Autonomous identification planning for mine countermeasures

Martin Syre Wiig, Thomas R. Krogstad and Øivind Midtgaard
Forsvarets Forskningsinstitutt (FFI, Norwegian Defence Research Establishment)
P O Box 25, NO-2017 Kjeller, Norway
E-mail: martin-syre.wiig@ffi.no

***Abstract*—This paper presents a concept and algorithms to detect, classify and identify mine-like objects within a single mission with an autonomous underwater vehicle. The autonomous mine hunting concept has been developed for the HUGIN series of vehicles. First, the operation area is surveyed either with a synthetic aperture sonar or a side-scanning sonar. During the survey, mine-like objects are detected and classified in the data using algorithms for automatic target recognition. When the survey is complete, a framework for autonomy initiates a fusion of the targets and starts the automatic planning of a mission plan for target identification. The autonomous mine hunting concept is a part of the development of a framework for advanced autonomy on HUGIN, the HUGIN autonomy layer. Implementation of this framework will reduce the risk of long-term AUV missions, and will provide intelligent vehicle behavior not only to re-inspect interesting objects and areas, but also to preserve vehicle safety, navigational accuracy and mission goals.**

Keywords—Autonomous underwater vehicles; path planning; mine counter measures; optical identification; automatic target recognition

I. INTRODUCTION

Autonomous underwater vehicles (AUVs) are rapidly becoming an important tool in mine countermeasure (MCM) operations, and have been used by the Royal Norwegian Navy (RNoN) for several years. The early adoption of AUVs in the RNoN Mine Warfare Flotilla was in part caused by tough national requirements for deep-water mine hunting. It was, however, realized early on that AUVs also could provide MCM capabilities in other areas[1]

Currently, a typical MCM mission that uses an AUV consists of a series of sequential operations. First, the area is surveyed using an AUV equipped with either a synthetic aperture sonar (SAS) or a side-scanning sonar. Once the survey is completed, the AUV is recovered and a human operator looks through the sensor data to detect and classify targets, possibly aided by automatic target recognition (ATR) algorithms.

Next, the targets are identified using either mine divers, a remotely operated vehicle (ROV) equipped with a camera or an AUV equipped with a camera. Of these, the most time efficient is using an AUV. However, planning a safe camera mission that ensures coverage of all targets can be very time consuming.

Finally, the targets identified as mines are neutralized using mine divers or a one-shot mine neutralization vehicle.

In this paper we demonstrate a mission concept that allows the sonar survey and photo identification part of an AUV MCM mission to be performed in a single dive. This requires automatic analysis of the gathered sonar data using ATR algorithms, described in Section III. An autonomy framework is required to tell the system which part of the mission constitutes the survey, and to initiate target fusion and path planning once the survey is complete. This framework is described in Section IV. Finally, a path planning algorithm for generating the target inspection mission plan is required. In this paper, a genetic algorithm described in Section V is used.

The system was tested during a sea trial outside Horten Norway, and it proved itself capable of successful inspection of a set of targets, providing the optical imagery required for target identification. The results are presented in Section VI. Finally, Section VII contains some concluding remarks and suggestions for further research.

II. THE HUGIN VEHICLE

The vehicle considered in this paper is the RNoN's HUGIN 1000 V0 vehicle, shown in Figure 1.

The HUGIN 1000 V0 is rated to 1000m depth, and is equipped with an Edgethech 4300-MPX sidescanning sonar, an EM3000 multi-beam echosounder and a Norsk Elektro Optikk Tilefish optical camera.

The vehicle is equipped with two pressure spheres for the two main computers, the payload processor and the control processor, as well as for other electronics. Additional pressure containers are added as needed for any payload sensor processing units. Otherwise, the system operates at the environmental pressure.

Figure 1: The HUGIN 1000 V0 vehicle

A. Software system

A description of the basic design of the HUGIN software system is given in [2] The system consists of:

- A control processor with a real-time operating system, which handles guidance and control, basic sensors, collision avoidance, error and command handling and communication.
- A navigation processor, which handles the IMU and is running the real-time aided inertial navigation system.
- A payload processor with a Win32 compatible OS, which handles all other payload sensors that are not part of the basic sensor suite.
- A top side HUGIN operator station that handles mission planning and monitoring, and communication with the vehicle.

