Filter
AWDD	Absolute Water Depth Difference	MXX	Mission XX
BPF	Bootstrap Particle Filter	PMF	Point Mass Filter
BT	Bottom-Tracking	PF(s)	Particle Filter(s)
CRLB	Cramer-Rao Lower Bound	SBL	Short Baseline
CTD	Conductivity Temperature Depth	RBPF	Rao-Blackwellized Particle Filter
DEM	Digital Elevation Model	RMSE	Root Mean Square Error
DT	Distance Travelled	SIS	Sequential Importance Sampling
DR	Dead Reckoning	SISR	SIS with Resampling
DVL	Doppler Velocity Logger	SIR	Sampling Importance Resampling
DynOPO	Dynamics of the Orkney Passage Outflow	SMC	Sequential Monte-Carlo
EKF	Extended Kalman Filter	TAN	Terrain-Aided Navigation
GNSS	Global Navigation Satellite System	TBN	Terrain-Based Navigation
GPS	Global Positioning System	ToF	Time of Flight
IMU	Inertial Measurement Unit	WT	Water-Tracking
INS	Inertial Navigation System	USBL	Ultra Short Baseline

1 Introduction

The localisation problem for AUVs becomes increasingly challenging as vehicle endurance increases. With the recent developments of long-range and deep-rated AUVs, such as Autosub Long Range (ALR) (Furlong et al., 2012) (see Figure 1), capable of operating for several months and travelling thousands of kilometres in a single deployment, there is an increased need for dedicated, accurate and robust on-board localisation techniques. Such methodologies will improve AUV autonomy, allowing vehicles to undertake long-endurance missions without the need for human intervention. ALR and similar long-range platforms open up a world of new AUV applications, including persistent monitoring and data collection in some of the most remote and inaccessible areas on Earth. One potential application is crossing under-ice the Arctic Ocean from Svalbard to the Barrow coast in Alaska, a mission with range exceeding 3000km and duration of the order of months rather than days or hours (Salavasidis et al., 2018).

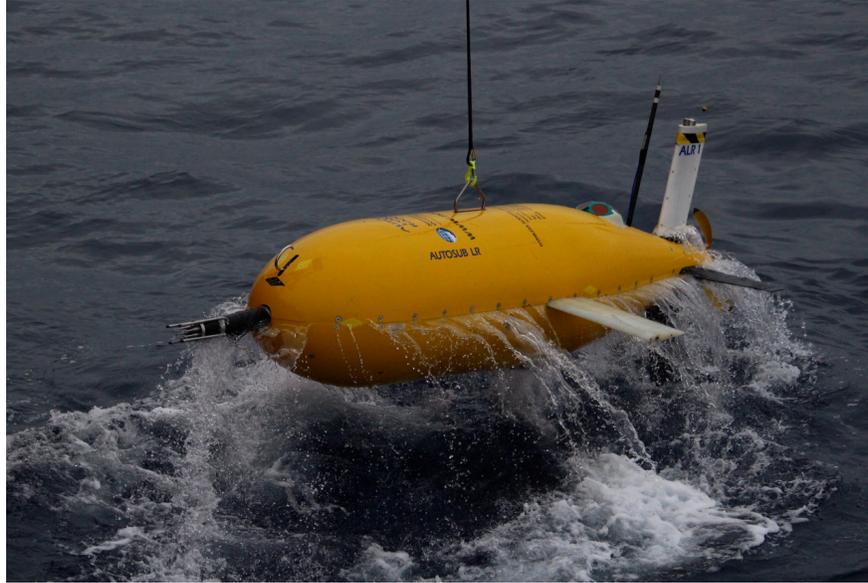


Figure 1: ALR recovery after a deep-water long-range mission in the Southern Ocean.

Long-range platforms are unable to fully exploit their endurance because of range limitations imposed by physical constraints (e.g. signal attenuation) experienced by many existing localisation methods. Typical AUV missions are of the scale of few kilometres within a pre-defined survey area. This is primarily because the AUV is required to remain within a constrained area, within range of an external support system (e.g. acoustic beacons) used to provide an effective external reference to bound the growth in inertial navigation positioning error (Paull et al., 2014). Alternatively, AUVs operating in open ocean may extend the operating range by surfacing regularly to obtain Global Navigation Satellite System (GNSS) fixes (e.g. using Global Positioning System - GPS) to bound the navigational drift. However, regular surfacing results in a trade-off between potential accuracy and time/energy consumption. This is an especially important consideration for deep-water missions and in areas where surface access is limited (e.g. when under ice or in regions of heavy ship traffic). Therefore, there is a demand for localisation methods, appropriate for prolonged operation in deep-water, which enable the AUV to navigate for thousands of kilometres without reliance on external position information. Such methods must rely entirely on on-board sensors, information available *a-priori* and efficient data processing.

Navigation methods which exploit geophysical information (e.g. bathymetry, magnetic and gravitational anomalies) are alternatives to widely-used acoustic techniques for bounding the inertial navigation error growth over time. However, in contrast to acoustic-aided methods, geophysical navigation can be fully implemented on-board the vehicle itself, without the need for external infrastructure. As a result, the use of these methods is particularly appealing for long-range AUV missions.

Terrain-Aided Navigation, also known as Terrain-Based Navigation (TBN), is a form of geophysical localisation where physical features of the seabed are exploited in order to localise an AUV within an *a-priori* Digital Elevation Model (DEM), or bathymetric reference map. To undertake this task, statistical estimators are used to generate estimates of the vehicle position provided a sequence of on-board water-depth measurements. Aerial military vehicles and land-attack cruise missiles have already demonstrated the great potential of terrain-aided navigation (Siouris, 2004; Ekütekin, 2007), supporting the conjecture that TAN can lead to the development of new generation underwater navigation systems. The accuracy of this method depends on the accuracy of motion sensors, the quality of available maps, the morphology of the seabed and the sensitivity of bathymetric sensors to changes in the AUV pose (Meduna, 2011).

Although TAN has drawn significant attention over the last two decades in underwater domain, real-world

implementations are limited in number and typically focus on short to very-short range applications (1-5 km), see (Carreno et al., 2010; Melo and Matos, 2017) for an overview of the current status. In contrast, this research develops a navigation method for extended time periods and deep-water operations. The TAN problem is formulated as a non-linear Bayesian estimation problem which is solved using Particle Filter (PF) estimator. The algorithm performance is evaluated using field-data collected by the ALR during deep long-range missions in the Southern Ocean (Naveira Garabato and Forryan, 2017).

1.1 Previous Research

Existing underwater localisation methodologies can be divided into three main categories (Paull et al., 2014) listed below:

- Inertial/Dead Reckoning
- Acoustic-based localisation
- Geophysical-based localisation

Each technique has its advantages and disadvantages. Although each technique can be used in isolation, the combination of multiple methods leads to higher estimation accuracy and robustness. Fusion and state estimation algorithms, such as a Kalman Filters (KF), PF and their variants, merge multiple techniques, resulting in a combined technique that exploits the strength of each component, whilst compensating for their weaknesses.

1.1.1 Dead-Reckoning

Dead reckoning/inertial navigation is a self-contained on-board localisation method where motion sensor measurements (e.g. from an Inertial Measurement Unit - IMU) are integrated with respect to time in order to propagate an initially known AUV position. Although this method can provide acceptable accuracy in short oceanographic missions, position estimates accumulate significant error (typically of the order of 1% of the distance travelled, although an expensive INS can even achieve a drift of 0.1% of the distance travelled (Paull et al., 2014)) due to sensor drift over time. To address this unbounded error growth, an external position reference is required. If this external information is efficiently coupled with the AUV's on-board sensors, a highly accurate positioning system can be formed (Paull et al., 2014).

1.1.2 Acoustic-based Localisation

One possibility for bounding the inertial navigation error growth is the use of underwater acoustics (Vickery, 1998). Acoustic-based support systems typically have coverage in the range of 100 m to 10 km with positioning accuracy from centimetres to few metres. However, acoustic signals can only propagate approximately 100km underwater with sufficient bandwidth for phase-coherent communication (Freitag et al., 2012; Caiti et al., 2010). The range primarily depends on the frequency of the signal, hence a trade-off must be made. High frequencies correspond to a higher propagation attenuation, but also increase positioning accuracy. Various acoustic-based localization methods have been developed over the years, although three schemes are the most frequently used: the Long Baseline (LBL), the Short Baseline (SBL) and the Ultra Short Baseline (USBL). The working principle of acoustic navigation lies in measuring bearing (in case of SBL and USBL) and the Time of Flight (ToF) of acoustic signals, which is then transformed to range measurements. For example, the standard mode of deploying AUVs is with a supporting system, e.g. a USBL head installed on a research ship hull (Mandt et al., 2001).

USBL is used to calculate the relative position between the ship-mounted transducers (USBL head) and the transponder on-board the AUV by acoustically measuring the range and bearing to the AUV. The relative

position can be easily extended to a global frame of reference, as long as the USBL head is globally referenced. If multiple acoustic transponders are available and installed prior to AUV deployment, then a LBL can be formed and used (Jakuba et al., 2008). In such systems, an AUV localises itself by multilaterating acoustic ranges obtained from globally-referenced transponders. Recent alternatives include: using an unmanned surface vehicle, which offers significant savings in expensive ship time (Salavasidis et al., 2016b; Pascoal et al., 2006; Pascoal et al., 2000); a moving LBL (Curcio et al., 2005); or network supported navigation systems (Munafò and Ferri, 2017). In principle, positioning could be also achieved using single range measurements, either from a fixed or moving beacon. However, in this particular case the observability issue should be considered and therefore range measurements from different directions must be obtained. A review of relevant background literature, including sensor technologies, is given in (Paull et al., 2014; Kinsey et al., 2006).

Although acoustic-based localisation is a well-established, accurate, and robust methodology, it does not align well with the challenges of long-range missions as acoustic-based support systems constrain the operational window of the AUV. In the LBL case, moreover, acoustic sensor networks are expensive and require considerable time and effort to deploy, survey, calibrate, and recover the beacons (resulting in further cost increases and delays until the mission can begin). If a surface aid is provided, then the operational cost for a dedicated assistive ship is high and impractical for some applications. Additionally, assistance from the surface might be also restricted by factors such as limited ship manoeuvrability or an inability to enter certain regions e.g. due to thick ice-cover. In order to prolong the mission duration, access remote areas and compensate for the inevitable inertial navigation drift, AUVs have to instead entirely rely on on-board sensors and efficient algorithms. Consequently, good candidates to bound the error growth are methods which use geophysical information (Teixeira, 2007; Teixeira et al., 2017b).

1.1.3 Geophysical-based Localisation

Geophysical navigation methods rely on observations of physical features (or feature variation) to acquire localising information. The measured geophysical properties of Earth could be used as references for localisation purposes. Given a map of the distribution of these properties, an algorithm can match sensor measurements with the reference map to determine the vehicle position within the map. A major advantage of these methods, when compared to acoustic-based techniques, is the lack of reliance on external signals and devices.

Given on-board bathymetric measurements and an *a-priori* bathymetric reference map, an AUV can be localised using terrain-aided navigation techniques. Various approaches have been developed to estimate the vehicle state. These approaches usually fall in two estimation classes: batch (correlation) and Bayesian (sequential) methods. Batch approaches periodically store terrain measurements and once a specific set of readings is collected, measurements are converted to a uniquely determined vehicle position by correlating the water-depth measurements with the reference map. The limitation of these methods, as stated in (Meduna, 2011), is that they do not take into account the motion uncertainty between consecutive terrain measurements. Therefore, they are successful only when combined with highly-accurate Inertial Navigation Systems (INS) or dense-range sonar sensors. In contrast, Bayesian methods estimate the state recursively, fusing measurements directly at observation time via navigation filters. Bayesian estimation is a widely-used framework for vehicle localisation because of its efficiency in fusing multiple sources of information, including both statistics and models of uncertainty.

Although TAN methods in the terrestrial domain date back to the early 1980s (Golden, 1980), the application of these methods to the underwater domain is relatively immature with outstanding challenges to be addressed. First, the underwater environment is dynamic and unstructured. Second, vehicle motion dynamics and sensors used to perceive underwater physical features increase the TAN complexity since, e.g., sensors are limited in range, accuracy and precision. Previous research (Nygren and Jansson, 2004) developed a method for terrain-aided navigation using a multi-beam echo-sounder for obtaining bathymetric information. An essential step in this method is to establish that the measurement likelihood function is

approximately Gaussian, given dense sonar measurements, and hence a KF can be used to fuse measurements from a multi-beam sonar and to provide an optimal solution in terms of Minimum Mean Square Error (MMSE), for linear systems with additive Gaussian, zero-mean, uncorrelated white noise. However, by definition TAN is a highly non-linear problem, where the non-linearity primarily arises from the unstructured underwater environment and the non-linear seafloor morphology. Therefore, the terrain measurements are characterised by non-linear functions of the vehicle position. If conditions for successful use of a KF do not hold (even approximately), as for the application considered in this paper (instead of an expensive multi-beam echo-sounder, a low information sonar is considered) and in (Donovan, 2011; Teixeira et al., 2016), non-parametric and non-linear Bayesian estimators need to be used. This has led to the application of particle filters (Arulampalam et al., 2002).

Particle filters are feasible in this area because of their ability to efficiently track multi-modal and non-Gaussian probability densities. This characteristic of PFs is of high importance for map-based localisation methods, where no unique mapping exists between depth measurements and AUV position. In (Anonsen and Hallingstad, 2006) a comparison between a Point Mass Filter (PMF) and a Bootstrap Particle Filter (BPF) for underwater TAN is given. Both are approximate realisations of Bayes filtering, with the difference that the PMF is a grid-based method, whereas the BPF is a sample-based sequential Monte Carlo method. Results showed that PMF slightly outperforms the BPF, though at the expense of higher on-board power consumption.

The developers of long-endurance vehicles typically seek to minimise hotel load by utilizing low-power processing boards with limited computational resources (Phillips et al., 2017; Phillips et al., 2012). This restricts the complexity of the localisation solution, a major limitation to the application of particle filters. In contrast to simple KF algorithms, PFs can be computationally very demanding when the state vector dimension increases. To reduce the computational requirements of the PF when a large number of parameters are to be estimated, the Rao-Blackwellized Particle Filter (RBPF) has been considered as navigation filter (Teixeira, 2007; Schon et al., 2006). An RBPF, also termed Marginalized Particle Filter (MPF), decomposes the state vector into linear and non-linear state sub-vectors. Then, the linear part of the state is estimated using a simple KF, while the non-linear parameters are estimated with a PF. In (Karlsson and Gustafsson, 2003) the performance of the MPF filter was compared to the widely used Sampling Importance Resampling (SIR) filter and the Cramer-Rao Lower Bound (CRLB), where the latter sets the lower limit for the error variance of a MMSE estimator. Both navigation solutions (MPF and SIR) performed very close to the CRLB. In addition, results showed that the state decomposition works well and that the MPF is both slightly more robust and more accurate in terms of Root Mean Square Error (RMSE), when compared to the SIR solution, for estimating a state vector of five elements.

To further increase the estimation accuracy and robustness of the MPF, (Tan et al., 2016) complemented bathymetry-based localisation with range measurements between multiple AUVs. TAN methods have also been reported for underwater gliders (Claus and Bachmayer, 2015). This work examined the potential of TAN for the Slocum glider using the SIR PF. The method was evaluated both off-line and in on-line field trials, and the results clearly demonstrated the potential of TAN for reducing the DR estimation error growth. In contrast to the vast majority of TAN research in the literature, where the reference map exists before the AUV deployment, (Morice et al., 2009) report successful demonstration of TAN on a real system, where the map used for terrain referencing was created using a multi-beam sonar during the mission. Once the map was constructed, a particle filter was used to fuse bathymetric data and estimate the error of a combined INS-Doppler Velocity Logger (DVL) navigation system. Also seeking to address underwater localisation and map building, (Fairfield and Wettergreen, 2008) report challenges related to on-line underwater mapping and propose a method to create a three-dimensional (3D) occupancy grid map using an octree-based data structure. From the localisation perspective, this research uses active localisation to enable higher filter accuracy and convergence after a simulated descent to 1000 m and 6000 m. These simulated dives use conservative error estimates, whilst data was collected from a physical platform during 1500 m dive field experiments. The active localisation method evaluates in simulation several candidate actions and the action that leads to the largest particle discrimination is selected (instead of estimating the filter entropy).

1.2 Contribution

This research presents a complete, low-complexity TAN algorithm designed for the Autosub Long Range, a deep-rated long-range underwater vehicle. The vehicle operates in proximity to the ocean bottom and in areas of known bathymetry. The algorithm is tested using real data from three deep (up to 3700 m) and long-range operations (up to 195 km in 77 hours) in the Southern Ocean (British Antarctic Survey, 2017). To assess the *a-posteriori* TAN performance, the TAN position estimates are compared against USBL measurements, which are considered as baseline positions, intermittently recorded from a support ship throughout the missions.

The demand for energy conservation to perform large-scale missions (e.g. over 3000 km) limits the available on-board energy and forces the navigation system to only use a small number of low-power sensors. As an alternative to using an energy-expensive dense-ranging multi-beam sonar and high grade INS, this paper reports the use of simple dead-reckoning sensors and altitude measurements from an Acoustic Doppler Current Profiler (ADCP) when in Bottom-Tracking (BT) mode. By using an ADCP to measure the distance to the bottom whilst estimating AUV velocity (Donovan, 2011; Teixeira et al., 2016), significant power can be conserved as there is no need for an additional active sensor dedicated to measure vehicle altitude.

The on-board estimation algorithm has also to conform to energy constraints and implementation simplicity. Therefore, the filter estimates only the two-dimensional (2D) position of the vehicle, avoiding the estimation of various sensor errors explicitly. In this way, the TAN algorithm is kept relatively simple in implementation (e.g. no need for tuning multiple parameters) and reasonably light in computational requirements. While ALR navigates itself using simple dead-reckoning for extended period of time, strong and variable water-currents may cause a significant drift in position estimates (e.g. while descending for more than 3.5 hours to depths greater than 3000 m). However, the 2D terrain-aided navigation is convergent on a localisation solution which agrees with the ship-based USBL to within 200 m in average for all three missions, once the algorithm receives a sequence of range-to-bottom measurements. Although bottom-tracking can occasionally be lost because the terrain is steep and rugged, results show that the algorithm is sufficiently robust to deal with low quality measurements and a sparse bathymetric map.

An extensive literature review produced no publications that report the results of testing TAN in truly deep-water operations (>3500 m) and for such long-range missions (over 190 km) without periodic surfacing to re-initialise the estimation process via GPS. Although TAN accuracy highly depends on the accuracy of motion sensors, the accuracy and density of sonar beams and the bathymetric map quality, the results given in this paper demonstrate that it is possible to navigate an AUV with a low resolution *a-priori* map, basic low-power DR sensors and a sequence of single water-depth measurements. Moreover, techniques for reducing the growth of the dead-reckoning error due to the use of low power/accuracy motion sensors are considered (e.g. on-line compass calibration).

The research reported in this paper extends our previous work aimed at developing TAN for long-range AUV operations, with a motivating interest in under-ice Arctic crossings (Salavasidis et al., 2018). The developed particle-filter based algorithm is tested, in an appropriate simulation environment, for use in missions where an AUV is tasked to cross the Arctic Ocean from Svalbard to the Barrow coast in Alaska, a mission with range over 3000 km. Given sufficient terrain variability, the PF-based TAN algorithm is able to address the positioning error growth, despite a sparse and low-quality bathymetric map. Towards the ultimate goal of an Arctic crossings using Autosub Long Range 1500 (ALR1500), the latest addition to the Autosub family of vehicles (Roper et al., 2017), this work gives significant new results on assessing the potential of TAN for bounding the inevitable positioning error growth that DR experiences during long-range missions.

