


Side Scan

SONAR IMAGE REGISTRATION

for AUV Navigation

by Peter Vandrish and
Andrew Vardy





PETER KING/MUN

MUN Explorer AUV platform with the Marine Institute's Holyrood Marine Base in the background.

Introduction

An Autonomous Underwater Vehicle (AUV) can be given sight through acoustic imaging systems. In this essay we discuss the usage of side scan sonar imaging for AUV navigation. In particular, we focus on image registration, which is the problem of identifying movement between two images. There has been much work on image registration in optical images. However, the application of image registration to side scan sonar images presents a number of challenges.

Memorial University of Newfoundland (MUN) is a leader in underwater vehicle operations and technologies. The Responsive Autonomous Underwater Vehicle Localization and Mapping (REALM) project led by Memorial University focuses on expanding the capabilities of the MUN Explorer AUV platform, in particular; and developing novel techniques for AUV navigation and surveying capabilities, in general. Navigation itself plays a very important role in enabling the autonomous operation of an AUV not only as a path planning subsystem, but also to ensure the safety and recovery of the vehicle. Traditional AUV localization techniques combine information from a variety of sensors. While on the surface, the AUV updates its position from the Global Positioning System (GPS). During a dive the AUV relies on its Inertial Navigation Unit (INU) to estimate its acceleration and integrate this information to update its position. In the vicinity of the seabed, a Doppler Velocity Log (DVL) can be used to estimate the speed over ground, which allows for increased precision in updating the vehicle's position. However, it is only GPS which provides an absolute sense of position, and even the position estimate from GPS will include some error. The INU and DVL can be used to update the position estimate, but it is inevitable that this estimate will drift from its true value. Therefore AUVs will surface periodically to obtain an absolute position fix from GPS.

A side scan sonar system provides an AUV with a framework for developing highly

sophisticated navigation systems. An AUV carrying a side scan sonar system within its payload can extract an enormous amount of information about its current surroundings in the form of sonar images. The objective of side scan sonar image registration is to accurately determine the mapping between the currently observed image, or sensed image, to a previously obtained reference image.

Application of Image Registration to AUV Navigation

Image registration can be used to provide an absolute sense of position with respect to a previously visited place. This has a number of potential applications. In Simultaneous Localization and Mapping (SLAM) the goal is to build a map of the environment while determining the vehicle's position within that map. Image registration can be used to determine both the change in the vehicle's position with respect to the existing part of the map, and also to determine how newly perceived parts of the environment should be added to the map.

The mobile robotics research community has been working on the problem of SLAM for two decades. Efficient and effective solutions have emerged. However, SLAM remains a computationally expensive process and it is not clear that large-scale navigation on an AUV will be possible with current algorithms. Therefore, in our project we are investigating an alternative strategy known as Qualitative Navigation. Qualitative Navigation is also an idea that has arisen from research on mobile robotics. It asserts that an autonomous robot can find its way throughout an environment without necessarily requiring a single global coordinate system for describing the locations of all landmarks. This contrasts with most work in autonomous robotics and in AUVs where the primary input to the navigational system is a description of landmarks and goals, all specified within a single global coordinate system. Our project focuses on autonomous route following of Memorial University's Explorer AUV within a Qualitative Navigation framework. An AUV capable of autonomously

following a specified route would have numerous applications. In particular, this is intended to address applications in infrastructure inspection and environmental monitoring.

Image registration forms a crucial part of our navigational strategy. The AUV is first guided along a route during the training phase. It may be necessary for the AUV to surface periodically to update its position estimate and therefore ensure that it is following the desired route. During the training phase the AUV will record the side scan sonar images encountered along the route. All of the steps described thus far can be managed with current technology. The next step is to follow the route again, autonomously and without surfacing. The AUV will continuously try to determine its position along the route. It is important to note that we only need to keep track of the percentage of the AUV's travel along the route. This is a 1-dimensional estimation problem that is much easier to handle than the problem of determining the AUV's position in 3-dimensional space. Once the AUV's position along the route is known, we will utilize side scan sonar image registration to determine the vehicle's displacement from the route. Once this displacement is known, the vehicle's trajectory can be altered to allow it to return to the desired route. Continuing in this manner, the AUV should be able to travel the entire route without needing to surface.