B. Payload sensors

The HUGIN 1000 V0 is equipped with a variety of sensors. For this experiment, it was equipped with a side-scanning sonar for object detection and classification. The sonar is an Edgetech MPX-4300. At 4 knots speed the sonar has a swat with of around 190m, a range resolution of 4-8 cm and an along-track resolution of 35-60cm [3].

Current MCM procedures require object identification to be performed using optical imaging, thus an optical camera is the main sensor for identification. The vehicle was equipped with a Norsk Elektro Optikk (NEO) Tilecam optical camera system, which has been specially designed for the HUGIN 1000 AUV through a collaborative development effort between NEO, FFI and Kongsberg Maritime. The system consists of a high resolution, high-sensitivity, digital still image camera and a synchronized strobe with multiple light emitting diodes (LEDs). This light source is ideal for AUVs, due to the LED's high energy efficiency and the use of pulsed, rather than continuous lightning. Also, the small LEDs can be built into a narrow panel fitted into the vehicle hull[4].

The camera and strobe are mounted with maximum separation distance on the AUV, thus reducing the loss of image contrast due to backscattered light from particles in the water, the setup is illustrated in Figure 2. The system uses LEDs with blue-green light to minimize absorption in the water and a monochromatic camera. Figure 3 shows an example photo cylindrical object that was classified as mine-like during a previous demonstration[5].

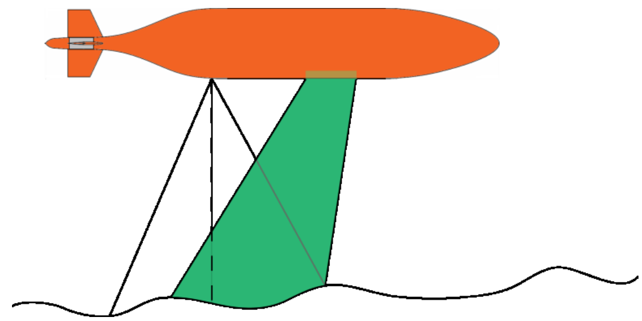

Figure 2: Illustration of the camera and strobe mounting on HUGIN AUVs

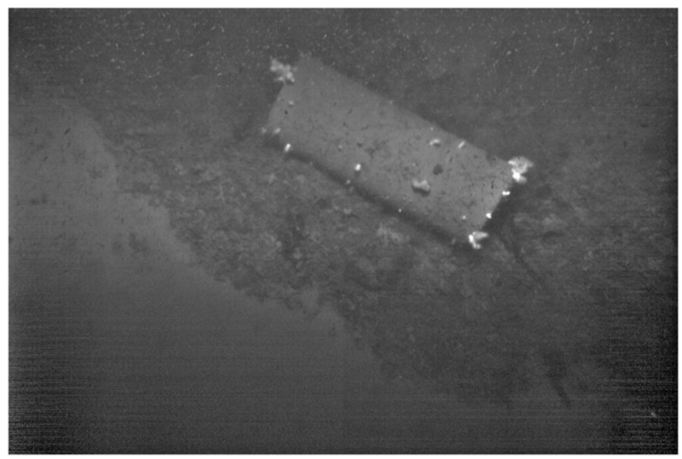

Figure 3: Photo of an unknown cylinder classified as mine-like

III. AUTOMATIC TARGET RECOGNITION

SITAR is an automatic target recognition (ATR) system for detection and classification of seafloor mines in synthetic aperture sonar (SAS) and side-scan sonar (SSS) imagery. It has been developed by FFI over the last decade, and can be used both for in-mission and post-mission analysis. The concept for in-mission operation is shown in Figure 4. The sonar processing converts sonar raw data into SSS or SAS image files, in the case of this paper to SSS files. An ATR control modules monitors the processing output, registers new image files and sequentially sends them to the ATR processing (SITAR) after each survey line. The control module can also distribute ATR results to the topside operator station on the host vessel using acoustic, radio or satellite link.