1.3 Paper Format

The remainder of the paper is structured as follows: Section 2 gives the required background on the ALR vehicle, its localisation sensor suite, the existing DR-based navigation method and the three missions in the

Southern Ocean. Section 3 defines the process and measurement models. Section 4 formulates the terrain-aided navigation as a recursive Bayesian estimation problem and provides an overview of the Sequential Importance Sampling (SIS) algorithm. Section 5 gives the complete TAN algorithm designed for the ALR. In Section 6, the TAN results are compared to DR estimates and the USBL measurements over the three missions in the Southern Ocean. Section 7 summarizes the most important outcomes of this research and gives a short discussion on future work and important aspects to be considered when implementing TAN in real-time and for long-range missions.

2 Deep & Long Range Experiments

Science expeditions as part of the Dynamics of the Orkney Passage Outflow (DynOPO) research program (Naveira Garabato and Forryan, 2017; British Antarctic Survey, 2017) were undertaken in April 2017 in the Orkney Passage (see Figure 2), a submarine valley of the Southern Ocean that connects the Weddell Sea and the Atlantic Ocean. Previous research (Purkey and Johnson, 2010; Purkey and Johnson,

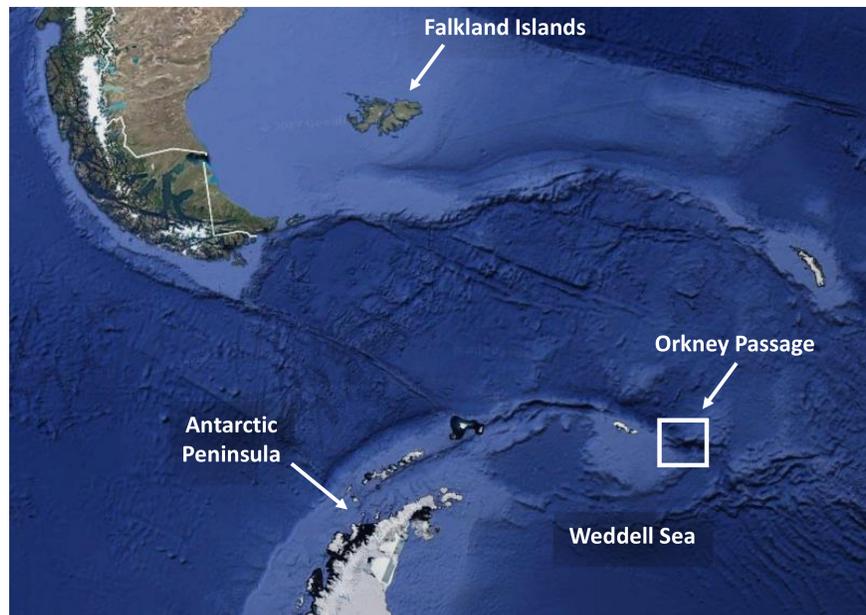


Figure 2: ALR operations at Orkney passage, a narrow gap in the ocean-floor ridge that extends north-east of the Antarctic Peninsula (Google Earth, 2017).

2013) suggests that warming of the deep-ocean (>2000 m) accounts for 30% of all global warming on Earth. Furthermore, deep-ocean warming appears to originate in the Southern Ocean and follow Antarctic Bottom Water (AABW) pathways. DynOPO operations consisted of three deep long-range missions (hereafter referred to as Missions 41, 42 and 44) aiming to map the flow of deep-water masses. The Orkney Passage is an area north-east of Antarctic Peninsula with water-depths of up to 5300 m, steep terrain morphology, and strong and time-varying water-currents (occasionally exceeding 0.6 m/s). Measurements of water temperature, salinity, water-currents, and turbulence close to the ocean bottom in Orkney Passage collected using the ALR vehicle have provided scientists with a new set of data for understanding the role of the bottom water flow in regulating the Earth's climate system. For the deployment and recovery of the ALR, Britain's polar ship, the RRS James Clark Ross (JCR), was used.

The following sub-sections give the relevant background on the ALR vehicle, the existing navigation system and the three DynOPO missions.

2.1 Autosub Long Range

ALR (Furlong et al., 2012) is a deep-rated, long-endurance flight-style AUV (see Figure 1) that is well-suited to the requirements of deep and long-range missions. The ALR is a propeller-driven, 3.6 m long vehicle and has a dry weight of approximately 660 kg. It is designed to have a 6000 m depth-rating and a range of up to several thousand kilometres. Such long-range capability is enabled by ensuring a low hotel load and travelling slowly (0.35–1.0 m/s).

Propulsion power is provided by a custom-designed propulsion unit which drives a two-bladed propeller. The vehicle is manoeuvred using two stern planes and a rudder, while the top fin is fixed and contains antennae for GPS, WiFi and Iridium satellite communication as well as an optional strobe to aid recovery. The moveable control planes are controlled using custom actuators, similar in structure to the propulsion motor, that are installed along the fore-aft axis of the vehicle. The output drive of the actuators goes through a worm and wheel gear to provide the 90° rotation required to drive the control planes. The control planes were sized so that the AUV would be hydrodynamically stable in forward motion, and thus the control planes do not need to be continuously adjusted to maintain heading.

2.1.1 ALR Navigation Sensor Suite

The dry-space and buoyancy for ALR is produced by two 6000 m depth-rated forged aluminium pressure spheres (see Figure 3). The forward sphere contains the batteries for the vehicle, which are lithium primary cells, while the rear sphere contains the control electronics. Given the restrictions in on-board power during long-range missions, the AUV is equipped with: a SiRFstarIV GSD4e GPS module to obtain GPS fixes when the ALR surfaces, a downward looking 300 kHz Teledyne RDI ADCP to estimate the 3D AUV velocity (bottom-relative or water-relative) expressed in the AUV body frame of reference; a 6-axis PNI TCM XB compass/attitude module based on magneto-inductive sensors; and, forward of the front sphere is a large internal flooded space which contains a Sea-Bird SBE 52-MP Conductivity, Temperature and Depth (CTD) probe. The CTD probe is a standard sensor for deriving the operating depth by measuring the hydrostatic pressure.

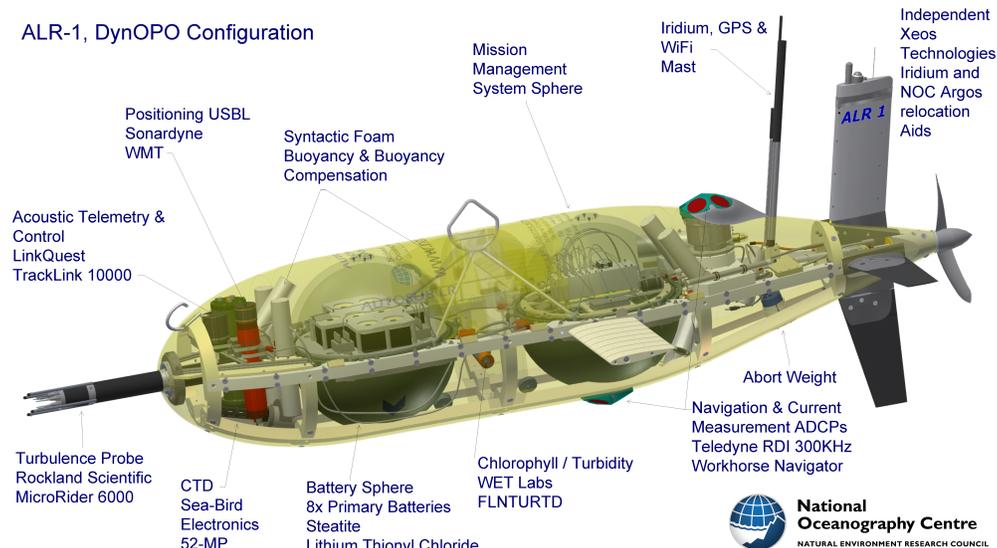


Figure 3: ALR equipped with both localisation sensors and science payload for the DynOPO missions. ALR's dry volume is provided by two 6000 m depth-rated spheres, additional buoyancy is provided by 7000 m rated syntactic foam.

The ADCP is mounted through a hole in the bottom of the rear sphere. It has a maximum reliable range

of 180 m for bottom-tracking and obtaining bottom-relative velocity estimates. For levelling the velocity estimates in a local horizontal reference frame, the ADCP module is integrated with tilt sensors (MEMS accelerometers), instead of using an IMU (resulting in significant power savings during long-range missions). Furthermore, to enable TAN, altitude data can be obtained when the ADCP is locked to the seabed.

2.1.2 ALR On-board Navigation Solution

To completely determine the pose of an AUV in the 3D underwater environment, six independent parameters (6-DoF) are required (three for linear motion and three for rotational). Four of the six parameters are directly measured by the on-board sensors. The AUV depth is measured with a high update rate and accuracy by the CTD probe. The vehicle attitude is acquired using the magnetic compass and the tilt sensors. Assuming that the vehicle is levelled horizontally and stabilized in roll and pitch, the inherently 3D localisation problem can be transformed to a 2D one, where the only unknowns are the two non-directly observable parameters constituting the position in Latitude/Longitude (LL).

An AUV can maintain an estimate of its own position using 2D kinematics and DR, which is a simple mathematical procedure for determining the position of an AUV, given a prior position estimate, the vehicle heading and the horizontally-levelled body velocity. The vehicle body velocity is then rotated to a local North-East (NE) geographical frame, using heading measurements, and integrated with respect to time:

$$\Delta \mathbf{x}_k = \Delta t \cdot R(\psi_k) \cdot \mathbf{u}_k, \quad (1)$$

where: $\Delta \mathbf{x} = [\Delta x^N, \Delta x^E]^T$ represents the horizontal displacement vector expressed in metres in a local NE reference frame; R is the rotation matrix, parametrized only by heading ψ_k , from the AUV body-fixed frame to the NE reference frame; Δt is the sampling time; and $\mathbf{u} = [u^f, u^s]^T$ represents the horizontally-levelled AUV body-frame velocity vector (containing the forward and starboard speed), which is assumed constant during Δt .

The calculated vehicle displacement is then used to propagate prior position estimates. However, the locally-defined NE navigation reference frame assumes the Earth to be a flat surface, which is an acceptable assumption for short-range missions. To avoid this simplification during long-range operations, the Earth is assumed a large spherical object and vehicle positions are expressed in LL. Given an initialisation in LL (typically using GPS when the AUV is on the surface) and the vehicle displacement $\Delta \mathbf{x}$, AUV position estimates are obtained using the following discrete-time kinematics:

$$\mathbf{x}_k = f_{LL}(\mathbf{x}_{k-1}, \Delta \mathbf{x}_k) = \begin{bmatrix} x_{k-1}^{lat} + \Delta x_k^N / R_e \\ x_{k-1}^{lon} + \Delta x_k^E / R_e \cos(x_{k-1}^{lat}) \end{bmatrix}, \quad (2)$$

where R_e represents the Earth radius in metres and $f_{LL}(\cdot)$ is a function that propagates the prior position, $\mathbf{x}_{k-1} = [x_{k-1}^{lat}, x_{k-1}^{lon}]^T$, in time given the vehicle displacement $\Delta \mathbf{x}_k$, whilst maintaining the radian expression. x_{k-1}^{lat} and x_{k-1}^{lon} are the vehicle latitude and longitude, respectively, estimated at time $k-1$ and expressed in radians.

2.1.3 Navigation Errors

Although the DR procedure provides sufficient short-term accuracy, uncompensated sensor errors inevitably induce an unbounded growth in the positioning error. The DR error is a combined effect of multiple error sources in velocity and heading measurements. Also, the geometry of the AUV path may have an effect on the estimation error. Field trials conducted at Portland Harbour, UK in October 2015 revealed that the DR error for the ALR vehicle is of the order of 1% of distance travelled in a straight line.

Doppler effect based velocity sensors, such as the ADCP, calculate the AUV body velocity vector by emitting acoustic pulses towards the seabed whilst the ADCP has bottom-lock (is within acoustic beam range of the seafloor). These transmitted pulses travel through the water column and get reflected by the sea bottom.

The AUV velocity is then calculated by measuring the Doppler shift of the returned echoes (Joyce, 1989). In the absence of bottom-lock (when operating in mid-water column), the ADCP module is able to calculate the AUV velocity relative to the local water masses (Water Tracking - WT mode) due to sound back-scattering by water molecules and particles. In the WT case, DR estimates drift due to the local water-currents and the positioning error growth is in proportion to the magnitude of water-currents. On the other hand, when operations take place near-bottom, the ADCP calculates the AUV velocity with respect to the stationary ocean bottom. The integration of bottom-relative velocity measurements results in a significantly reduced DR error growth rate. However, despite the possibility for bottom-tracking, there are further ADCP errors contributing to the positioning error growth. Bias effects in the ADCP, which are a function of the flying altitude, seabed properties and the inaccurate knowledge of the water sound speed, are examples of not fully understood ADCP error sources.

Besides errors depending upon inaccurate ADCP measurements, other major contributors to the navigation error are the misalignment (in heading) between the ADCP and the magnetic compass and systematic biases of the magnetic compass. In particular, causes for systematic errors in heading are the influences of magnetic materials in the vicinity of the AUV and offsets in the magnetometer and accelerometer measurement vectors. These compass-related biases are probably better understood than the ADCP biases. However, the actual vector sensed by the compass (the horizontal geomagnetic field) varies in an unpredictable manner (e.g. due to geomagnetic storms and currents in the ionosphere and magnetosphere) and might be considered to introduce bias or random errors depending upon the time scales of the observations (seconds to days).

2.1.4 Navigation Error Compensation

To compensate for errors in localisation due to inaccurate sensor measurements, the on-board ALR navigation system performs various on-line corrections and sensor calibrations. To reduce the effect of unknown water-currents, the ALR navigation algorithm estimates the velocity of the water flow local to the vehicle, whenever bottom-relative velocity is available. The AUV forward speed, calculated as a function of propeller rotation¹, is subtracted from the bottom-relative forward speed computed by the ADCP. This together with the bottom-relative starboard speed (provided by the ADCP) constitute, after rotating to a NE reference frame, an estimate of local water-currents. On-board the ALR, this water-current vector is then low-pass filtered and an estimate is maintained when bottom-relative velocity is no longer available. Water-current estimates are important when bottom-track speed is not constantly available (e.g. due to steep downwards slopes). These estimates allow the DR to maintain adequate accuracy, although bottom-relative velocity may be occasionally unavailable.

As far as the systematic biases of the magnetic compass are concerned, the misalignment error produces a navigation error proportional to the radial distance run from the starting position. On the other hand, the compass magnetometer errors produce errors which can be roughly approximated, for a complex mission, as proportional to the total distance run (rather than the radial distance travelled in the case of the alignment error). This is because the compass magnetometer offsets produce errors which are sinusoidally dependent upon the AUV heading. It is important therefore to measure and correct for the compass magnetometer offsets early in a dive. The biases and magnetic influences in the magnetometers can be measured by manoeuvring the AUV through a range of headings. For the ALR, a horizontal calibration is carried out as an automatic part of the dive sequence. Although this horizontal calibration does not measure the full 3D magnetometer offset, it actually measures the combined effect of both magnetometer and accelerometer biases, which affect the horizontal accuracy of the compass. In addition, the ALR has a built-in geomagnetic model correction for magnetic declination variation.

The determination of the misalignment between the compass and the ADCP in deep-water missions is not

¹Previous field operations have shown that the ADCP is likely to give poor quality water-relative speed estimates, particularly at depths greater than 1500 m, due to a sporadic lack of biological matter to provide back-scatter in clear oceanic water. Therefore, to avoid potentially poor water-relative speed estimates, good quality bottom-track data, obtained from deep ocean missions, has been used to calibrate the AUV propeller RPM against bottom-track speed. The calibrated coefficients are then used to estimate the ALR forward speed, instead of using the ADCP water-tracking measurements.

as easy as in shallow-water missions. In shallow water, the AUV obtains GPS fixes from the water surface, dives to a shallow depth (where it reliably obtains bottom-relative speed) and then surfaces to obtain more GPS fixes. The alignment is measured by comparing the GPS-defined track direction, and the dead reckoned track direction. In deep-water this is not possible, and typically an external acoustic aid is required to locate the AUV position when it is deep and within bottom-tracking range.

Finally, the net effect of sensor biases can be further reduced by designing (when possible) closed vehicle paths. Moreover, due to the closed path geometry, bias sensor errors cancel themselves during reciprocal trajectories. All navigational sub-systems and functionalities implemented on-board the ALR are shown schematically in Figure 4.

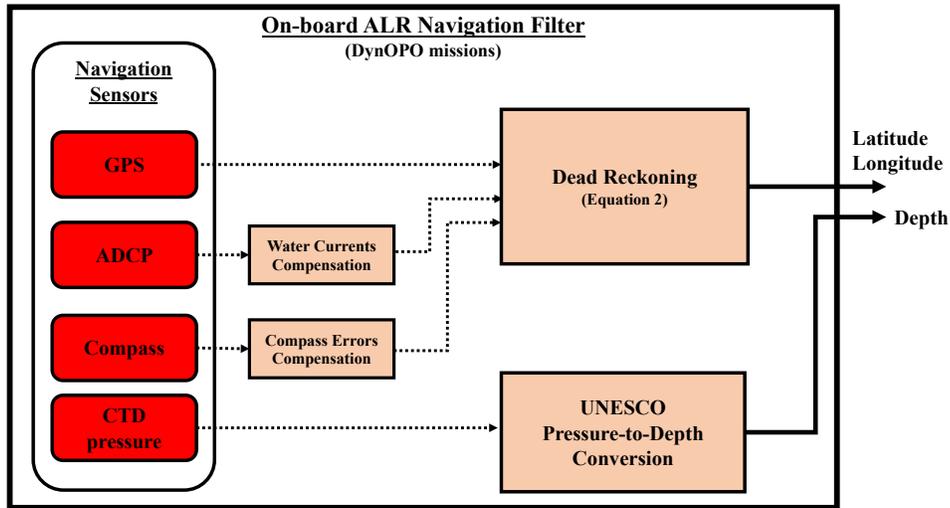


Figure 4: Schematic representation of the ALR’s on-board navigation system based on dead-reckoning.

2.2 Missions Details

Since the aim of the DynOPO project is to investigate the flow of AABW into the Atlantic Ocean through the Orkney Passage, the ALR sensor suite is further enhanced with the MicroRider 6000 instrument package for turbulence microstructure measurement from Rockland Scientific and a second Teledyne RDI 600 kHz ADCP module mounted facing upwards. This additional ADCP, used in combination with the downward-facing ADCP, allows the AUV to collect enhanced quality science data by further refining the estimates of local water-currents (note that data from the additional ADCP is not used for navigation purposes). Water temperature and salinity is provided by the existing CTD probe. Figure 3 shows the complete set of sensors and systems installed on the ALR, for both localisation and science.