Side Scan Sonar

Side scan sonar systems produce high resolution image representations (or sonograms) of acoustic backscatter; see Figure 1. They operate by emitting a narrow (in the fore-aft direction) acoustic pulse perpendicular to the along-track direction vector and out each side of the sonar. The pulse spreads out over a wide angle in the across-track plane to ensonify [flood with sound waves] the seabed. The width along the seafloor covered by the acoustic pulse (ping) is referred to as the swath width. In the unprocessed received signal, regions along the swath which produce strong backscatter will



Figure 1: A typical side scan sonar image.

show up lighter within the image scanline. Protruding objects on the seabed will produce light regions followed by a dark shadow. There are numerous factors which come to impact the accuracy of image interpretation. They include changes to the velocity, orientation (roll, pitch, and yaw) and sway of the vehicle, acoustic noise, beam pattern, and the absorptivity and structure of the reflecting surfaces, for example. As a ping propagates away from the sonar, the beam width in the along-track direction expands. This in turn has the effect of a decreased along-track image resolution at points sufficiently far from the sonar. Multiple collinear landmarks running parallel to the along-track at a long range from the sonar could appear merged together in the final image. Obviously, one might think that this problem can be sidestepped by simply accepting data from within a shorter range. However, a similar problem exists which affects the across-track resolution. Depending on the pulse duration, multiple collinear landmarks along the scanline closer to the sonar could appear merged in the final image. Hence, there is a tradeoff between the resolutions in the across-track versus along-track directions. This is one challenge which could affect the reliability of image registration.

In terms of backscatter, large protruding objects will cast a dark "shadow" adjacent to the object. This occurs since no acoustic energy is capable of reaching this region. The horizontal length of this shadow will vary depending on where the object is in the across-track direction relative to the sonar. The shadow may also appear differently depending on the vantage point of the sonar. Thus, the

shadow may be a misleading indication of the height and shape of the object. In addition, areas of seabed which produce a very high specular reflection could also be mistakenly interpreted as a shadow. Another aspect which influences the amount of backscatter seen at the sonar is the structure and absorptivity of the seabed. Regions of seabed which are coarse will tend to reflect more energy back to the acoustic sensor and they will therefore show up lighter. Areas which absorb much of the energy will show up darker in intensity. These effects depend on the seabed material, the measure of coarseness, the frequency of the transmitted pulse, and the angle of incidence, among other factors. One consequence of this is that these regions show up as texture-like features in the resulting image.

Geometric image distortions can arise on account of the dynamics of the vehicle. Most sonar transducers, whether contained within an AUV's payload or part of a towfish construction, will usually undergo changes in orientation, velocity and sway. Some of these dynamics are more present in a towfish configuration since surge and heave in the towing ship will have a large impact on the towfish. For simplification, in this work we consider the AUV to be constrained by three degrees of freedom. They include the freedom of translatory motion in the x-y directions and the freedom to rotate within a 2D plane parallel to the surface at a constant depth. This constraint fits within our application of route following. The depth of the vehicle can be made constant during the training phase and it should be straightforward for the vehicle to maintain itself at this same depth while autonomously following the route.

Image Registration

Image registration techniques can be categorized as being either feature-based or area-based. It is also possible to further sub-categorize the techniques depending on whether they operate on the images as a whole or partition the images into regions of interest. Feature-based techniques, as the name implies,

attempt to identify interesting structure within an image such that the feature can be assigned a unique and robust description. This step is known as feature extraction. Image registration then involves a feature matching process which attempts to match features based on their descriptions. Once a set of candidate matches is constructed, geometric methods are applied to determine the transformation which maps the set of features observed in one image to the set of features observed in the second image. If the number of matches is high, and a transformation is found which maximizes the number of overlapping feature points between images, then the transformation is accepted. Area-based registration techniques attempt to register images by operating on pixel intensity values directly. A similarity metric is used to evaluate how accurate the transformation parameters map the sensed image onto the reference image. The metric may be based on correlation or minimum distance. Some techniques perform a transformation on the image signal prior to registration. An effective example of this is known as phase-correlation which performs the entire registration in the frequency domain. We have determined that two well known image registration techniques apply very well to side scan sonar images. One is based on a feature extractor called the scale-invariant feature transform (SIFT), while the other is the aforementioned phase-correlation which operates on the images globally and in the frequency domain.