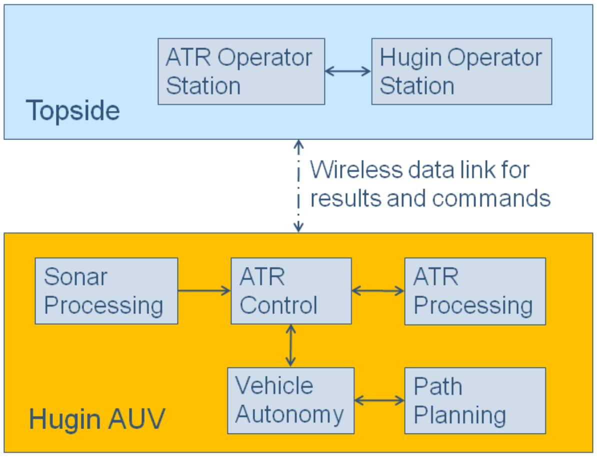

Figure 4: In-mission ATR software overview

SITAR applies multiple detection filters on the full sonar data set to search for mine-sized object responses[6]. Each detected contact is assigned a confidence value, indicating its degree of mine-likeness. Multi-filter contacts corresponding to the same seafloor object are then fused using Dempster-Shafer's theory of evidence. Contacts with fused confidence value above a (low) threshold are then input to the classifier. This procedure provides a high probability of target detection with a moderate false alarm rate.

The SITAR classifier contains a large set of template images, showing reference responses for all specified target types, target orientations, sonar ranges and sonar altitudes[7]. Currently, three mine shapes are included: Cylinder, truncated cone and wedge. For each detected contact, the classifier selects the appropriate subset of templates corresponding to the actual sonar range and altitude, and compares these templates with the measured sonar image. Based on the template similarity scores, confidence values are assigned to all mine classes and also to a non-mine (clutter) class.

SITAR is applied after each line to the sonar images produced during that line. After the survey is complete, SITAR is again applied, this time to the full sonar data set to achieve better performance. Survey completion can either be marked in the mission plan, or decided by the autonomy layer. This approach allows both immediate interruption of a survey to inspect mine-like objects rapidly, or complete survey where the most interesting objects in an area are all inspected.

IV. VEHICLE AUTONOMY

A module for enhanced vehicle autonomy, the HUGIN autonomy layer (HAL), has been developed at FFI [8]. HAL implements a hybrid control architecture [9, 10] to implement mission plan parsing, automatic mission planning of high level mission goals, and *in-situ* mission adaptation to counter environmental changes.

The structure of HAL is shown in Figure 5. A mission module is at the top of a hierarchical structure, and is responsible for the highest level of mission planning and user interaction. The module receives a set of mission goals, which is used to generate a set of mission tasks that constitute an intermediate layer of planning. The tasks form the lower part of the hierarchy in which the mission part is on top. Together, the mission module and the tasks form the deliberate layer of the autonomy system.

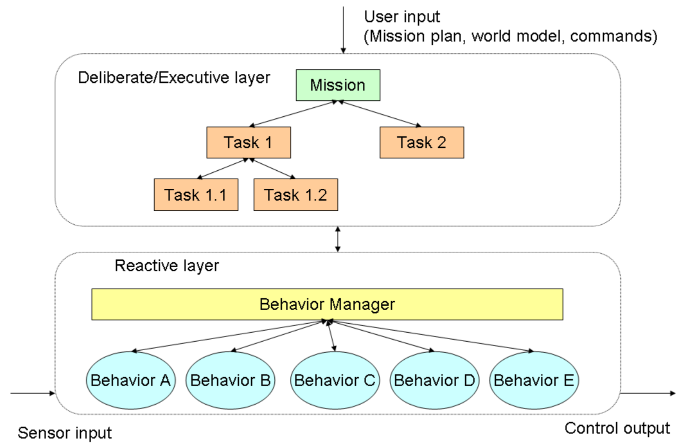

Figure 5: Structure of the HUGIN autonomy layer - HAL

The lower, reactive layer of HAL is the layer closest to the traditional guidance and control system of HUGIN. The reactive layer consists of a set of behaviors; parallel processes that each are trying to control the vehicle to perform a single function. Typical examples of behaviors are collision avoidance and moving towards a target. The output of the behaviors may be conflicting, and must be coordinated in a rational manner that maximizes task accomplishment and vehicle safety. By employing different sets of behaviors, the different tasks are thus able to achieve their goals. By ensuring that the tasks interacting with the reactive layer have as simple and straight forward goals as possible, the issue of an overly complex reactive layer is avoided.

A. *Autonomous surveys*

A typical mission plan when performing a side scan survey will be a set of parallel track lines in what is known as a lawnmower pattern. HAL allows such a survey to be defined as a task. This allows the user to define the start and end of the surveying part of a mission, which serves as input to the ATR system as described in Section III.