To fulfil the requirement for near-bottom measurements, the ALR was demanded to fly at 90 m above the seabed with constant forward speed of 0.7 m/s relative to the sea bottom. In terms of localisation, the ALR was self-navigated, using GPS fixes when on the surface and dead-reckoning (see Section 2.1.2) when submerged. The AUV path was specified by pre-defined waypoints and the vehicle was guided between them using a track following algorithm (McPhail and Pebody, 1998). The maximum achieved depth was approximately 3700 m and the longest mission was approximately 195 km, a range achieved over more than three days. The average descent rate over 3000 m was 0.5 m/s (0.6 m/s near surface to 0.4 m/s at 3000 m); the average ascent rate was 0.62 m/s (0.55 m/s shallow to 0.73 m/s deep). During the descent phase and without bottom-lock, DR estimates are expected to significantly drift due to water-currents.

Thanks to the on-site research ship, the position of the ALR was intermittently recorded during the mission. This was achieved using a USBL head, mounted on the ship hull, and a 6G WMT Sonardyne transponder mounted on the ALR. These USBL measurements were combined with the GPS position of the ship in order to calculate the absolute AUV position. However, USBL outliers are also present, and hence the *Hampel* outlier identifier (Hampel, 1974) was used to identify and remove these outliers. Given the absolute USBL estimates, which are treated as the baseline position of the ALR, and the ALR’s on-board sensor logs, evaluation of the potential of the TAN system for long-range operation becomes feasible².

Figure 5 shows the ALR path for the three missions (M41, M42 and M44) as computed in real-time using the existing ALR navigation algorithm and the ship-based USBL measurements. The mission start is when the vehicle has dived to 2 m depth and the end point is the surfacing location. Furthermore, note that M42 was aborted due to mechanical issues before reaching the final destination. Therefore, the end of M42 is the location at which the vehicle aborted the mission. Table 1 lists details of each mission.

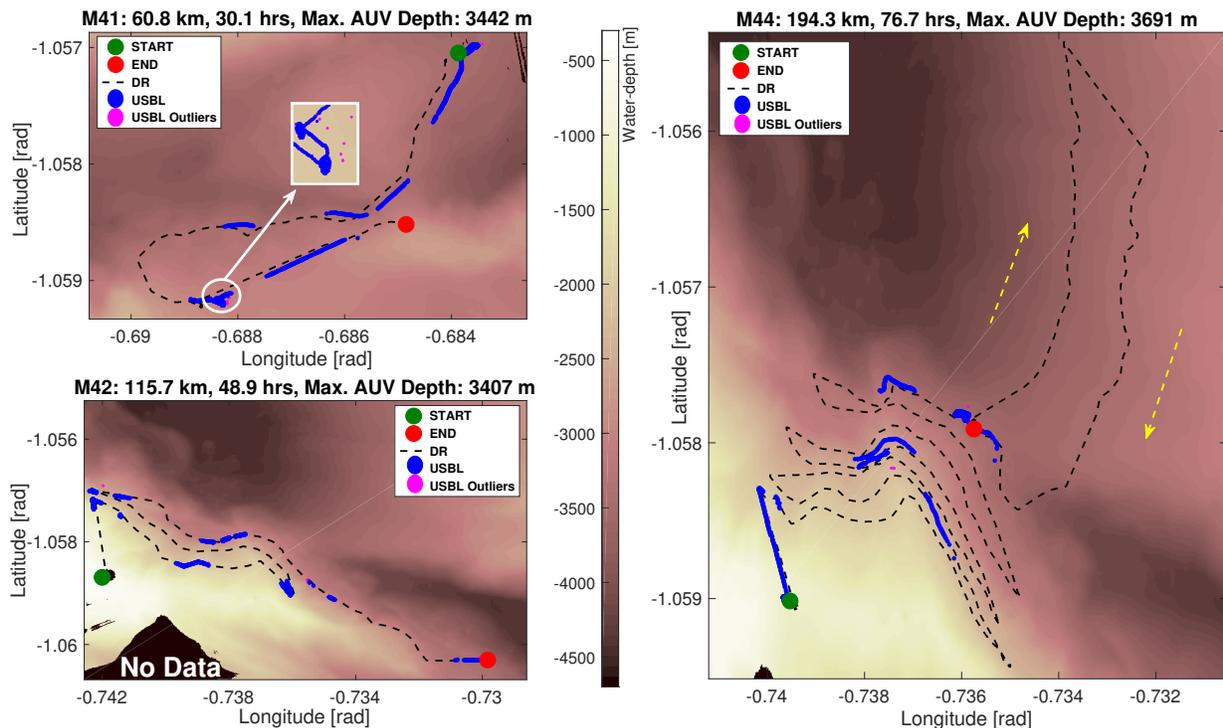


Figure 5: ALR position estimates using dead-reckoning (existing on-board navigation solution) and USBL measurements for the three DynOPO missions. The USBL outlier detection process can be seen in M41. The arrows in M44 show the direction of the vehicle.

For M41 and M44, the RMSE appeared to be significantly higher³ (1088.5 m and 1198.7 m at the last recorded USBL measurement respectively) compared to M42 (309 m) when using DR navigation. This is because M41 and M44 accumulated large positioning error due to strong and time-varying water-currents (occasionally over 0.6 m/s) during both descent and ascent. In contrast, M42 position estimates drifted only during descent, as the mission was aborted at depth. In more detail, although M41 is the shortest mission, the respective DR navigation experienced the largest drift relatively to its range (2.1% of DT) among all missions. There are two reasons for this, firstly the ALR descended to the depth of 3483 m at the beginning of the mission, while at the end it ascended from the depth of 2920 m. In practice, this means that ALR

²Note that the USBL measurements were not used during the missions for updating the real-time position estimates.

³Note that visual comparison between DR and USBL estimates cannot entirely capture the estimation error. Rather than comparing only path deviation, trajectories (including the time factor) have to be considered for proper evaluation.

Table 1: Details of the three DynOPO missions

Parameters	Mission 41	Mission 42	Mission 44	Units
Total duration	30.1	48.9	76.7	hrs
Total duration in bottom-tracking	19.4	48	72.3	hrs
Total Distance Travelled (DT)	60.8	115.7	194.3	km
Distance travelled in bottom-tracking	48	114.1	183.6	km
Unbroken bottom-tracking	YES	NO	NO	-
Maximum AUV depth	3442	3406	3691	m
Start water-depth (descent)	3483	478	1136	m
End water-depth (ascent)	2920	-	3505	m
Number of USBL measurements	4352	2567	4161	-
Number of detected USBL outliers	104	38	84	-
Distance travelled until last USBL	51.6	115.7	191.7	km
DR error (at last USBL)	1088.5	309	1198.7	m
DR error w.r.t DT% (at last USBL)	2.1	0.267	0.625	DT%

operated under the effect of water-currents for approximately 6.4 km, whereas the corresponding range for M44 is approximately 4.6 km and 478 m for M42 (which is only the descent phase as the mission was aborted before the ascent). The second reason for the higher error in M41 is that the misalignment error between the ADCP and compass was calculated based on M41 data. Post-processing data from M41 gave an alignment error of 0.018 radians, which was subtracted from the heading estimate for the subsequent missions (M42 and M44).

3 Problem Formulation

This section formulates the terrain-aided navigation problem. In contrast to most TAN studies, where water-currents and bias terms in vertical dimensions (tides and pressure sensor errors) are included in the estimation process, this implementation considers only two unknown parameters representing the 2D vehicle position. This 2D position is expressed in spherical coordinates since the LL convention avoids the simplification of the Earth being flat, conforms with the USBL measurements and allows a direct interface with the bathymetric map.

Given the ALR flies in proximity to the bottom, an estimation of water-currents is not required. However, it is to be expected that when operating over steep slopes, or in the event of sudden drop-outs, the bottom-track navigation will be lost. To address the effect of water current-induced drift during these periods while keeping the complexity of the navigation filter low, the real-time approach described in Section 2.1.4 for estimating water-currents local to the vehicle is used. Therefore, although the 2D PF does not explicitly estimate water-currents, they are predicted over a short time-window.

In terms of depth offsets, as long-endurance operations took place in open ocean, where typical tidal range is up to 1m (Stammer et al., 2014), and the reference bathymetric map was compensated for tidal variation during its construction, through reference to a given tide level, the variation in vertical dimension is considered negligible.

3.1 Process Model

The AUV motion model is assumed to be a discrete Markov process:

$$\mathbf{x}_k = f_{LL}(\mathbf{x}_{k-1}, \Delta\mathbf{x}_k + \mathbf{m}_k), \quad (3)$$

where: $f_{LL}(\cdot)$ is as defined in Eq. (2); \mathbf{x}_k is the Markovian state vector representing the 2D AUV position $[x_k^{lat}, x_k^{lon}]^T$ expressed in LL at time k ; $\Delta\mathbf{x}_k$ is the displacement caused by the vehicle motion; and \mathbf{m}_k is a process noise following a normal distribution $\mathcal{N}(0, \Sigma_m)$, accounting for sensor errors and unmodelled vehicle dynamics. Two process noise intensities are used: a) high intensity $\mathcal{N}(0, \Sigma_m^H)$, when the AUV is in descent or ascent phase, and b) low intensity $\mathcal{N}(0, \Sigma_m^L)$, when TAN is active and the AUV is in BT range (ignoring instances with sudden loss of bottom-track navigation). This distinction aims to capture the drift due to unknown water-currents during the AUV descent and ascent. Note that it is more correct to assume water-currents as bias errors in a short time-window rather than aiming to capture their effect with normally distributed process noise. However, since the inclusion of water-currents in the state vector is to be avoided, the covariance matrix of the process noise should be selected carefully and large enough not to underestimate the potential long-term effect of water-currents during descent and ascent.

3.2 Bathymetric Map & Water Depth Measurement Model

The effectiveness of the terrain-aided navigation is closely related to the quality of water-depth measurements and of the bathymetric reference map (Meduna, 2011). Since Bayesian filters sequentially perform correlation between on-board measurements and measurement models, this section gives the water-depth model and the associated uncertainty.

Typically, before a conventional AUV deployment, the operating area is surveyed and ship-based multi-beam bathymetry is acquired, processed (e.g. outliers are removed), and gridded to construct a bathymetric reference map. As such, a series of surveys were conducted using the JCR research ship and a ship-based DEM of the Orkney Passage (see Figure 2) was subsequently generated. The JCR is equipped with a Kongsberg E122 multi-beam echo-sounder for bathymetric surveys. Since this is a deep-water area (up to 5300 m) with steep terrain, the creation of a surface-made high resolution and high accuracy reference map is not possible (Jensen et al., 2011). Consequently, the reference map is gridded into squares of 50 m in longitude and latitude ($\approx 0.052^\circ/0.025^\circ$ respectively) and the centre of each grid cell takes the average value of the multi-beam readings that fall within that cell. In areas with a lack of measurements, gaps of up to 10 grid cells wide have been filled using spline interpolation.

Inevitably this bathymetric map is affected by errors. Therefore, a level of confidence is assigned to each grid cell. The uncertainty model used in this research for the reference map due to the errors in the ship-based bathymetry is similar to (Claus and Bachmayer, 2015; Lagadec, 2010; Salavasidis et al., 2018; Marks and Smith, 2008) for characterising errors in bathymetric maps and/or sounding sensors:

$$\sigma_{cell} = 0.5\sqrt{1 + (0.023z_{cell})^2}. \quad (4)$$

This formula is derived from the 95% confidence level of the maximum acceptable vertical uncertainty in sounding sensors (parametrized by the water-depth of the cell, z_{cell}) for deep bathymetric measurements, according to the S44 International Hydrographic Organization (IHO) standards (Order 2). This assigns a normal distribution with space-varying standard deviation σ_{cell} to the reference map, assuming uncorrelated errors in sound travel time (1 m error) and sound velocity (expressed as an error of 2.3% of the true range-to-bottom).

The averaging procedure to calculate the cell's water-depth introduces an additional error in the reference map, hereafter referred to as gridding error. Whenever the terrain morphology is locally smooth, the bilinear interpolation, which is used to calculate the water-depth of an arbitrary location in the map, reduces the gridding error (Meduna, 2011). However, in this research we do not attempt to model the gridding error explicitly (which is a function of terrain characteristics). Instead, it is assumed that the gridding error follows a space-invariant normal distribution $\mathcal{N}(0, \sigma_g^2 = 10^2 \text{ m}^2)$. Therefore, the total map uncertainty is the combination of the uncertainty due to the errors in the ship-based sonar measurements and the errors due to the map gridding:

$$\sigma_{map}^2 = \sigma_{cell}^2 + \sigma_g^2. \quad (5)$$

To calculate the water-depth over the AUV path, measurements from the ADCP and pressure sensor are combined. Depth estimates can be calculated by converting hydrostatic pressure measurements from the CTD probe to the operating depth. This is done by applying the UNESCO pressure-to-depth formula (Fofonoff and Millard Jr, 1983). In the case of the altitude, a common ADCP configuration consists of four beams pointing downwards and displaced by 30° from vertical to calculate 3D velocity. In addition to its use as a motion sensor, the ADCP can be also treated as a four-beam sonar. Each beam provides range-to-bottom readings based on measured round-trip travel time of emitted acoustic signals. Given the ADCP configuration, these range measurements then can be transformed to a single AUV altitude by averaging of as many beams ranges (1 to 4) that are valid (ADCP firmware determines whether ranges are acceptable based on signal-to-noise ratio). This averaging procedure gives a smoother AUV altitude, averaging out fluctuations due to AUV attitude changes.

If the vehicle operating depth d_k is added to the measured altitude r_k , the water-depth z_k at the vehicle's location is computed (see Figure 6). Since sensors are always noisy, a normal distribution $\mathcal{N}(0, \sigma_\omega^2)$ aims to capture the water-depth measurement error. The intensity of the water-depth noise is assumed to vary with altitude and depth (independently of each other) as:

$$\sigma_\omega^2 = (b \cdot r_k)^2 + (c \cdot d_k)^2, \quad (6)$$

where the two terms correspond to noise intensities due to errors in the altitude and depth measurements, and b and c are constants that define, respectively, the proportional relation between the sensor measurements and noise intensities. For the application considered in this research, the two uncertainty terms are defined empirically from many past experiments conducted using these sensors. The maximum expected error in altitude and depth correspond to 1% of the measured altitude and 0.1% of the measured depth. Therefore, b and c are calculated using the 2- σ bounds of the expected errors ($b = 5 \cdot 10^{-3}$ and $c = 5 \cdot 10^{-4}$). Note that small errors could also be introduced if the vertical separation between the sensors (ADCP and CTD) is not considered and compensation introduced.

The terrain-aided navigation filter requires a likelihood function that models the errors in the reference map and water-depth measurements. The following discrete time water-depth model incorporates these errors additively:

$$z_k = h(\mathbf{x}_k) + v_k(\sigma_v), \quad (7)$$

where $h(\cdot)$ is a non-linear function that relates the reference map at position \mathbf{x}_k to the water-depth measurement z_k . The additive term $v_k \sim \mathcal{N}(0, \sigma_v^2)$ is the measurement noise aiming to capture the aforementioned error sources:

$$\sigma_v^2 = \sigma_{map}^2 + \sigma_\omega^2, \quad (8)$$

which is the overall uncertainty model assuming that there are three dominant and independent error sources affecting the measurement model: a) uncertainty due to inaccurate ship-based bathymetric measurements σ_{cell} (see Eq. (4)), b) spatio-temporally invariant uncertainty associated with the map gridding error σ_g (second term in Eq. (5)), and c) uncertainty varying with altitude and depth σ_ω (see Eq. (6)). Given the above measurement model, the filter likelihood function is defined as: $p(z_k|\mathbf{x}_k) = p(z_k - h(\mathbf{x}_k)) = p(v_k)$.

In the general case, if significant tidal variation is expected, an additional uncertainty term that accounts for tides must be included.

4 Bayesian-based Terrain-Aided Navigation using Particle Filters

In this section, the required background on Bayesian estimation and particle filters used to estimate a dynamically evolving system state is given. For a detailed treatment in a general setting see e.g. (Arulampalam et al., 2002).

Terrain-aided navigation is formulated as a non-linear Bayesian estimation problem using a Particle Filter (PF). The goal is to compute recursively the posterior probability density of the AUV state given water-depth

measurements, the likelihood function and the prior density. Particle filters are recursive, non-parametric and non-linear techniques to deal with non-linear motion and/or measurement models without requiring strong assumptions, e.g., linearisation.

The water-depth model given in Eq. (7) is inherently a non-linear function. An intuitive explanation of the seafloor non-linearity is that: as multiple occurrences of similar bathymetry are present in the reference map and therefore no unique mapping exists between water-depth measurements and AUV position (see Figure 6), the seafloor is a non-linear surface. The seafloor needs to be completely flat (an unrealistic scenario) for this assumption not to hold. Therefore, a multi-modal representation of the position probability density is required. PFs are able to track a multi-modal hypothesis and solve the global localisation problem as they can represent arbitrary probability densities.

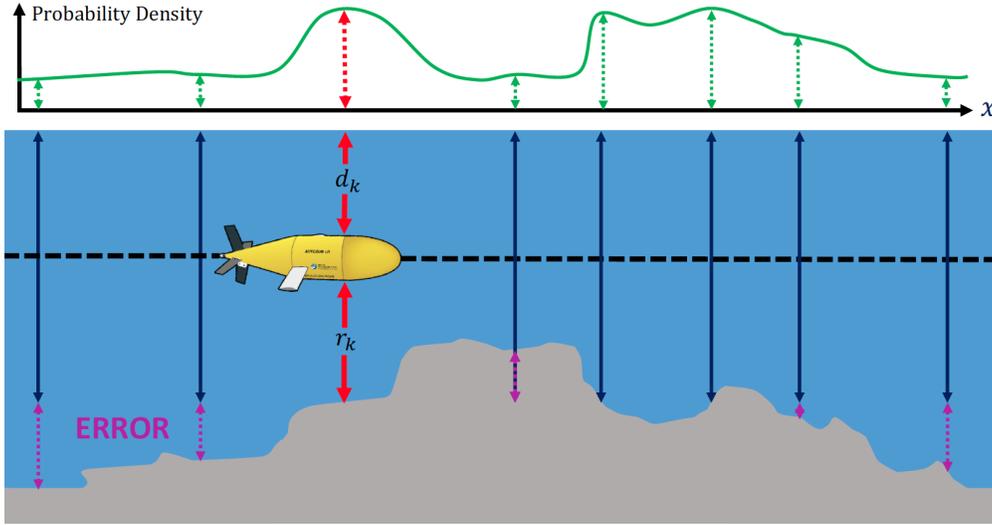


Figure 6: On-board water-depth measurements are computed by summing up the vehicle altitude r_k , provided by the ADCP, and the operating depth d_k , measured by the CTD probe. As a unique mapping between water-depth measurements and AUV position cannot be guaranteed (1D representation), multiple areas of similar water-depths lead to the need for estimation techniques that can handle multi-modal distributions (e.g. particle filters).

Given sufficient terrain variability and frequent bathymetric measurements, it is likely that a multi-modal probability density will eventually converge to a unimodal one. In such an event, the filter is assumed convergent and position estimates can be calculated using MMSE estimates (the mean of the posterior), without introducing significant estimation bias (as would happen in the case of a multi-modal posterior).