SIFT was developed by David Lowe at the University of British Columbia. It has found wide application in computer vision research




Figure 2: Side scan sonar reference and sensed images with SIFT feature points overlaid.





Figure 3: Examples of seamlessly registered side scan sonar images.

and is well known as one of the most robust feature extraction methods available today. The accuracy of SIFT is dependent on the number of correct matches obtained between a sensed and reference image pair. The number of correct matches will, in turn, depend on the similarity between the images. Unfortunately, changes in the position of the AUV will change the position of the acoustic shadows of seabed objects which may cause false matches and may prevent true matches from being found. Figure 2 shows the features extracted from a pair of side scan sonar images using SIFT. Figure 3 provides the registration of those images based on the set of matched features.

Phase-correlation is a well known registration technique based on the Fourier transform. It is characterized as highly efficient and robust to non-uniform illumination variations. Coupling this technique with a number of image pre-processing stages makes it perform very well for side scan sonar images, including when the images contain a considerable amount of structure. Phase-correlation operates by firstly computing the Fourier transforms of two input signals; in this case, the signals are two dimensional images. The cross-power spectrum is determined by the multiplication of the

normalized spectrum of the first signal by the normalized complex conjugate of the second spectrum. The inverse Fourier transform of the cross-power spectrum is a two dimensional signal with a dominant peak located at a point which indicates the relative displacement between the two signals.

Conclusions

The real-time registration of side scan sonar images obtained by an AUV could provide a rich foundation for the development of sophisticated navigation algorithms. Not only can registration be used to provide feedback in the absence of a GPS signal, it may also be used as a standalone navigation solution. However, side scan sonar systems present a number of challenges to image registration algorithms when compared to optical imaging systems. Facing these challenges, SIFT and phase-correlation provide robust solutions given that the stability of the platform is assumed. Extending these concepts to the wider variety of side scan sonar image degradations is just around the corner. A complete navigation system based on the unification of computer vision and sonar imaging will continue to be the goal of our work at Memorial University of Newfoundland. ∞

Acknowledgements

This work has been supported by ACOA's Atlantic Innovation Fund, Fugro Geosurveys Inc. of St. John's, Canada, and DRDC Atlantic of Dartmouth, Canada.



Peter Vandrish is completing his M.Eng. under the supervision of Andrew Vardy, Octavia A. Dobre, and Dan Walker at Memorial University of Newfoundland. He obtained a bachelor of computer engineering from MUN in 2009. His research interests include the development of autonomous underwater vehicle navigation and computer vision algorithms.



Andrew Vardy is the Principal Investigator for the Responsive AUV Localization and Mapping project. He is an Assistant Professor jointly appointed with the Department of Computer Science and the Faculty of Engineering and Applied Science at Memorial University of Newfoundland. He completed his first degree in Electrical Engineering from Memorial University in 1999. He then studied Evolutionary and Adaptive Systems at the University of Sussex, earning an M.Sc. in 2000. In 2005 he completed a PhD in Computer Science at Carleton University. His research interests include robot navigation, spatial representation, and autonomous underwater vehicles.

The advertisement features a central 3D rendering of a yellow cylindrical mooring device with a blue top section containing various sensors and a grey base. The background is a blue, wavy water surface. The text is arranged as follows:

ROMOR
OCEAN SOLUTIONS

C-ROM
Compact Recoverable
Ocean Mooring

Reliable Recovery of
Subsea Instrumentation

www.romoroceansolutions.com