A major issue when planning survey missions is to estimate the sonar minimum and maximum range. The range can vary depending on a variety of factors such as the sea state, sea bed type, bathymetry, water depth and vehicle altitude. An operator must often use heuristics to determine the ranges, and errors in range estimation result in either holes in the coverage or a higher than necessary time consumption.

Several HUGIN vehicles are equipped with the Kongsberg Maritime HISAS 1030 synthetic aperture sonar[11], where range can be estimated using interferometric coherence[12]. Interferometric coherence is a measure of the similarity between the complex data from vertically displaced sonar receivers. A low level of similarity in a range bin might imply a low signal to noise ratio, and that the range bin is outside of the acceptable range interval of the sonar.

A type of survey task that has been designed is the survey with adaptive linespacing, where coherence values are estimated along survey lines, thus allowing achieved range and coverage to be estimated. This estimate is in turn used to calculate the distance to the next survey line.

The adaptive survey algorithm is run repeatedly until the entire mission area is covered with the desired amount of overlap. If the in-mission ATR algorithm has been running in the background, the autonomy system will at this point fuse the detections and classifications of each survey line and start the path planning described in Section V.

B. Interruptions

During a mission, especially a long duration mission, the need for a temporary interruption of a task to perform a more immediate task is sometimes needed. An example of this is a navigational fix, where the current goal of HUGIN is temporarily abandoned to improve the navigational accuracy. If the vehicle is in shallow waters, HAL will solve this by pausing the current task, move to the surface to acquire a GPS fix and then resume task in such a manner that there are no holes in the coverage. An evaluation of the procedure is performed during the fix, and if the GPS is found to be broken or no improvement in navigational accuracy was achieved the procedure will not be repeated.

A similar procedure is the sound velocity profile (SVP) fix. Before a HUGIN mission an up to date SVP is uploaded to the vehicle. However, the actual SVP may change over time and as the vehicle moves to different areas. If the autonomy system detects that the loaded SVP differs from the measured sound velocity, it will initiate the SVP fix and read the sound velocity at several different depths. As with the GPS fix, the current task will be paused during the fix, and resumed afterwards in a manner that ensures no holes in the coverage.

In an exploratory mission, in which the goal is to determine whether or not an area contains any mines rather than to find all of them, the adaptive survey can be aborted when one or more minelike objects with a certain confidence has been found on a survey line. The path planning algorithm can then be started using the objects found on the survey line as input.

V. PATH PLANNING USING A GENETIC ALGORITHM

A. Problem definition

When the survey mission has concluded and the respective sonar data have been processed by the ATR, a list of classified targets is generated. Each target in this list must be imaged optically to provide the operator with sufficient information to identify the target. Hence, the autonomy system must use the target list and produce a mission plan within some given time limit. The plan itself must not only be feasible with respect to bathymetry, coast-lines, and vehicle manoeuvrability but also image the objects in a time optimal manner.

The trivial answer to this planning problem is to find the shortest-path which visits all of targets at least once. This is recognized as the travelling salesman problem (TSP), one of the most widely studied combinatorial optimization problems. The problem statement may seem trivial, and yet it remains one of the most challenging problems in operational research. Though the TSP problem is NP-complete and not solvable in polynomial time, several exact and approximate solutions exist for problems of limited size [13] .

In this paper we consider path-planning for an under-actuated AUV, i.e. we can only control speed and heading in the horizontal plane. This means the vehicle has a certain minimum turning radius, which must be considered by the path-planning approach. In addition, to ensure that the vehicle has converged to the waypoint line when traveling over the target, we need to ensure that the waypoint line passing over the target has the necessary lead in and lead out, Figure 6. When defining our optimization problem this must be taken into account.