4.1 Bayesian Formulation

Using the Markov property of the process model, Bayes filters determine recursively the posterior probability density over the state space conditioned on all the available information, $p(\mathbf{x}_k|z_{1:k})$. Assuming a known initial position density $p(\mathbf{x}_0)$, the two steps of Bayesian estimation (prediction and update) are as follows (Arulampalam et al., 2002):

$$p(\mathbf{x}_k|z_{1:k-1}) = \int_{-\infty}^{\infty} p(\mathbf{x}_k|\mathbf{x}_{k-1})p(\mathbf{x}_{k-1}|z_{1:k-1})d\mathbf{x}_{k-1}, \quad (9)$$

$$p(\mathbf{x}_k|z_{1:k}) = \frac{p(z_k|\mathbf{x}_k)p(\mathbf{x}_k|z_{1:k-1})}{p(z_k|z_{1:k-1})}, \quad (10)$$

where $p(z_k|z_{1:k-1}) = \int_{-\infty}^{\infty} p(z_k|\mathbf{x}_k)p(\mathbf{x}_k|z_{1:k-1})d\mathbf{x}_k$ is a normalizing constant. The prediction step, Eq. (9), calculates the prior density using the probabilistic description of system dynamics $p(\mathbf{x}_k|\mathbf{x}_{k-1})$. Once bathymetric measurements are available, the update step, Eq. (10), calculates the posterior density using Bayes rule and the probabilistic description of the likelihood function $p(z_k|\mathbf{x}_k)$.

4.2 Sequential Importance Sampling

The BF described in Section 4.1 is a conceptual solution which provides a probabilistic technique for data fusion and recursive state estimation. Due to the integral in Eq. (9), a general analytical solution is rarely possible. However, given some assumptions, a BF can be implemented in various ways. BF realisations require specification of the likelihood function $p(z_k|\mathbf{x}_k)$, the motion model $p(\mathbf{x}_k|\mathbf{x}_{k-1})$ and the representation of the probability density over the state vector. As one example, a KF provides the optimal Bayesian solution (in terms of MMSE) if the conditions for the use of a KF hold (i.e. linear process and measurement models with additive Gaussian, zero-mean, uncorrelated white noise). However, these conditions are strong and restrictive since most systems are non-linear and non-Gaussian. Although linearisation can be performed to extend the KF applicability to some terrain-aided navigation scenarios, the Extended Kalman Filter (EKF) is unlikely to perform well in areas with steep terrain morphology (Melo and Matos, 2017).

Particle filters, also termed Sequential Monte-Carlo (SMC) methods, are sample based techniques for approximating the BF solution. In simple terms, a PF algorithm makes multiple guesses (particles) of where the AUV might be and then assigns weights to these guesses based on a likelihood function $p(z_k|\mathbf{x}_k)$, which provides a statistical measure of conformity (or matching) between on-board water-depth measurements and the existing reference map. The more weight a particle has in a region, the higher the probability that the state vector falls within that region. Using this set of weighted samples, it is possible to construct the state posterior probability density, $p(\mathbf{x}_k|z_{1:k})$.

Let N represent the number of particles used by the PF algorithm. Then \mathbf{x}_k^i is the i -th particle and w_k^i is the weight associated with the i -th particle at time-step k . To successfully track the vehicle state, particles should be drawn from the probability density of the current state, also termed a target density, given all the measurements up to that point, $p(\mathbf{x}_k|z_{1:k})$. However, such density is not available to be sampled to generate particles. Instead, the importance sampling principle (Doucet et al., 2000) permits the samples to be drawn from an approximation of the target density, called importance (or proposal) density. If arbitrary samples \mathbf{x}_k^i are drawn from the appropriately chosen importance density $q(\mathbf{x}_k|z_{1:k})$ and the target density $p(\mathbf{x}_k|z_{1:k})$ can be evaluated point-wise (or up to proportionality), then the posterior density can be approximated as:

$$p(\mathbf{x}_k|z_{1:k}) \approx \sum_{i=1}^N w_k^i \delta(\mathbf{x}_k - \mathbf{x}_k^i). \quad (11)$$

In this equation, δ is the impulse function and w_k^i is the normalized importance weight of the i -th particle given by Eq. (12):

$$w_k^i \propto \frac{p(\mathbf{x}_k^i|z_{1:k})}{q(\mathbf{x}_k^i|z_{1:k})} \propto w_{k-1}^i \frac{p(z_k|\mathbf{x}_k^i)p(\mathbf{x}_k^i|\mathbf{x}_{k-1}^i)}{q(\mathbf{x}_k^i|\mathbf{x}_{k-1}^i, z_k)}. \quad (12)$$

Moreover, the approximation in Eq. (11) tends to equality as the number of particles tends to infinity. The recursive algorithm that draws particles from a proposal density and updates the weights based on Eq. (12) is the well known SIS algorithm (Arulampalam et al., 2002; Doucet et al., 2001).

One possible selection for the importance (proposal) density is the prior density $q(\mathbf{x}_k|\mathbf{x}_{k-1}, z_k) = p(\mathbf{x}_k|\mathbf{x}_{k-1})$. In this case particles are weighted recursively using only the likelihood density:

$$w_k^i \propto w_{k-1}^i p(z_k|\mathbf{x}_k^i), \quad (13)$$

where $p(z_k|\mathbf{x}_k^i)$ is the likelihood function of the observation z_k given the particle's location \mathbf{x}_k^i .

The main advantage of this algorithm is its simplicity in implementation. However, a common problem with SIS is the so-called sample impoverishment, or degeneracy (Carpenter et al., 1999), where after a few iterations only very few particles contribute to the final state estimation. The reason for this is that the weight is concentrated over very few particles, whilst the weight of the remaining particles is negligible. In order to avoid this, a metric to detect degeneracy and a method to decrease its effect is needed. Such a metric is the effective sample size (Nordlund, 2002):

$$N_{eff} = \frac{1}{\sum_{i=1}^N (w_k^i)^2}, \quad (14)$$

where w^i is the normalised weight of the i -th particle. According to (Nordlund, 2002), a threshold which may indicate degeneracy is whenever the effective sample size decreases below $N_{th} = \frac{2N}{3}$.

To mitigate the degeneracy effect, a resampling step in the SIS particle filter (SIS+Resampling) has been introduced (Doucet et al., 2000). The concept behind the resampling is to replace particles that have a small weight with others that have a larger weight. Although this technique reduces the degeneracy effect, it can lead to another problem, termed particle collapse, where particles are concentrated in a very small region (Arulampalam et al., 2002). In this case, particles cannot effectively cover the state space in order to capture unmodelled system dynamics. This problem becomes critical and leads to estimation divergence when insufficiently small process noise intensity is assumed. One way for addressing the particle collapse is to apply additional noise (so-called jitter) immediately after executing a resampling step, as reported in (Claus and Bachmayer, 2015; Salavasidis et al., 2016a).

4.3 Sequential Importance Sampling Resampling Algorithm

The two stages of the recursive SISR PF algorithm are:

Prediction: Particles are propagated according to the state transition model (Eq. (3)):

$$\mathbf{x}_k^i = f_{LL}(\mathbf{x}_{k-1}^i, \Delta \mathbf{x}_k + \mathbf{m}_k^i), \quad (15)$$

where i denotes the index of the particle to be propagated.

Update: Once bathymetric measurements are available, the particle weight is calculated according to Eq. (13) and the $\mathcal{N}(z_k - h(\mathbf{x}_k^i) | 0, \sigma_{v,k}^2)$ likelihood function, see Section 3.2 and Eq. (8).

The SISR algorithm specified above is used in the remainder of this work and Algorithm 1 summarises its implementation. Each iteration the algorithm receives as input: the particles with the corresponding weights from the preceding iteration; the AUV displacement $\Delta \mathbf{x}_k$, calculated by the motion sensors; and the observation z_k . The number of particles N and the resampling threshold N_{th} are constants and hence they are not treated as inputs to the SISR algorithm. Outputs of the algorithm are the posterior particles with the associated weights and the filter point estimates $\{\hat{\mathbf{x}}_k, P_k\}$. Particles are propagated using the process model and Eq. (15) in step 4. Then, step 5 assigns weights to newly drawn particles using the likelihood function and Eq. (12). On completion of these steps, the algorithm normalises the particle weights and calculates the point estimates (steps 7-9). Finally, steps 11-14 perform the particle resampling and weight initialisation, given that the resampling condition holds and the number of effective particles is below the resampling threshold N_{eff} . For the resampling procedure, the so-called Systematic Resampling method (Kitagawa, 1996; Carpenter et al., 1999) is used.

5 Terrain-Aided Navigation Implementation

This section presents the entire terrain-aided navigation algorithm developed for the ALR, linking all the components described in the previous sections. The basic navigation idea is to perform TAN when the ALR

Algorithm 1 Sequential Importance Sampling Resampling

Inputs: $\{\mathbf{x}_{k-1}^i\}_{i=1}^N, \{w_{k-1}^i\}_{i=1}^N, \Delta \mathbf{x}_k, z_k$
Outputs: $\{\mathbf{x}_k^i\}_{i=1}^N, \{w_k^i\}_{i=1}^N, \hat{\mathbf{x}}_k, P_k$

```

1: function SISR
2:   for  $i = 1 \rightarrow N$  do
3:      $\mathbf{m}_k^i \sim \mathcal{N}(0, \Sigma_m)$  ▷ Process Noise Samples
4:      $\mathbf{x}_k^i \leftarrow f_{LL}(\mathbf{x}_{k-1}^i, \Delta \mathbf{x}_k + \mathbf{m}_k^i)$  ▷ Particle Propagation
5:      $w_k^i \leftarrow w_{k-1}^i p(z_k | \mathbf{x}_k^i)$  ▷ Particle Weight
6:   end for
7:    $w_k^i \leftarrow \frac{w_k^i}{\sum_{i=1}^N w_k^i}, \text{ for } i = 1 \rightarrow N$  ▷ Weight Normalization
8:    $\hat{\mathbf{x}}_k \leftarrow \sum_{i=1}^N w_k^i \mathbf{x}_k^i$  ▷ Position Point Estimate
9:    $P_k \leftarrow \sum_{i=1}^N w_k^i (\mathbf{x}_k^i - \hat{\mathbf{x}}_k)(\mathbf{x}_k^i - \hat{\mathbf{x}}_k)^T$  ▷ Position Covariance Estimate
10:   $N_{eff} \leftarrow (\sum_{i=1}^N (w_k^i)^2)^{-1}$  ▷  $N_{eff}$  Calculation
11:  if  $N_{eff} < N_{th}$  then
12:     $\{[\{\mathbf{x}_k^i\}_{i=1}^N]\} \leftarrow \mathbf{Resampling} [\{\mathbf{x}_k^i, w_k^i\}_{i=1}^N]$  ▷ Particle Resampling
13:     $w_k^i \leftarrow 1/N, \text{ for } i = 1 \rightarrow N$  ▷ Weight Re-initialisation
14:  end if
15: end function

```

operates within bottom-tracking range and altitude measurements are available. The recursive nature of Bayesian estimation techniques requires the algorithm to be initialised. Therefore, Section 5.1 addresses the particle initialisation. Subsequently, Section 5.2 describes the TAN algorithm in detail, providing also a schematic representation of the entire navigation system (Figure 7).

5.1 Particle Initialisation

Particles can be initialised either at the beginning of the mission, when the vehicle is on the surface, or once the AUV in BT mode and an altitude measurement is acquired for the first time. In the first scenario, the on-board GPS device provides an estimate of the AUV position with an associated uncertainty. Particles are then initialised following a normal distribution centred at the GPS position estimate and a covariance associated with the GPS measurement. After the initialisation, the particles are propagated through time according to the process model, Eq. (15), with the high intensity process noise $\mathcal{N}(0, \Sigma_m^H)$ (since the TAN is inactive and the ALR is in WT mode) until the ocean bottom is approached. However, to conserve computational power, the second initialisation scenario is preferred, where there is no need to propagate particles during the prolonged descent time. Particles are initialised only once the ocean bottom is within the ADCP range. To do so, an estimate of the ALR position and of the corresponding position uncertainty is needed. Therefore, a surface-acquired GPS position and the associated uncertainty are propagated through time, using a process equivalent to executing the prediction phase of a Kalman filter. Therefore, the predicted state is obtained by getting the expected value of the process model (Eq. (3)), which is equivalent to the DR model (Eq. (2)), and the predicted state covariance is computed by propagating the initial GPS uncertainty as:

$$P_k^{KF} = F_{k-1} P_{k-1}^{KF} F_{k-1}^T + L_{k-1} \Sigma_m^H L_{k-1}^T, \quad (16)$$

where P^{KF} is the position error covariance matrix and F, L are Jacobian matrices of f_{LL} with respect to \mathbf{x} and \mathbf{m} respectively, evaluated in each case at $\mathbf{x} = \hat{\mathbf{x}}_{k-1}$. Once the ocean bottom is within range, the particle weight is initialised to $1/N$ and particles are spread using a normal distribution centred at the predicted state estimate with covariance P_k^{KF} .

5.2 TAN Algorithm

A high level overview of the entire navigation system is shown in Figure 7. At each time-step Δt (where Δt

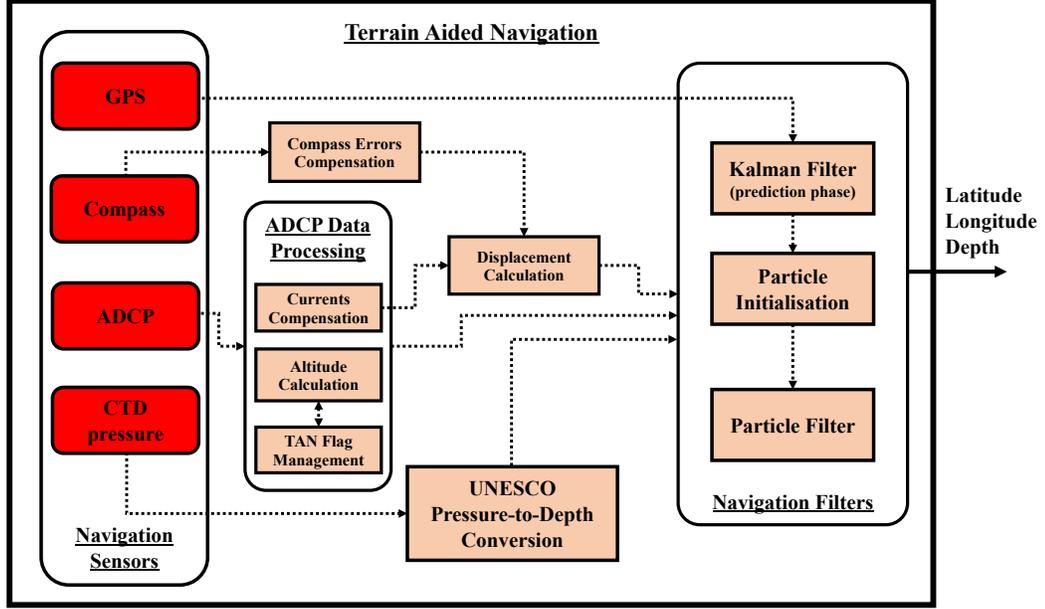


Figure 7: Schematic representation of the terrain-aided navigation system.

is dictated by the lowest frequency among the motion sensors) of the navigation algorithm, the navigation filter receives multiple input/outputs arguments. These arguments are presented in Table 2.

Table 2: Input and output arguments of the terrain-aided navigation algorithm

Inputs	$\{\mathbf{x}_{k-1}^i\}_{i=1}^N$	Prior particles, $\mathbf{x} \in \mathfrak{R}^2$	[rad]
	$\{w_{k-1}^i\}_{i=1}^N$	Prior Particle Weights	[-]
	$\hat{\mathbf{x}}_{k-1}$	Prior Position Estimate	[rad]
	$\Delta \mathbf{x}_k$	Horizontal Displacement	[m]
	r_k	Altitude	[m]
	d_k	Operating Depth	[m]
	N	Number of particles	[-]
	Σ_m^H	High Intensity Process Noise Covariance	[m ²]
	Σ_m^L	Low Intensity Process Noise Covariance	[m ²]
	P_k^{KF}	Kalman Filter Covariance	[rad ²]
N_{th}	Resampling Threshold	[-]	
Outputs	$\{\mathbf{x}_k^i\}_{i=1}^N$	Posterior Particles	[rad]
	$\{w_k^i\}_{i=1}^N$	Posterior Particle Weights	[-]
	$\hat{\mathbf{x}}_k$	Position Estimate	[rad]
	P_k	Particle Covariance	[rad ²]
	P_k^{KF}	Kalman Filter Covariance	[rad ²]

Algorithm 2 shows the required steps for implementing an iteration of the TAN on-board the ALR⁴. To properly execute the terrain-aided navigation filter, three flags are introduced (but, for simplicity, excluded

⁴Note that Algorithm 2 is only a simplified pseudocode used for demonstration.

from being input-output arguments of the navigation pseudo-code) and managed by an associated data processing block called *TAN Flag Management*. Two flags, denoted *InitFlag* and *TANactive*, are required to avoid the initialisation of the particle filter multiple times, which has to be performed only once and after the TAN activation (first obtained altitude measurement). Initially, both flags are set to zero and the initially estimated AUV position, based on GPS readings, is propagated forward in time (steps 2–5). Once bathymetric data is obtained, the *TAN Flag Management* block sets both *InitFlag* and *TANactive* flags to 1. Now, instead of propagating the position information forward in time, the particle filter initialisation is triggered. Given the previously predicted position, steps 6–10 initialize the particle filter. Immediately after the particle initialisation, the *InitFlag* flag is un-set and remains in this condition forever. The *TAN Flag Management* block also determines whether the calculated altitude is within the acceptable/reliable range. If so, the TAN update phase is allowed by setting an associated flag, denoted as *BathFlag*. However, if suddenly an unreliable measurements arrives or the AUV loses the bottom-lock (e.g. due to steep downwards slopes), the *TAN Flag Management* block un-sets the *BathFlat* and the algorithm does not execute the PF update phase.

6 Terrain-Aided Navigation Experiments

The TAN approach developed in Section 5 is implemented in MATLAB using the field-data collected during ALR operations at Orkney Passage (see Section 2 for mission specifications). The post-processing algorithm incorporates logged sensor readings and the ship-constructed bathymetric map (see Section 3.2) of the area. In order to numerically evaluate the TAN performance, terrain-aided navigation results, which are obtained after executing a set of Monte Carlo (MC) runs, are compared to the on-board real-time navigation solution (see Section 2.1) and the USBL measurements collected by the JCR.

6.1 Evaluation Metrics

For convenience and easier interpretation of results, filter estimates and the associated covariances are expressed in metres and square metres, respectively. Therefore, positions in LL (e.g. filter estimates $\hat{\mathbf{x}}$ or particle location \mathbf{x}^i) are transformed to a time-varying local NE reference frame, by solving Eq. (2) inversely:

$$\Delta \mathbf{x} = f_{NE}(\mathbf{x}_i, \mathbf{x}_j) = \begin{bmatrix} (x_i^{lat} - x_j^{lat}) R_e \\ (x_i^{lat} - x_j^{lat}) R_e \cos(x_j^{lat}) \end{bmatrix}, \quad (17)$$

where f_{NE} transforms two positions (\mathbf{x}_i and \mathbf{x}_j) expressed in LL into the horizontal separation $\Delta \mathbf{x}$ of these positions, referenced in a local NE frame originated at \mathbf{x}_j . For example, to calculate the estimation error in metres at time k , given the existence of USBL measurement \mathbf{x}_k^{USBL} , the filter estimate $\hat{\mathbf{x}}_k$ is transformed to a local NE reference frame originated at the USBL position. By using Eq. (17), the horizontal separation/error vector between the USBL measurement and the TAN position estimate can be computed:

$$\mathbf{e}_k = [e_k^N, e_k^E]^T = f_{NE}(\hat{\mathbf{x}}_k, \mathbf{x}_k^{USBL}), \quad (18)$$

where e_k^N and e_k^E are the estimation errors at time-step k in north and east direction, respectively. Note that due to the noise in USBL measurements, \mathbf{e}_k is in essence only an estimate of the navigation error. Similarly, the particle distribution can be represented in a NE reference frame (e.g. originated at the filter estimate $\hat{\mathbf{x}}_k$) in order for the weighted particle covariance to be expressed in square metres at any time-step.