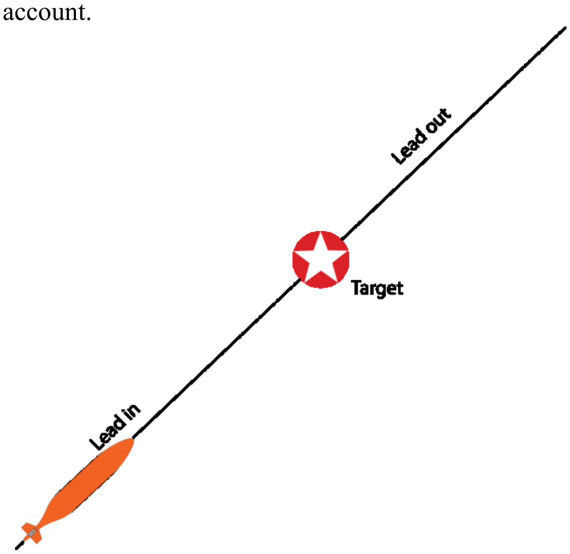

Figure 6: Vehicle path over target

B. Genetic Algorithms

Genetic algorithms are search procedures based on the mechanics of natural selection and genetics [14]. To solve a problem using genetic algorithms, we first encode solutions to the problem as an artificial chromosome or chromosomes. These chromosomes can be represented by strings of genes, which hold represent information about the solution, 0s and 1s, parameter lists or more complex structures. In the initialization of the algorithm a given number of chromosomes, usually denoted the individuals are chosen at random to form the population on which the algorithm will operate.

To know which solution best solves the problem we need some way of distinguishing good solution from bad. This can by anything from a simple mathematical expression to complex computer simulations, but must result information about the relative *fitness* of the individuals in the population.

The genetic algorithm then obtains a solution by gradually evolving and improving the population by means of genetic operators. In a typical genetic algorithm three such operators are used: (1) Selection, (2) Reproduction and (3) Mutation. The selection process selects which individuals should be allowed to reproduce. This can be accomplished in a variety of ways, but with the main idea of selecting more fit individuals over worse ones.

The second genetic operator, reproduction, recombines parts of two individuals' chromosome to form a new, possibly better individual and thereby evolve the population towards some local or global optimum. These operators, usually referred to as crossover operators, must be designed with the specific problem encoding in mind.

Reproduction by itself constitutes no more than a random shuffling of the components present in the initial population. We therefore also include a third genetic operation, called mutation. Mutation acts on a single individual by changing on or more traits.

If designed properly the genetic operators will gradually improve the initial population towards the optimal solution.

C. Genetic algorithm for TSP

For the travelling salesman problem the solution is encoded as a string of two properties: *Target ID* and *lead in enabled.* For such a string to be a valid TSP tour each target ID must appear only once. For a 5 target problem this encoding could be

$$\begin{matrix} 3 & 2 & 0 & 4 & 1 \\ 0 & 0 & 1 & 1 & 0 \end{matrix}$$

Where the first line encodes the target order and the second line lead in enabled or disabled. Lead in enabled affects how the tour distance is calculated by the fitness function as seen in Figure 7.

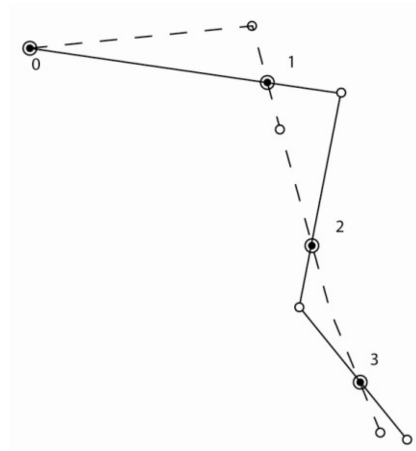

Figure 7: Illustration of the effect of having lead in enabled (dashed) and disabled (solid) for target 1. By adding a lead in point, which calculated as a point shift 100 m along the line connecting target 1 and 2, the vehicle turning is reduced, which will be reflected in the path's calculated fitness.

The fitness function is calculated as the inverse of the route cost, i.e. higher fitness amounts to a better route. The route cost reflects both the route distance and additional route weights which are added to discourage sharp turns (<90º) and encourage targets which are visited along a straight line. Moreover, when calculating route distances between two targets, we also consider which target we are coming from and which target we are going to. Since this affects the calculation of lead in and lead out point as illustrated in Figure 7.

In our implementation of the genetic algorithm we use the weighted roulette wheel selection method. This selection process involves generating a vector of scaled partial fitness sums, where each element is defined by

$$p_i = \frac{\sum_0^{i-1} f_i}{\sum_0^{n-1} f_i} \tag{1}$$

Then a random number between 0 and 1 is generated and the first element larger than this number is selected.