The evaluation metrics used in this work are similar to (Teixeira et al., 2017a):

Time-indexed RMSE of a single MC run

$$RMSE_k^1 = \sqrt{(e_k^E)^2 + (e_k^N)^2}, \quad k = 1, \dots, k_{max} \quad (19)$$

Algorithm 2 SISR-based Terrain-Aided Navigation

Inputs: $\{\mathbf{x}_{k-1}^i\}_{i=1}^N, \{w_{k-1}^i\}_{i=1}^N, \hat{\mathbf{x}}_{k-1}, P_{k-1}^{KF}, \Delta \mathbf{x}_k, d_k, r_k, \Sigma_m^H, \Sigma_m^L, N, N_{th}$ **Outputs:** $\{\mathbf{x}_k^i\}_{i=1}^N, \{w_k^i\}_{i=1}^N, \hat{\mathbf{x}}_k, P_k$

```
1: function TAN
   %%%% Position Propagation Phase Before Particle Initialisation %%%%
2:   if TANactive = FALSE then
3:      $\hat{\mathbf{x}}_k \leftarrow f_{LL}(\hat{\mathbf{x}}_{k-1}, \Delta \mathbf{x}_k)$  ▷ Position Propagation
4:      $P_k^{KF} \leftarrow F_{k-1} P_{k-1}^{KF} F_{k-1}^T + L_{k-1} \Sigma_m^H L_{k-1}^T$  ▷ Covariance Propagation via High Intensity Noise
5:   end if
   %%%% Particle Filter Initialisation %%%%
6:   if InitFlag = TRUE then
7:      $\mathbf{s}^i \sim \mathcal{N}(0, P_{k-1}^{KF})$ , for  $i = 1 \rightarrow N$  ▷ Get Samples
8:      $\mathbf{x}_{k-1}^i \leftarrow \hat{\mathbf{x}}_{k-1} + \mathbf{s}^i$ , for  $i = 1 \rightarrow N$  ▷ Particle Initialisation
9:      $w_{k-1}^i \leftarrow 1/N$ , for  $i = 1 \rightarrow N$  ▷ Weight Initialisation
10:  end if
   %%%% Particle Filter %%%%
11:  if TANactive = TRUE then
12:     $\mathbf{m}_k^i \sim \mathcal{N}(0, \Sigma_m^L)$ , for  $i = 1 \rightarrow N$  ▷ Samples from Low Intensity Process Noise
13:     $\mathbf{x}_k^i \leftarrow f_{LL}(\mathbf{x}_{k-1}^i, \Delta \mathbf{x}_k + \mathbf{m}_k^i)$ , for  $i = 1 \rightarrow N$  ▷ Particle Propagation
14:     $w_k \leftarrow w_{k-1}$  ▷ Particle Weight
15:    if BathFlag = TRUE then
16:       $z_k \leftarrow r_k + d_k$  ▷ Water-Depth Measurement
17:       $\hat{z}_k^i \leftarrow h(\mathbf{x}_k^i)$ , for  $i = 1 \rightarrow N$  ▷ Particle Water-Depth via bi-linear interpolation
18:       $\sigma_{v,k}^i \leftarrow \text{Eq. (8)}$ , for  $i = 1 \rightarrow N$  ▷ Measurement Uncertainty
19:       $w_k^i \leftarrow w_{k-1}^i \cdot \mathcal{N}(z_k - \hat{z}_k^i | 0, (\sigma_{v,k}^i)^2)$ , for  $i = 1 \rightarrow N$  ▷ Particle Weight Update
20:       $w_k^i \leftarrow \frac{w_k^i}{\sum_{j=1}^N w_k^j}$ , for  $i = 1 \rightarrow N$  ▷ Weight Normalisation
21:    end if
22:     $\hat{\mathbf{x}}_k \leftarrow \sum_{i=1}^N w_k^i \mathbf{x}_k^i$  ▷ Position Point Estimate
23:     $P_k \leftarrow \sum_{i=1}^N w_k^i (\mathbf{x}_k^i - \hat{\mathbf{x}}_k)(\mathbf{x}_k^i - \hat{\mathbf{x}}_k)^T$  ▷ Position Covariance Estimate
24:     $N_{eff} \leftarrow \left( \sum_{i=1}^N (w_k^i)^2 \right)^{-1}$  ▷  $N_{eff}$  Calculation
25:    if  $N_{eff} < N_{th}$  then
26:       $[\{\mathbf{x}_k^i\}_{i=1}^N] \leftarrow \text{Systematic Resampling} [\{\mathbf{x}_k^i, w_k^i\}_{i=1}^N]$  ▷ Perform Systematic Resampling
27:       $w_k^i \leftarrow 1/N$ , for  $i = 1 \rightarrow N$  ▷ Weight Re-initialisation
28:    end if
29:  end if
30: end function
```

Time-indexed RMSE of a set of MC runs

$$RMSE_k^R = \sqrt{\frac{1}{R} \sum_{i=1}^R \left[(e_{k,i}^E)^2 + (e_{k,i}^N)^2 \right]}, \quad k = 1, \dots, k_{max} \quad (20)$$

Average RMSE of a single MC runs

$$RMSE^1 = \sqrt{\frac{1}{k_{max}} \sum_{i=1}^{k_{max}} \left[(e_i^E)^2 + (e_i^N)^2 \right]}, \quad (21)$$

Average RMSE of a set of MC runs

$$RMSE^R = \frac{1}{R} \sum_{j=1}^R \sqrt{\frac{1}{k_{max}} \sum_{i=1}^{k_{max}} \left[(e_{i,j}^E)^2 + (e_{i,j}^N)^2 \right]}, \quad (22)$$

where R is the number of independent MC runs and $k = [0, \dots, k_{max}]$ is the discrete time-step of each mission. The time-indexed RMSE represents an estimate of the positioning error over mission time, either for a single run $RMSE_k^1$ (the superscript 1 denotes that the estimated error is calculated over a single MC run), or averaged over the entire set of R MC runs, $RMSE_k^R$. On the other hand, the average RMSE represents the estimated mean positioning error calculated over the entire mission length either for a single run, $RMSE^1$, or for the full set of R MC runs, $RMSE^R$.

USBL measurements, which are treated as baseline locations of the ALR, are not always available and, when available, are not noise-free (especially when there is a high horizontal separation between the ship and the AUV). To account for this, the Absolute Water Depth Difference (AWDD) error metric is also introduced. The AWDD measures the absolute difference between the measured and the predicted water-depth for a given position. The time-indexed and average AWDD are defined as follows:

Time-indexed AWDD of a single MC run

$$AWDD_k^1 = |z_k - \hat{z}_k|, \quad k = 1, \dots, k_{max} \quad (23)$$

Time-indexed AWDD of a set of MC runs

$$AWDD_k^R = \frac{1}{R} \sum_{i=1}^R |z_{k,i} - \hat{z}_{k,i}|, \quad k = 1, \dots, k_{max} \quad (24)$$

Average AWDD score of a single MC runs

$$AWDD^1 = \frac{1}{k_{max}} \sum_{i=1}^{k_{max}} |z_i - \hat{z}_i|, \quad (25)$$

Average AWDD score of a set of MC runs

$$AWDD^R = \frac{1}{R} \frac{1}{k_{max}} \sum_{j=1}^R \sum_{i=1}^{k_{max}} |z_{i,j} - \hat{z}_{i,j}|, \quad (26)$$

where z is the on-board water-depth measurement and \hat{z} is the predicted water-depth for a given position. The meaning of superscripts and subscripts are as in the RMSE metric case.

As the AWDD is a function of the measured and predicted water-depth, it also depends on the inherent noise in z and the accuracy in both localisation and reference map. Considering these dependences, the AWDD cannot be used as a direct metric for assessing the localisation quality. For instance, it is possible that an accurate position estimate produces higher AWDD, when compared to a less accurate one, due to potential inaccuracy in the reference map and/or water-depth measurement. Moreover, AWDD is likely to be an inefficient metric when evaluated at position estimates obtained by the TAN algorithm. In particular, this may be the case when the TAN filter is stuck in a local minima due to terrain symmetries and, although the TAN estimates are poor, it achieves low AWDD values.

Due to the lack of the ground-truth position of the ALR during the entire length of the missions and the potential inaccuracies in the USBL measurements, the AWDD is used to characterise the quality of the available USBL measurements through the mismatch between the measured and predicted water-depth. Since USBL measurements are treated as baseline locations of the vehicle, the corresponding AWDD is to be expected low. In the opposite scenario, either the reference map is locally poorly mapped (and/or the water-depth measurement is inaccurate) or the USBL estimate is of low quality.

6.2 Filter Parameters

To use the TAN solution for the navigation problem, parameters of the particle filter have to be specified, i.e., the number of particles N and the intensity of the process noise (Σ_m^H and Σ_m^L). These parameters are typically defined off-line by appropriately tuning the algorithm based on available data (Claus and Bachmayer, 2015). However, the aim here is to develop an on-line system that does not rely on post-processing existing field-data.

Recalling the PF derivation presented in Section 4.2, it is known that an increased number of particles allows the filter to more accurately represent the posterior density function (Arulampalam et al., 2002). However, this comes at the expense of increased on-board power requirements and hence an extremely large number of particles would be infeasible due to power constraints in long-range applications. The existing ALR on-board computer is able to handle a N greater than 10000 in real-time. However, considering the low dimensional state, a number of particles beyond 1000–5000 would not result in a significant increase in filter accuracy (Claus and Bachmayer, 2015; Salavasidis et al., 2018). To assure filter robustness, $N=5000$ is selected, which is a relatively large number that: a) allows particles to more effectively cover the state space when the filter is relatively uncertain (e.g. during the initialisation of the particles), b) enables the filter to achieve faster convergence to USBL measurements and, c) allows the filter to maintain this convergence throughout the entire mission length. It is important to highlight that filter convergence is assumed when the posterior probability density is unimodal (or nearly unimodal) and the filter estimate uncertainty satisfies a pre-selected uncertainty threshold (Meduna, 2011).

The appropriate selection of the process noise covariance matrix (Σ_m^H and Σ_m^L) is of high importance for the PF algorithm. The process noise intensity in Eq. (3) refers to the scale of the expected and normally distributed mismatch between the process model, and the actual AUV dynamics. One limitation of this 2D modelling is that known sources of systematic errors are not explicitly modelled and hence estimated. An example of violating the assumption in Eq. (3) that only random errors are present is when the AUV operates in water-tracking mode during descent/ascent and systematic errors are introduced by water-currents (assuming that water-currents do not vary significantly in short time intervals).

However, despite the 2D formulation, significant drift due to systematic errors is not to be expected during the nominal near-bottom operation, as the ADCP provides bottom-relative velocity. Moreover, the DR error correction techniques used (see Section 2.1.4) contribute further to reduce the effect of systematic errors (e.g. 2D offset in the magnetic compass or drift caused due to sudden drop-outs in the ADCP). However, it would be wrong to believe that the process model is free of systematic errors in the steady state of bottom-

tracking mode. For example, although the magnetic compass is calibrated on-line, this is a 2D calibration whose performance degrades in the event of significant vehicle pitch and/or roll. If one of these cases hold for extended period of time, a potentially insignificant short-term bias error will produce a non-negligible accumulated drift which the assumed process noise will underestimate. This is likely to become even more severe and lead to filter divergence when the vehicle crosses low informative terrain, poorly mapped regions and obtains noisy water-depth measurements.

From the SIS perspective, the process noise intensity parametrizes the proposal density, which essentially defines the area where the particles are to be sampled and the filter searches for the actual vehicle position. Therefore, the intensity of the process noise must be carefully selected to take into account unmodelled vehicle dynamics (Σ_m^H during descent and Σ_m^L for near-bottom operation) and to provide a sufficient, but not too large, particle spread. Avoiding an unnecessarily large process noise intensity is important for mainly two reasons: a) the number of particles might be insufficient to represent the probability distribution over the state, and b) it is likely that the presence of terrain similarities and of noise in bathymetric measurements would lead the filter to become trapped in local minima⁵.

For near-bottom operation, the process noise intensity is defined as:

$$\Sigma_m^L = \Delta t^2 \begin{bmatrix} \sigma_N^2 & 0 \\ 0 & \sigma_E^2 \end{bmatrix} = \Delta t^2 \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad (27)$$

and expressed in square metres. This noise intensity is considered sufficient to capture the unmodelled vehicle dynamics and is low enough to avoid divergence due to the terrain similarities. Furthermore, it is important to use a slightly higher noise intensity than expected to re-gain the particle diversity after the particle resampling step, if performed.

In terms of the process noise during descent, Σ_m^H depends on the time that the AUV remains under the effect of water-currents. M41 began in an area with approximately 3400 m water-depth. Assuming an average drift caused by water-currents of the order of 0.15–0.2 m/s (which can be obtained from spatio-temporal historical data) and descent duration of 3.5 hours (including the procedure for the compass calibration during descent, see Section 2.1.4), a Σ_m^H of the following intensity is assumed sufficient to capture the drift during descent in order to effectively initialise the particle filter:

$${}^{M41}\Sigma_m^H = \Delta t^2 \begin{bmatrix} 36 & 0 \\ 0 & 36 \end{bmatrix}, \quad (28)$$

which is expressed in square metres and used for propagating the position uncertainty, before the AUV approaches the ocean bottom and the *TAN Flag Management* block triggers the PF initialisation process.

Since the water-depth at the beginning of M42 and M44 is shallower (approximately 500 m and 1150 m respectively) compared to M41, the required time to reach the ocean bottom is shorter (approximately 1 h and 1.5 h respectively, including the procedure for the compass calibration). Hence, the bias effect of local water-currents in M42 and M44 can be captured with a lower intensity process noise, compared to the intensity of the process noise in M41. Given this, ${}^{M42}\Sigma_m^H$ and ${}^{M44}\Sigma_m^H$ are as follows (in square metres):

$${}^{M42}\Sigma_m^H = {}^{M44}\Sigma_m^H = \Delta t^2 \begin{bmatrix} 16 & 0 \\ 0 & 16 \end{bmatrix}, \quad (29)$$

which are used for propagating the position uncertainty during descent in the corresponding missions.

In terms of the initial ALR position uncertainty, it is assumed that the GPS error is normally distributed around the estimated GPS position with covariance (expressed in square metres):

$$\Sigma_{INIT} = \begin{bmatrix} 25 & 0 \\ 0 & 25 \end{bmatrix}. \quad (30)$$

⁵The terrain similarity problem can be addressed to some extent (Teixeira et al., 2017a), but is out of scope of this research.

In addition to the covariance matrices, Table 3 summarises the ALR motion demands (see Section 2.2) and the required parameters for setting up the TAN algorithm.

6.3 Terrain-Aided Navigation Results

This section evaluates the performance of the TAN algorithm using data from M41, M42 and M44. The primary aim is to test the performance of the filter in terms of estimation accuracy, uncertainty and repeatability⁶. Therefore, $R = 20$ MC runs are executed for each mission and the results are summarised in Table 4. It is important to highlight that mission waypoints were defined along bathymetric contour lines for safety reasons. This fact makes the TAN problem extremely challenging since the AUV flew in parallel to the contours, and thus over low-informative along-track terrain morphology for prolonged time intervals (see specifically M42 and M44).

Mission 41

Figure 8 shows the AUV path using TAN, as calculated averaging 20 MC runs, compared to DR position estimates and USBL measurements during M41⁷. Although M41 is the shortest mission (see Table 1), DR position estimates experience the largest error due to the long descent to a depth of 3400 m, during which the process model is updated using only water-relative speed estimates. In addition, DR-based positioning further drifts due to the misalignment between the compass and ADCP throughout the entire mission (see Section 2.1.3). Nevertheless, not only is the TAN convergence maintained in all 20 runs, but the TAN estimation error is significantly reduced compared to the DR error. The average RMSE for all MC runs⁸, $RMSE^{20}$, is 158.9 m when the ALR operates near-bottom, whilst the corresponding DR error is 1168.5 m.

Figure 9.(a) shows the time-indexed RMSE ($RMSE_k^{20}$) evolution over the Distance Travelled (DT) for TAN and DR, whenever a USBL measurement is available. Although DR estimates experience a growth in the estimated error as time progresses, TAN error estimates are bounded given that the vehicle operates in bottom-tracking mode and water-depth measurements are available. In addition, to demonstrate the ability of the filter to re-produce accurate position estimates, the associated upper and lower time-indexed error bounds are plotted over the DT. These error bounds correspond to the highest (upper error bound) and lowest (lower error bound) observed $RMSE^1$ among all MC runs. Figure 9.(b) shows when the ALR is in BT mode and when USBL measurements are available.

During the initialisation of the filter, the large accumulated DR error (see Figure 9.(a)) is inevitably inherited

Table 3: ALR motion details and TAN filter set-up

Motion & PF parameters	Value	Unit
Demanded forward speed	0.7	m/s
Demanded altitude	90	m
Acceptable beam range	[0,180]	m
DR propagation frequency, (Δt)	1	sec
Reference map resolution	50	m
Measurement model parameters, (σ_g, b, c)	$10, 5 \cdot 10^{-4}, 5 \cdot 10^{-3}$	m,-,-
Bottom-tracking process intensity, (Σ_m^L)	$diag[\Delta t^2, \Delta t^2]$	m ²
Descent process noise intensity - M41, (Σ_m^H)	$diag[36 \cdot \Delta t^2, 36 \cdot \Delta t^2]$	m ²
Descent process noise intensity - M42 & M44, (Σ_m^H)	$diag[16 \cdot \Delta t^2, 16 \cdot \Delta t^2]$	m ²

⁶The filter results are considered repeatable when all MC runs produce similar $RMSE_k^1$ for each k time-step.

⁷Note that Figure 8 does not contain time information. In order to have a complete picture of the estimation error, time factor has to be considered.

⁸Note that the average RMSE does not entirely reflect the TAN performance since USBL measurements are not uniformly distributed in time. A significant number of USBL measurements are concentrated in the PF initialisation areas where TAN typically experiences its largest estimation error.

Table 4: Estimation error results during bottom-tracking for M41, M42 and M44

Metrics	Mission 41			Mission 42			Mission 44			Units
	TAN	DR	USBL	TAN	DR	USBL	TAN	DR	USBL	
$RMSE^{20}$	158.9	1168.5	-	105.8	355.3	-	152.4	400.1	-	m
$AWDD^{20}$	3.2	65.1	8.9	2.8	51.6	9.4	3.1	68.9	8.4	m
$RMSE_k$, at $k = k_{BT}^{USBL}$	51	1088.4	-	59.5	309	-	126.1	398.1	-	m
$RMSE_k$, at $k = k_{BT}^{USBL}$	0.09	2.1	-	0.05	0.267	-	0.06	0.21	-	DT%

$\dagger k = k_{BT}^{USBL}$ corresponds to the time of the last USBL measurement during bottom-tracking.

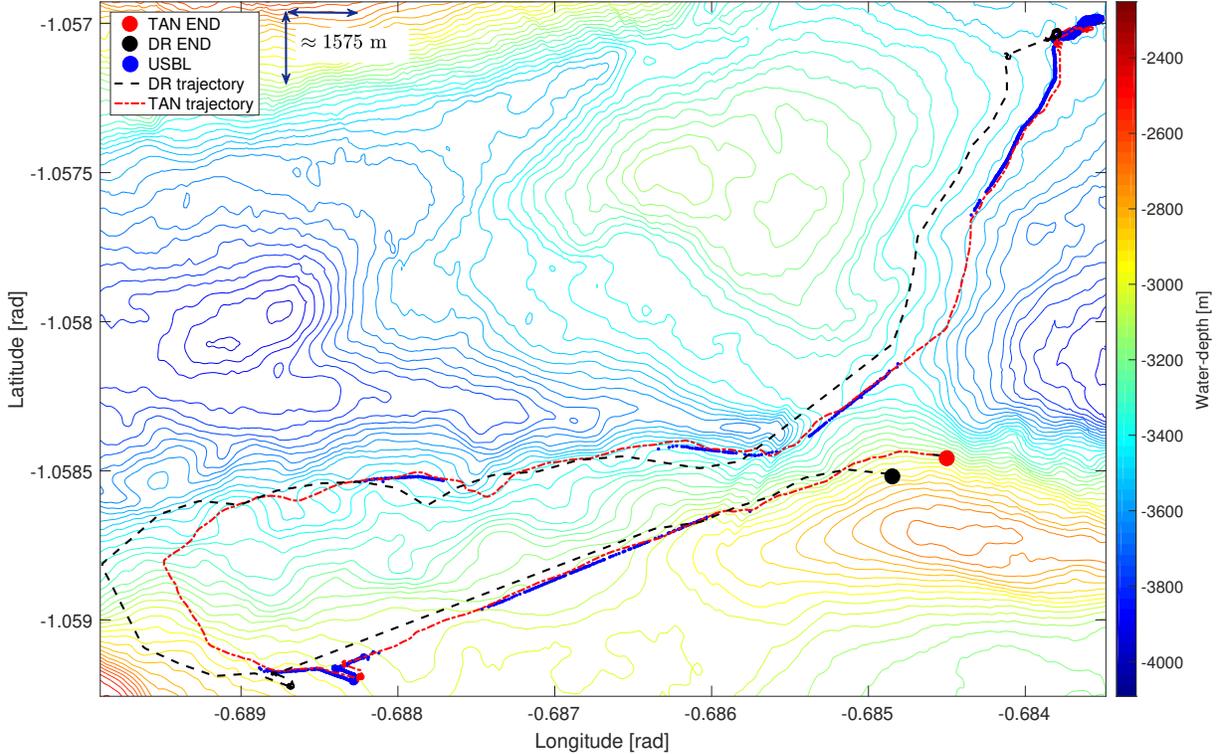


Figure 8: M41 - Estimated AUV path using TAN (red) and DR (black). There are 4248 available USBL measurements (blue), covering approximately 3.92% of the mission.

by the TAN-based estimates, which results in a high initial TAN estimation error. However, TAN is able to rapidly reduce the initial error to an error lower than 100 m. This demonstrates that ${}^{41}\Sigma_m^H$ is high enough in order for the predicted position uncertainty P_k^{KF} to capture the drift caused by water-currents during descent. Furthermore, it can be seen that the inherent stochastic nature of the particle filter initialisation does not result in significant variation between MC runs. The filter produces repeatable results as the error bounds are convergent during initialisation (7–9 km of the DT). This indicates that the number of particles is sufficiently large to effectively cover the initialisation area (which typically corresponds to the largest area to be covered by particles as the filter experiences the largest uncertainty during long descents).

The filter estimates are not only repeatable during the initialisation but throughout the entire mission. The overall small error range between the error bounds implies that all MC executions produce similar results, which demonstrates the ability of the filter to produce repeatable estimates. The only part of the mission with non-negligible variability between MC runs is at 23–24 km of DT. At this range, the local mean of the TAN error corresponds to 180 m, whilst the best MC run gives a mean error below 100 m and the worst MC execution experiences an error of approximately 340 m. As the available USBL measurements seem precise

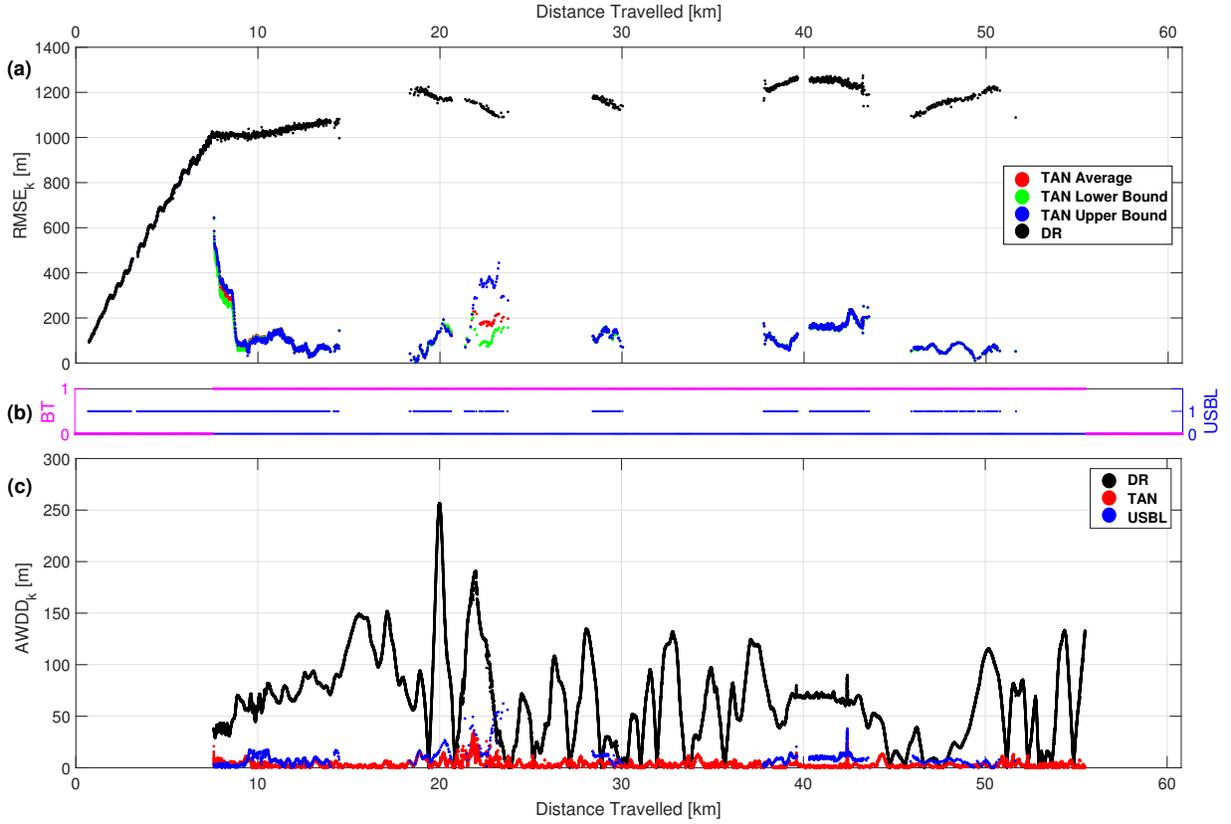


Figure 9: M41 - a) Time-indexed RMSE over distance travelled for TAN (red) and DR (black), along with lower (green) and upper (blue) TAN error bounds of 20 MC runs, b) BT mode and USBL measurement flags (with 1 indicating availability), and c) time-indexed AWDD over distance travelled for TAN (red), DR (black) and USBL (blue).

(low spatial variation, see Figure 8), this slight estimation variability is attributed to potential errors in the bathymetric map or noise in water-depth measurements.

Differences between the measured and predicted water-depth along the vehicle path, as estimated by DR, TAN and USBL measurements, can be seen in Figure 9.(c). The average AWDD ($AWDD^{20}$) for DR position estimates is 65.1 m, whereas for TAN and USBL it is 3.2 m and 8.9 m, respectively. However, at the region where the filter estimates appear less repeatable (23–24 km), none of the positioning methods (DR, TAN and USBL) provide water-depth estimates that accurately match with the observed water-depth. The locally evaluated $AWDD^{20}$ exceeds 30 m, 15 m and 130 m for USBL, TAN and DR, respectively. Given that the USBL measurements are consistent and hence probably reliable, the increased AWDD indicates that either the on-board measurements are noisy or the bathymetric map is locally inaccurate.

Besides the average RMSE ($RMSE^{20}$), it is interesting to examine the time-indexed RMSE at k_{BT}^{USBL} , which is the time of the last recorded USBL measurement whilst the ALR is within bottom-tracking range. The TAN $RMSE_k^{20}$ at $k = k_{BT}^{USBL}$ is 51 m, which corresponds to 0.09% error of the DT, while at the same time the DR error is 1088.4 m, corresponding to 2.1% of the DT.

Figure 10 shows estimates of the TAN error for each direction with the respective two-sided $3\text{-}\sigma$ confidence bounds (or intervals) over the DT. At the beginning of the mission, while the AUV is in WT mode, the filter confidence is calculated by propagating the initial uncertainty of the vehicle position (Eq. (16)) in time,

equivalent to performing the prediction phase of a KF. Once the AUV is in BT mode, the particle spatial distribution is approximated by a bivariate normal distribution uncorrelated in each direction. However, this assumption is occasionally very strong as the particle posterior distribution is rarely an exact uni-modal distribution. Hence, it is likely that the computed $3\text{-}\sigma$ bounds underestimate the actual filter uncertainty. In addition, to also take into account inaccuracies in USBL positioning (e.g. due to the noise in the USBL head and the GPS position of the ship), an uncertainty associated with each USBL measurement is calculated. It is assumed that the noise for both GPS and USBL follows a zero-mean Gaussian distribution. As the GPS noise is independent of the USBL noise, the total error covariance is computed by summing the USBL error covariance and the GPS error covariance. The USBL error is assumed equal to 1% of the measured slant range, as specified in the sensor’s datasheet. The GPS noise intensity is set equal to Σ_{INIT} , similar to the AUV’s GPS uncertainty. Given these assumptions, the $3\text{-}\sigma$ bounds of the estimation error are computed and represented by attaching error bars to each error estimate. Besides the small region at 23–24 km of

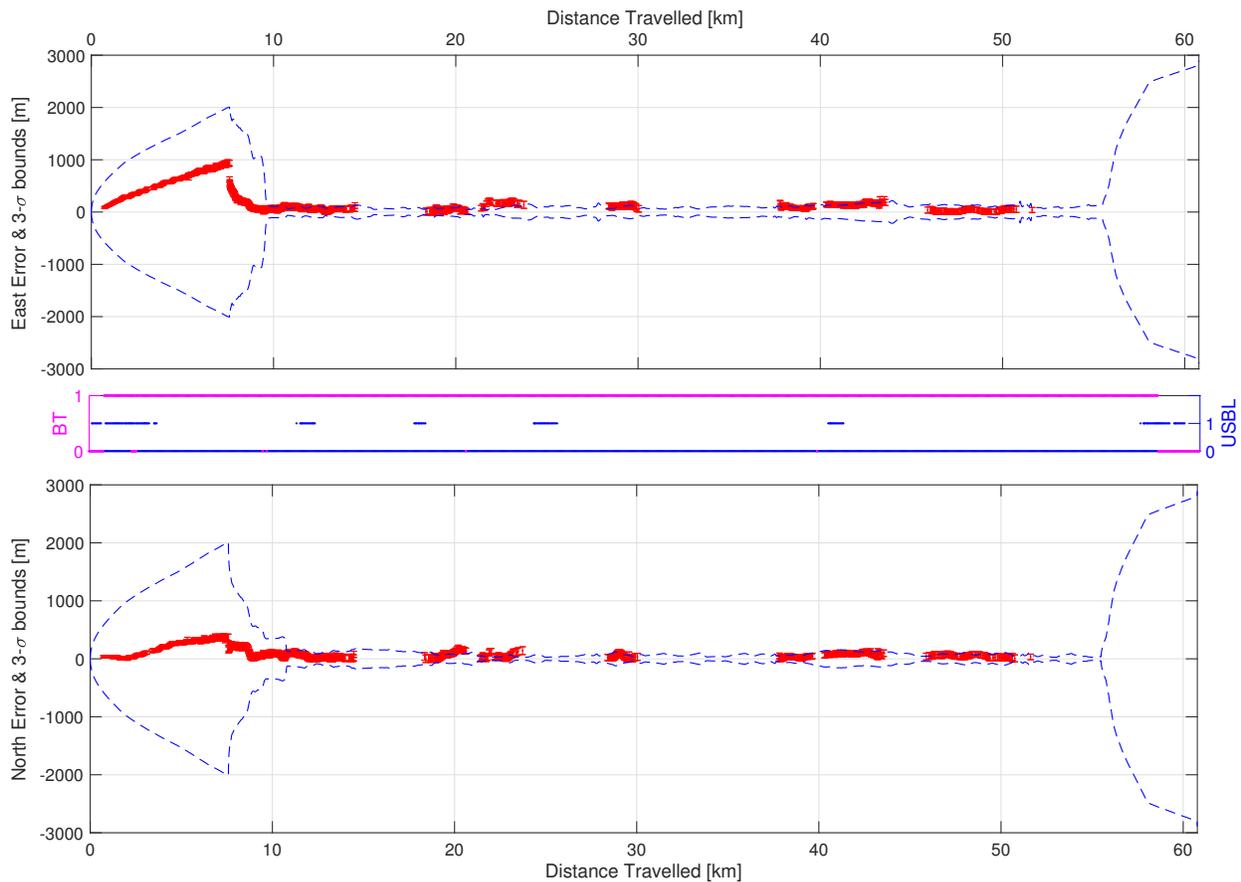


Figure 10: M41 - TAN time-indexed RMSE (red) and the associated $3\text{-}\sigma$ bounds (blue) over distance travelled and north and east direction. The uncertainty in USBL measurements is represented by error bars attached to each TAN error estimate.

DT where the filter appears to underestimate the positioning error (the same region where the filter gives less repeatable estimates, see Figure 9.(a)), error estimates for the rest of the mission almost always fall within the uncertainty bounds, despite the simplifying assumption of an uncorrelated Gaussian particle distribution. Note also that, as it was expected the filter experiences the highest position uncertainty during the initialisation due to the prolonged descent phase.

Mission 42

Figure 11 shows the AUV path as estimated using TAN (averaged over 20 MC runs), DR and USBL measurements for M42. In contrast to M41, where the vehicle is required to dive directly to a depth of

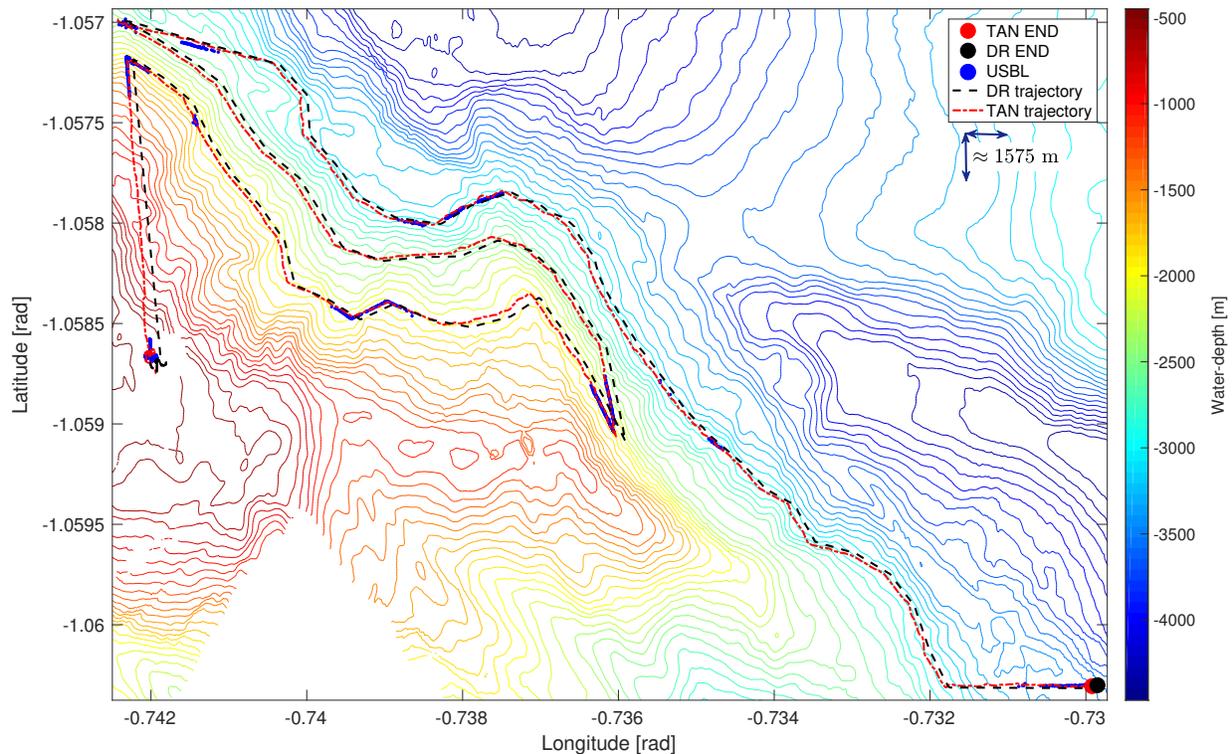


Figure 11: M42 - Estimated AUV path using TAN (red) and DR (black). There are 2529 available USBL measurements (blue), covering approximately 1.44% of the mission.

approximately 3400 m, in M42 the vehicle descends to an initial depth of approximately 500 m. This results in a significantly lower drift caused by water-currents during descent. Figure 12.(a) shows the DR and TAN time-indexed RMSE ($RMSE_k^{20}$) evolution over the DT, including the upper and lower observed TAN error bounds of all MC runs. The vast majority of the DR error is accumulated during the descent time and afterwards remains in the range of 300–400 m. In addition to the short descent time, another reason for low DR error is that DR estimates are compensated for the misalignment error between the ADCP and compass, which was calculated using data from M41.

The TAN algorithm is able to reduce the DR error to lower than 110 m throughout the entire mission. However, a slight increase in error is observed at approximately 15 km of the DT, where bottom-track is lost due to steep terrain morphology. The average TAN RMSE, $RMSE^{20}$, when the ALR operates near-bottom is 105.8 m, whilst the corresponding DR error is 355.3 m. In terms of $RMSE_k^{20}$ at $k = k_{BT}^{USBL}$, the estimated TAN error is 59.5 m, which corresponds to 0.05% error of the DT, whereas the DR error is 309 m, corresponding to 0.23% of the DT. It is also important to notice that the time-indexed RMSE bounds are almost always convergent and hence the filter gives almost identical results for each MC run. The ability of the filter to re-produce position estimates indicates that the number of particles is large enough to effectively cover the state space, particularly during the filter initialisation where the inherent stochastic nature of the filter does not significantly vary the filter results.

Figure 12.(c) shows the time-indexed AWDD over the DT, as calculated using TAN, DR and USBL position estimates. The average AWDD ($AWDD^{20}$) for DR is 51.6 m, whereas TAN estimates and USBL measurements achieve 2.8 m and 9.4 m, respectively. Given the TAN's low time-indexed AWDD and RMSE, where

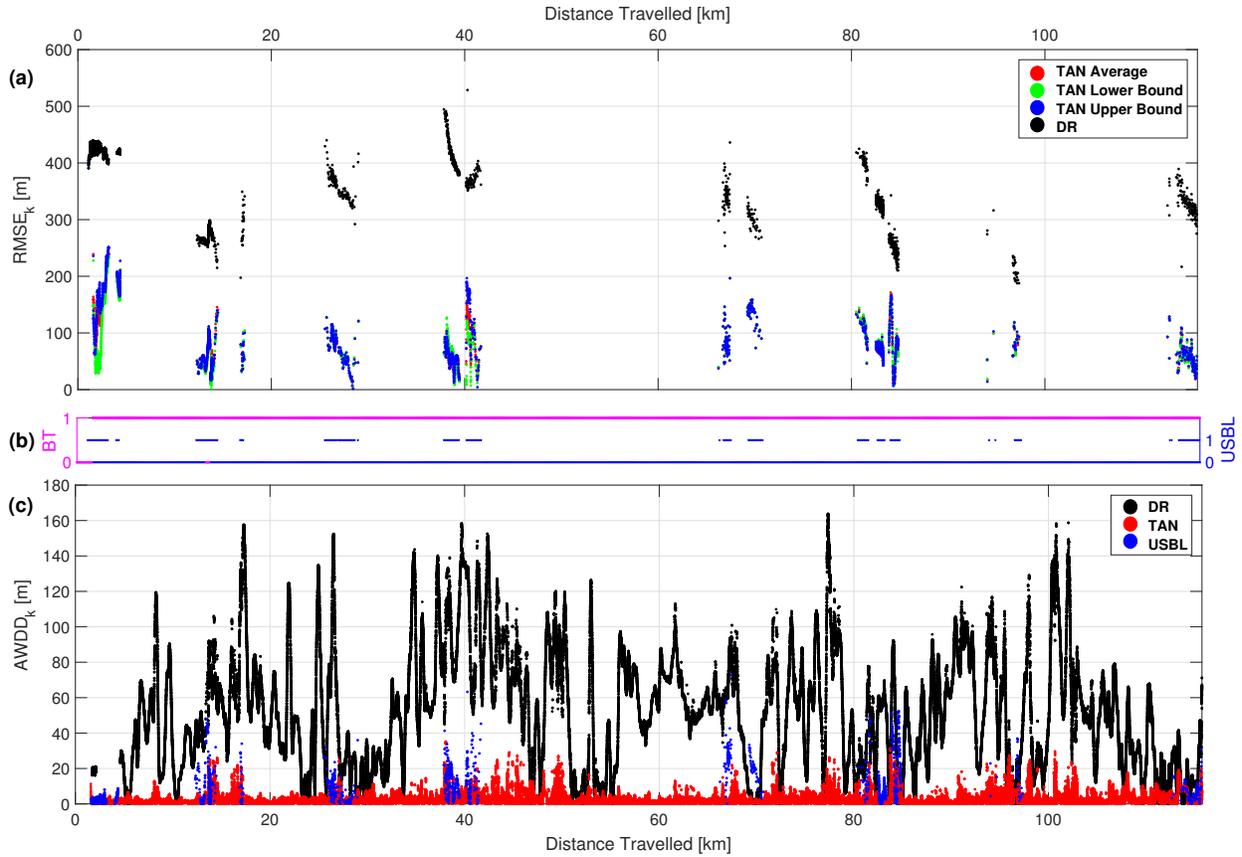


Figure 12: M42 - a) Time-indexed RMSE over distance travelled for TAN (red) and DR (black), along with lower (green) and upper (blue) TAN error bounds of 20 MC runs, b) BT mode and USBL measurement flags (with 1 indicating availability), and c) time-indexed AWDD over distance travelled for TAN (red), DR (black) and USBL (blue).

the latter is sometimes lower than 110 m, it is likely that the largest part of the estimated TAN error is due to the noise in USBL positioning. The increase in the time-indexed AWDD (see Figure 12.(c)) for the USBL measurements when high TAN time-indexed RMSE is observed (e.g. at 41 km, 72 km and 84 km of the DT), further adds to the conjecture that the estimates of the TAN error are greatly impacted by USBL errors.

Figure 13 shows the estimates of the TAN error for each direction with the respective two-sided confidence bounds over the DT. In the vast majority, the estimates of the TAN error fall within the $3\text{-}\sigma$ confidence bounds. Note also that the selected process noise intensity for the vehicle descent phase is sufficient to capture the error induced by the water-currents during descent.

Mission 44

The performance of the TAN algorithm is further evaluated using the longest mission (see Table 1). Due to the mission range (≈ 200 km) and the AUV path along bathymetric contour-lines (low-informative terrain morphology along-track) for prolonged intervals, this mission is considered as the most challenging for TAN. Figure 14 shows the AUV path as estimated using TAN (averaged over 20 MC runs), DR and USBL measurements.

To assess the performance of the TAN algorithm in relation to the mission time (which is not captured

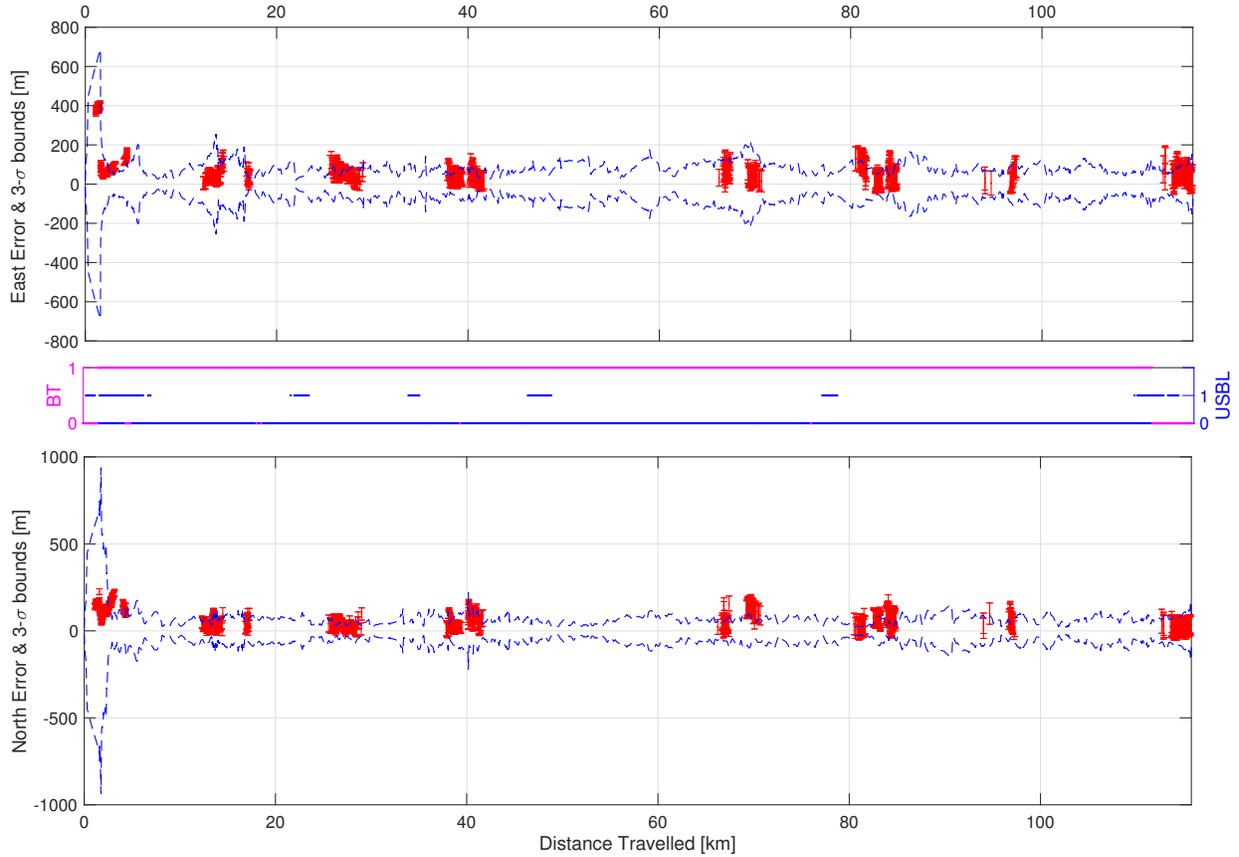


Figure 13: M42 - TAN time-indexed RMSE (red) and the associated 3- σ bounds (blue) over distance travelled and north and east direction. The uncertainty in USBL measurements is represented by error bars attached to each TAN error estimate.

in the 2D representation of the vehicle path), Figure 15.(a) shows the DR and TAN time-indexed RMSE ($RMSE_k^{20}$) evolution over DT, including the upper and lower observed TAN error bounds of all MC runs. Despite the extended duration of the mission, DR estimates do not experience high growth in error as a result of the geometry of the vehicle path (due to the closed path that cancels out systematic errors). The DR positioning error is primarily accumulated during the descent to a depth of approximately 1000 m. Once within bottom-tracking range, the vehicle is unaffected by water-currents and the rate of the increase of the DR error is significantly reduced. However, when the vehicle ascends from approximately 3200 m depth at the end of the mission, DR estimates drift again due to strong water-currents (occasionally over 0.6 m/s), which results in 1198.7 m error at surface.

Once water-depth measurements are obtained, the TAN algorithm is able to bound the positioning error below 200 m for almost the entire mission (see Figure 15.(a)), except for at the very beginning of the mission (approximately 6–7 km of the DT) when bottom-track is lost and the TAN error experiences an increase of 120 m. The average TAN RMSE, $RMSE^{20}$, when the ALR operates near-bottom is 152.4 m, while the corresponding DR error is 400.1 m. In terms of $RMSE_k^{20}$ at $k = k_{BT}^{USBL}$, the estimated TAN error is 126.1 m, which corresponds to 0.06% error of the DT, whereas the DR error is 398.1 m, corresponding to 0.21% error of the DT. It is also important to notice the ability of the filter to produce repeatable estimates. All 20 runs produce identical results as the time-indexed RMSE bounds are almost always convergent. This indicates that the number of particles is sufficient to produce robust results for M44, similarly to M41 and

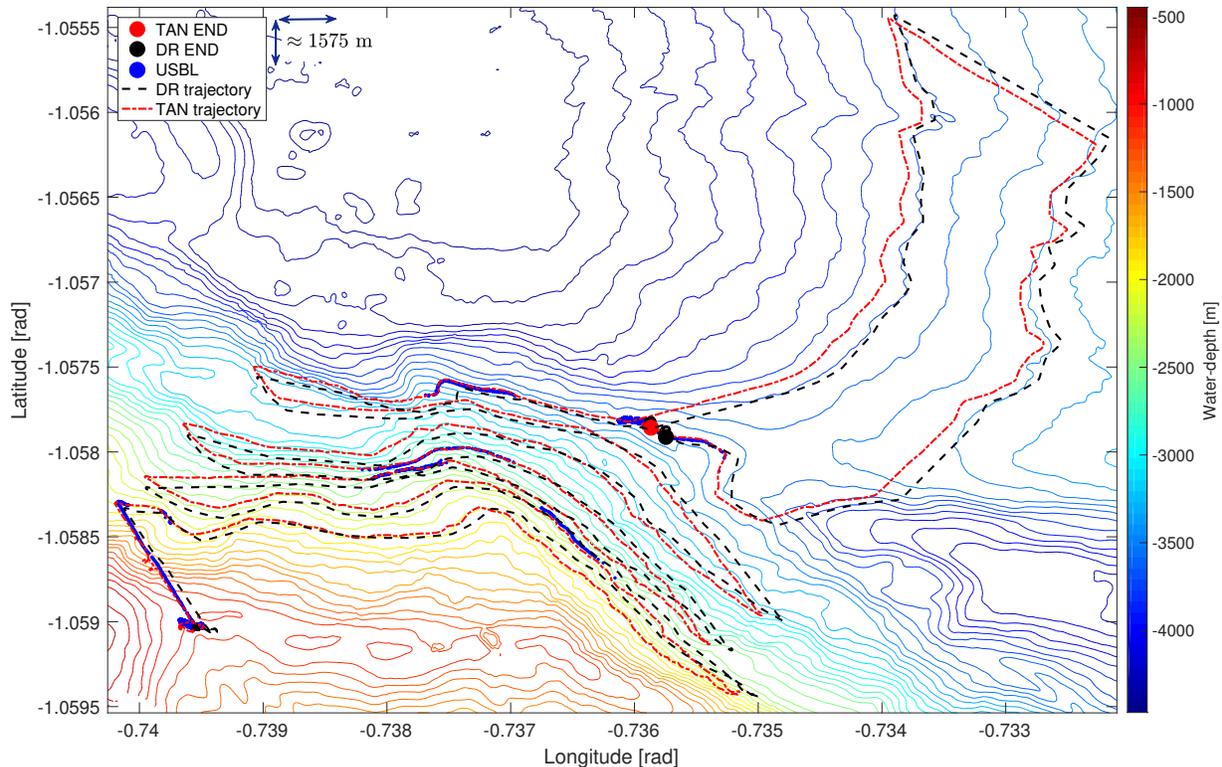


Figure 14: M44 - Estimated AUV path using TAN (red) and DR (black). There are 4077 available USBL measurements (blue), covering approximately 1.48% of the mission.

M42.

Figure 15.(c) shows the time-indexed AWDD over the DT, as calculated using TAN, DR and USBL position estimates. The average AWDD ($AWDD^{20}$) for the DR is 68.9 m, whereas TAN estimates and USBL measurements achieve 3.1 m and 8.4 m, respectively. Given the TAN's low time-indexed AWDD and RMSE, where the latter sometimes is lower than 50m, it is likely that the largest part of the estimated TAN error is driven by the noise in USBL positioning. For example, the high time-indexed AWDD for the USBL measurements at approximately 38 km and 82 km of the DT (see Figure 15.(c)) implies that these USBL measurements are potentially inaccurate.

Figure 16 shows the estimates of the TAN error for each direction with the respective two-sided confidence bounds over DT. Despite the strong assumption that particles are normally distributed and uncorrelated in each direction, the error estimates mostly fall within the $3\text{-}\sigma$ bounds. Note also that the selected process noise intensity during vehicle descent is sufficient to capture the error induced by water-currents during descent.

7 Conclusion

A PF-based TAN algorithm for the ALR long-range AUV has been presented. The algorithm compares water-depth measurements against an *a-priori* low-resolution bathymetric reference map to estimate the 2D position of the ALR. TAN results, obtained using field data from the Orkney Passage, demonstrate the potential of using geophysical localisation methods for bounding the growth in dead-reckoning error, thus enabling long-endurance missions.

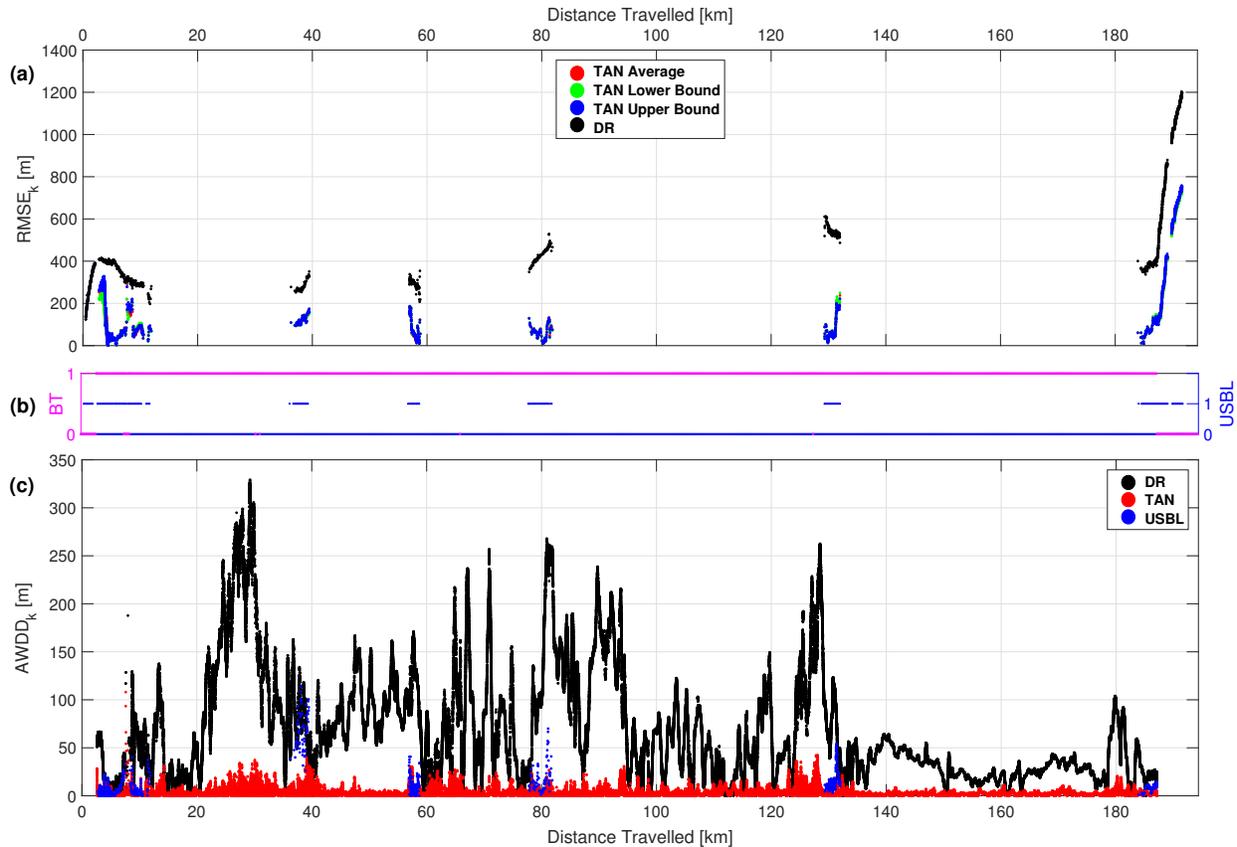


Figure 15: M44 - a) Time-indexed RMSE over distance travelled for TAN (red) and DR (black), along with lower (green) and upper (blue) TAN error bounds of 20 MC runs, b) BT mode and USBL measurement flags (with 1 indicating availability), and c) time-indexed AWDD over distance travelled for TAN (red), DR (black) and USBL (blue).

The TAN algorithm has been designed and built to accommodate the needs for extended duration underwater operations by providing a light and tractable navigation solution for a vehicle operating within bottom-tracking range. The proposed TAN system uses a small number of low-power navigation sensors, requiring no additional ranging sensors as single altitude measurements are obtained from an ADCP, already in-use to measure the bottom-relative speed of the ALR, significantly reducing the energy-cost of the presented approach.

The estimation efficiency and robustness of the algorithm were evaluated using field-data from three deep (up to approximately 3700 m) and long-range (up to 195 km in 77 hours) ALR operations. These were science missions, conducted as part of the DynOPO project, to measure water flow, temperature and underwater turbulence at the Orkney Passage, a region of the Southern Ocean (April 2017).

The parameters of the particle filter were defined based upon environmental characteristics and available on-board processing power. TAN results were compared to position estimates both on-board in real-time through dead-reckoning and via USBL measurements, used as baseline locations, intermittently recorded from a support ship during the missions. Despite the challenges, including the motion of the vehicle along low-informative terrain morphology for prolonged time intervals, post-processing results demonstrated that TAN holds potential to extend underwater missions to hundreds of kilometres without the need for surfacing to re-initialise the estimation process. Although the ALR was equipped with a limited suite of motion sensors and only had access to a low resolution reference map, the TAN algorithm maintained convergence and a

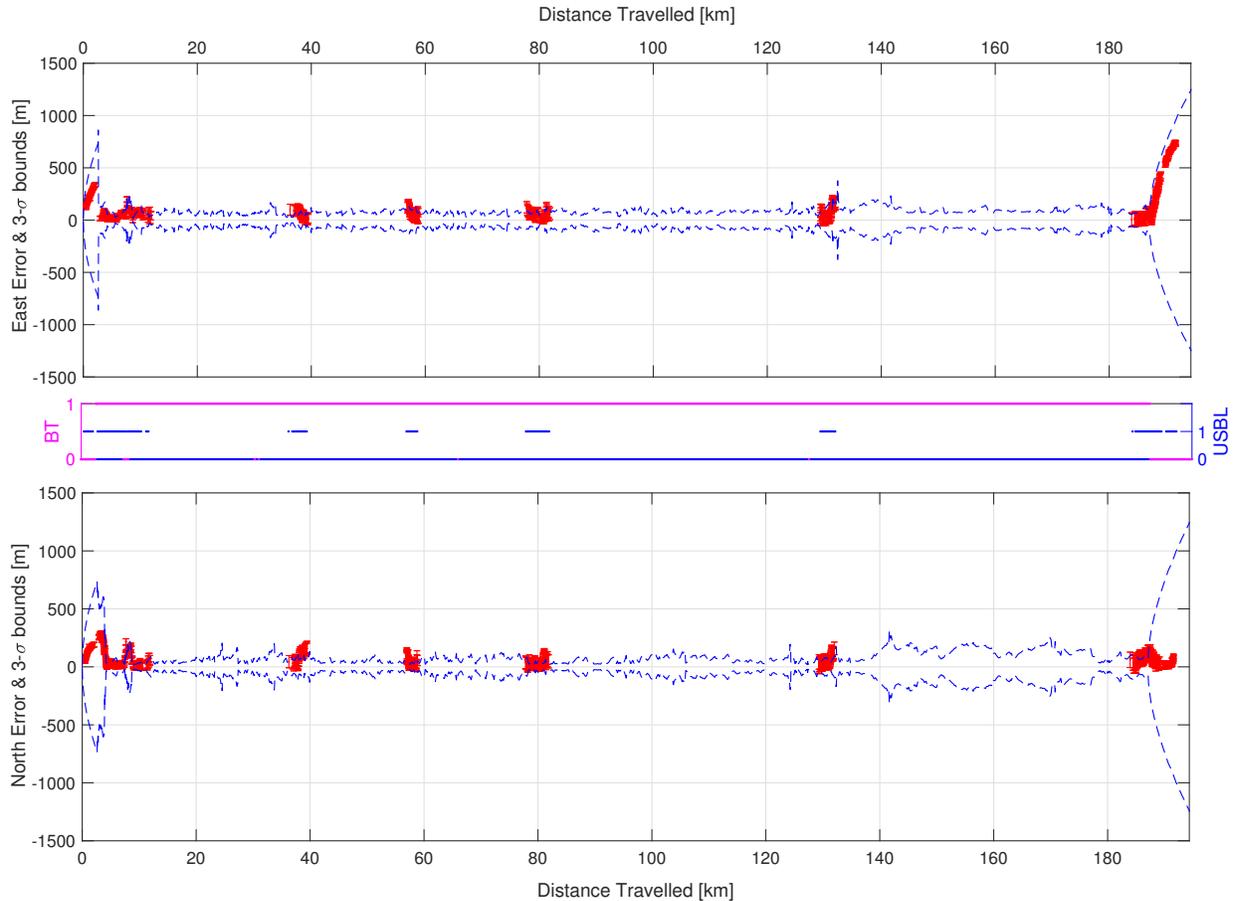


Figure 16: M44 - TAN time-indexed RMSE (red) and the associated 3- σ bounds (blue) over distance travelled and north and east direction. The uncertainty in USBL measurements is represented by error bars attached to each TAN error estimate.

bounded estimation error in all three missions over all performed Monte Carlo runs.

7.1 Future Work

Finally, there are some important aspects that must also be considered before implementing TAN in real-time for long-range missions.

Probably the most important property for a reliable TAN system is the PF robustness against unexpectedly noise-corrupted sensor measurements (e.g. acoustic signal interference) and/or regions of poorly mapped terrain. In either such events, a TAN system might become prone to divergence, which eventually may result in the need to either reset and re-initialise the PF algorithm or to surface and obtain a GPS fix. However, surfacing is not an available option for operations in the 12% of the world's oceans covered by ice. Hence, extremely challenging missions, such as the crossing of the Arctic ocean (Salavasidis et al., 2018), must have an enhanced robustness navigation system with the ability to compensate for potential sensor errors (e.g. automatic sensor re-calibration, similar to the techniques used by the existing ALR navigation system), to ignore temporarily poor sensor readings, reject wrong position estimates and, if divergence suspected, to re-initialise the PF algorithm.

The rejection of inaccurate position estimates and the maintenance of sufficient navigation accuracy is also of high importance for an AUV control system. Sample-based techniques often compute a multi-modal posterior density function. Given this multi-modal distribution, position estimates calculated by the mean of the posterior (MMSE estimates) might be highly biased, causing significant jumps in the position estimate. It is therefore crucial for the AUV control system to eliminate poor position estimates and discontinuities in the estimated trajectory, as these may result in a non-optimal control performance or even to control instability. As a consequence, actions to prevent or discount inaccurate position estimates are required. Such actions could be performed by a mechanism that monitors the filter convergence and confidence over the position estimates. This mechanism essentially provides a method to only accept TAN estimates when the state posterior probability density is approximately unimodal and the filter sufficiently confident over the position estimates. For example, immediately after the filter initialisation, the filter estimates are typically highly uncertain (see Figure 10) and it is likely that the posterior distribution is multi-modal. In this instance, the aforementioned monitoring mechanism would reject the TAN estimates and the navigation system would instead resort to an alternative on-board solution, avoiding poor and jumpy estimates. Therefore, instead of directly using a stand-alone TAN, an integrated solution is required. This might be a tightly-coupled TAN-DR navigation system which switches to the DR solution when TAN estimates are rejected, resulting in a significantly smoother vehicle trajectory utilised by the control system.

A similar mechanism is considered in (Meduna, 2011; Ånonsen et al., 2013). In the latter, the navigation system consists of a loosely coupled INS-TAN for the HUGIN AUV. In this case, the unimodal distribution test becomes crucial since a KF filter is used, which assumes Gaussian error models, for the INS updates. To further test the validity of the TAN estimates, the HUGIN navigation system passes the estimates through additional checks. These checks include calculating statistics over the obtained water-depth measurements and the reference map in order to prevent feeding the employed INS with poor TAN estimates.

It is also important to utilise a technique for detecting false convergence of the filter. One method might be by monitoring the summation of the weight of the particles, which might indicate a wrong convergence if it gets extremely small values for prolonged periods of time. In such an event, a re-initialisation of the particles might be necessary to enable the filter to correctly re-convergence (Teixeira et al., 2017a).

Acknowledgement

This work was supported by the ROBOCADEMY (FP7 Marie Curie Programme ITN Grant Agreement Number 608096) and NERC Oceanids programme. Authors would also like to thank the officers, crew, and scientific party of RRS James Clark Ross cruise JR16005 (DynOPO) for assistance in acquiring the data used here.

References

- Ånonsen, K. B., Hagen, O. K., Hegrenæs, Ø., and Hagen, P. E. (2013). The HUGIN AUV terrain navigation module. In *Oceans-San Diego, 2013*, pages 1–8. IEEE.
- Anonsen, K. B. and Hallingstad, O. (2006). Terrain Aided Underwater Navigation using Point Mass and Particle Filters. In *Proceedings of the IEEE/ION Position Location and Navigation Symposium*, pages 1027–1035.
- Arulampalam, M. S., Maskell, S., Gordon, N., and Clapp, T. (2002). A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on signal processing*, 50(2):174–188.
- British Antarctic Survey (2017). Dynamics of the Orkney Passage (DynOPO). <https://www.bas.ac.uk/project/dynamics-of-the-orkney-passage-outflow/>. Accessed: 24-October-2017.
- Caiti, A., Crisostomi, E., and Munafo, A. (2010). Physical Characterization of Acoustic Communication Channel Properties in Underwater Mobile Sensor Networks. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Carpenter, J., Clifford, P., and Fearnhead, P. (1999). Improved particle filter for nonlinear problems. *IEE Proceedings-Radar, Sonar and Navigation*, 146(1):2–7.
- Carreno, S., Wilson, P., Ridao, P., and Petillot, Y. (2010). A survey on terrain based navigation for AUVs. In *OCEANS 2010*, pages 1–7. IEEE.
- Claus, B. and Bachmayer, R. (2015). Terrain-aided Navigation for an Underwater Glider. *Journal of Field Robotics*, 32(7):935–951.
- Curcio, J., Leonard, J., Vaganay, J., Patrikalakis, A., Bahr, A., Battle, D., Schmidt, H., and Grund, M. (2005). Experiments in moving baseline navigation using autonomous surface craft. In *OCEANS, 2005. Proceedings of MTS/IEEE*, pages 730–735. IEEE.
- Donovan, G. T. (2011). Development and testing of a real-time terrain navigation method for AUVs. In *OCEANS 2011*, pages 1–9. IEEE.
- Doucet, A., De Freitas, N., and Gordon, N. (2001). An introduction to sequential Monte Carlo methods. In *Sequential Monte Carlo methods in practice*, pages 3–14. Springer.
- Doucet, A., Godsill, S., and Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and computing*, 10(3):197–208.
- Ekütekin, V. (2007). Navigation and control studies on cruise missiles. *Middle East Technical University: PhD Thesis*.
- Fairfield, N. and Wettergreen, D. (2008). Active localization on the ocean floor with multibeam sonar. In *OCEANS 2008*, pages 1–10. IEEE.
- Fofonoff, N. P. and Millard Jr, R. (1983). Algorithms for the computation of fundamental properties of seawater. *UNESCO Technical Papers in Marine Science 44*.
- Freitag, L., Koski, P., Morozov, A., Singh, S., and Partan, J. (2012). Acoustic communications and navigation under Arctic ice. In *Oceans, 2012*, pages 1–8. IEEE.
- Furlong, M. E., Paxton, D., Stevenson, P., Pebody, M., McPhail, S. D., and Perrett, J. (2012). Autosub Long Range: A long range deep diving AUV for ocean monitoring. In *Autonomous Underwater Vehicles (AUV), 2012 IEEE/OES*, pages 1–7. IEEE.
- Golden, J. P. (1980). Terrain contour matching (TERCOM): A cruise missile guidance aid. *Image processing for missile guidance*, 238:10–18.

- Google Earth (2017). Orkney Passage, Antarctic Peninsula, $59^{\circ} 13' 20''\text{S}$, $49^{\circ} 57' 35''\text{W}$, 3,504.29 km. <https://earth.google.com/web/@-58.24792076,-54.57702361,-3401.62319381a,3512896.3445127d,35y,0h,0t,0r>. Accessed: 10-August-2017.
- Hampel, F. R. (1974). The influence curve and its role in robust estimation. *Journal of the American statistical association*, 69(346):383–393.
- Jakuba, M. V., Roman, C. N., Singh, H., Murphy, C., Kunz, C., Willis, C., Sato, T., and Sohn, R. A. (2008). Long-baseline acoustic navigation for under-ice autonomous underwater vehicle operations. *Journal of Field Robotics*, 25(11-12):861–879.
- Jensen, F. B., Kuperman, W. A., Porter, M. B., and Schmidt, H. (2011). *Computational ocean acoustics*. Springer Science & Business Media.
- Joyce, T. M. (1989). On in situ calibration of shipboard adcps. *Journal of Atmospheric and Oceanic Technology*, 6(1):169–172.
- Karlsson, R. and Gustafsson, F. (2003). Particle filter for underwater terrain navigation. In *Statistical Signal Processing, 2003 IEEE Workshop on*, pages 526–529. IEEE.
- Kinsey, J. C., Eustice, R. M., and Whitcomb, L. L. (2006). A survey of underwater vehicle navigation: Recent advances and new challenges. In *IFAC Conference of Manoeuvring and Control of Marine Craft*, volume 88, pages 1–12.
- Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of computational and graphical statistics*, 5(1):1–25.
- Lagadec, J. (2010). Terrain Based Navigation using a Particle Filter for Long range glider missions-Feasibility study and simulations. *Masters thesis*.
- Mandt, M., Gade, K., and Jalving, B. (2001). Integrating DGPS-USBL position measurements with inertial navigation in the HUGIN 3000 AUV. In *Proceedings of the 8th Saint Petersburg International Conference on Integrated Navigation Systems, Saint Petersburg, Russia*, pages 28–30.
- Marks, K. and Smith, W. (2008). An uncertainty model for deep ocean single beam and multibeam echo sounder data. *Marine Geophysical Researches*, 29(4):239–250.
- McPhail, S. D. and Pebody, M. (1998). Navigation and control of an autonomous underwater vehicle using a distributed, networked, control architecture. *Underwater Technology*, 23(1):19–30.
- Meduna, D. K. (2011). *Terrain relative navigation for sensor-limited systems with application to underwater vehicles*. PhD thesis, Stanford University.
- Melo, J. and Matos, A. (2017). Survey on advances on terrain based navigation for autonomous underwater vehicles. *Ocean Engineering*, 139:250–264.
- Morice, C., Veres, S., and McPhail, S. (2009). Terrain referencing for autonomous navigation of underwater vehicles. In *Oceans 2009-Europe*, pages 1–7. IEEE.
- Munafò, A. and Ferri, G. (2017). An Acoustic Network Navigation System. *Journal of Field Robotics*, 34(7):1332–1351.
- Naveira Garabato, A. and Forryan, A. (2017). RRS James Clark Ross Cruise JR16005: Dynamics of the Orkney Passage Outflow (DynOPO) cruise report. University of Southampton, National Oceanography Centre, Southampton, U.K.
- Nordlund, P.-J. (2002). *Sequential Monte Carlo Filters and Integrated Navigation*. PhD thesis, Department of Electrical Engineering, Linköpings University.
- Nygren, I. and Jansson, M. (2004). Terrain navigation for underwater vehicles using the correlator method. *IEEE Journal of Oceanic Engineering*, 29(3):906–915.

- Pascoal, A., Oliveira, P., Silvestre, C., Sebastião, L., Rufino, M., Barroso, V., Gomes, J., Ayela, G., Coince, P., Cardew, M., et al. (2000). Robotic ocean vehicles for marine science applications: the european asimov project. In *Oceans 2000 MTS/IEEE Conference and Exhibition*, volume 1, pages 409–415. IEEE.
- Pascoal, A., Silvestre, C., and Oliveira, P. (2006). Vehicle and Mission control of single and multiple autonomous marine robots. *IEE Control Engineering Series*, 69:353–386.
- Paull, L., Saeedi, S., Seto, M., and Li, H. (2014). AUV navigation and localization: A review. *Oceanic Engineering, IEEE Journal of*, 39(1):131–149.
- Phillips, A., Haroutunian, M., Man, S., Murphy, A., Boyd, S., Blake, J., and Griffiths, G. (2012). *Nature in engineering for monitoring the oceans: comparison of the energetic costs of marine animals and AUVs*, pages 373–405. 77. The Institution of Engineering and Technology.
- Phillips, A., Haroutunian, M., Murphy, A. J., Boyd, S., Blake, J., and Griffiths, G. (2017). Understanding the power requirements of autonomous underwater systems, Part I: An analytical model for optimum swimming speeds and cost of transport. *Ocean Engineering*, 133:271–279.
- Purkey, S. G. and Johnson, G. C. (2010). Warming of global abyssal and deep Southern Ocean waters between the 1990s and 2000s: Contributions to global heat and sea level rise budgets. *Journal of Climate*, 23(23):6336–6351.
- Purkey, S. G. and Johnson, G. C. (2013). Antarctic Bottom Water warming and freshening: Contributions to sea level rise, ocean freshwater budgets, and global heat gain. *Journal of Climate*, 26(16):6105–6122.
- Roper, D. T., Phillips, A. B., Harris, C., Salavasidis, G., Pebody, M., Templeton, R., Vikraman, S. S. A., Smart, M., and McPhail, S. (2017). Autosub Long Range 1500: an ultra-endurance AUV with 6000Km range. In *MTS/IEEE Oceans’17-Aberdeen*. IEEE.
- Salavasidis, G., Harris, C., McPhail, S., Phillips, A. B., and Rogers, E. (2016a). Terrain-Aided Navigation for long range AUV operations at Arctic latitudes. In *Autonomous Underwater Vehicles (AUV), 2016 IEEE/OES*, pages 115–123. IEEE.
- Salavasidis, G., Harris, C. A., Rogers, E., and Phillips, A. B. (2016b). Co-operative Use of Marine Autonomous Systems to Enhance Navigational Accuracy of Autonomous Underwater Vehicles. In *Conference Towards Autonomous Robotic Systems*, pages 275–281. Springer.
- Salavasidis, G., Munafò, A., Harris, C., McPhail, S. D., Rogers, E., and Phillips, A. B. (2018). Towards Arctic AUV Navigation. In *IFAC Conference on Control Applications in Marine Systems*. Elsevier.
- Schon, T. B., Karlsson, R., and Gustafsson, F. (2006). The marginalized particle filter in practice. In *Aerospace Conference, 2006 IEEE*, pages 11–pp. IEEE.
- Siouris, G. M. (2004). *Missile guidance and control systems*. Springer Science & Business Media.
- Stammer, D., Ray, R., Andersen, O. B., Arbic, B., Bosch, W., Carrère, L., Cheng, Y., Chinn, D., Dushaw, B., Egbert, G., et al. (2014). Accuracy assessment of global barotropic ocean tide models. *Reviews of Geophysics*, 52(3):243–282.
- Tan, Y. T., Chitre, M., and Hover, F. S. (2016). Cooperative bathymetry-based localization using low-cost autonomous underwater vehicles. *Autonomous Robots*, 40(7):1187–1205.
- Teixeira, F. C. (2007). *Terrain-Aided Navigation and Geophysical Navigation of Autonomous Underwater Vehicles*. PhD thesis, Dynamical systems and ocean robotics lab, Lisbon.
- Teixeira, F. C., Quintas, J., Maurya, P., and Pascoal, A. (2017a). Robust particle filter formulations with application to terrain-aided navigation. *International Journal of Adaptive Control and Signal Processing*, 31(4):608–651.

- Teixeira, F. C., Quintas, J., and Pascoal, A. (2016). AUV terrain-aided navigation using a Doppler velocity logger. *Annual Reviews in Control*, 42:166–176.
- Teixeira, F. C., Quintas, J., and Pascoal, A. (2017b). Robust methods of magnetic navigation of marine robotic vehicles. *IFAC-PapersOnLine*, 50(1):3470–3475.
- Vickery, K. (1998). Acoustic positioning systems. A practical overview of current systems. In *Autonomous Underwater Vehicles, 1998. AUV'98. Proceedings of the 1998 Workshop on*, pages 5–17. IEEE.