We also implement two crossover operators: *Partially matched crossover* (PMX) [14] and *Greedy Subtour Crossover* (GSX) [15]. These are common operators for permutation encoded genetic algorithm, which require the recombined offspring to be a valid permutation, i.e. a permutation where all alleles in the chromosome are unique.

Finally, three mutation operators are implemented: *Greedy exchange mutation*[16] *, move mutation* and *mutate lead in.*

VI. SEA TRIAL RESULTS

The in-mission ATR and the path planning algorithms were tested outside Horten in Norway in June 2012. Six exercise mines (one cylinder and five truncated cones) where deployed by the RNoN before the tests in a pattern shown in Figure 8. In addition, there are several other minelike objects, mainly cylinders and oil drums, in the area.

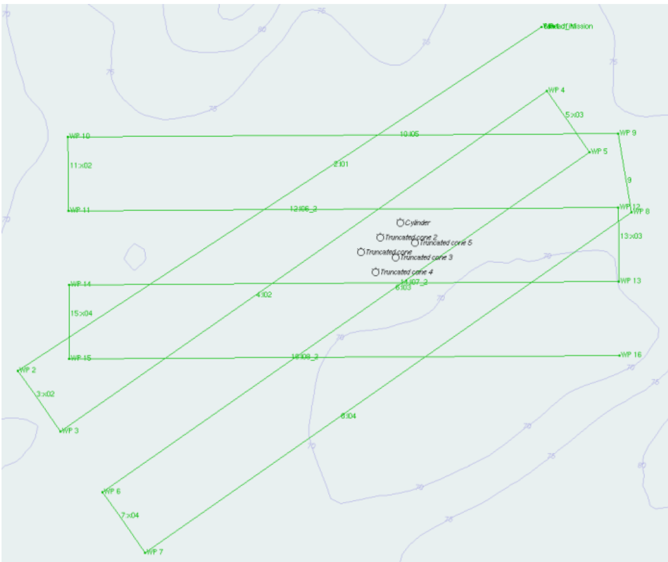

Figure 8: The deployed targets and the resulting survey plan. The distance between the survey lines is 75m.

The algorithms were installed on RNoN's vessel HUGIN 1000 V0, described in Section II. The test consisted of two parts. The first part was a survey where the vehicle moved along two lawnmower patterns at an angle to each other, shown in Figure 8. The lawnmower patterns were planned to give a good coverage of the area with the deployed mines as well as the immediate surroundings. The plan also avoided terrain that was considered unsafe for the low-altitude inspection part of the dive.

The second part of the dive was an inspection part where the vehicle acquired optical imagery the targets. The planning algorithm were configured to ignore targets that were classified as clutter, had a classification confidence less than 20% or where outside the area of interest. Figure 9 and Figure 10 show the detected targets, the area of interest and the resulting mission plan, which was executed at 4-6m altitude. Only four of the six exercise mines were detected by the ATR algorithm, likely due to the low resolution of the sonar.

Between the dives the vehicle surfaced to establish radio contact with the operators, who confirmed that the generated inspection plan was safe for the vehicle to execute.

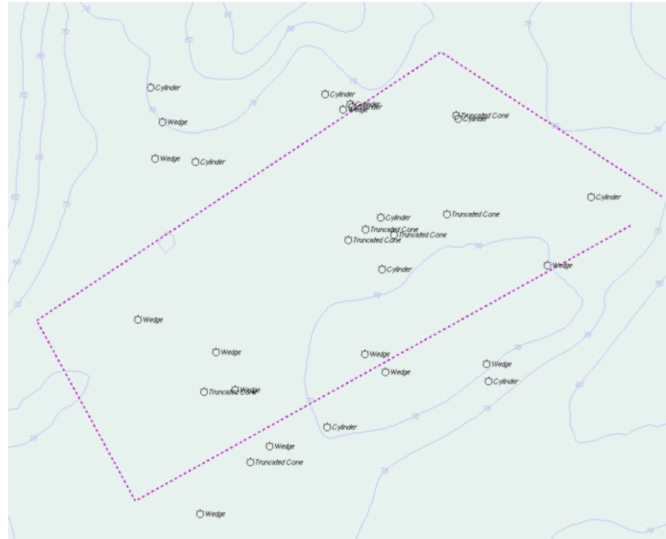

Figure 9: Detected mine-like objects and area of interest

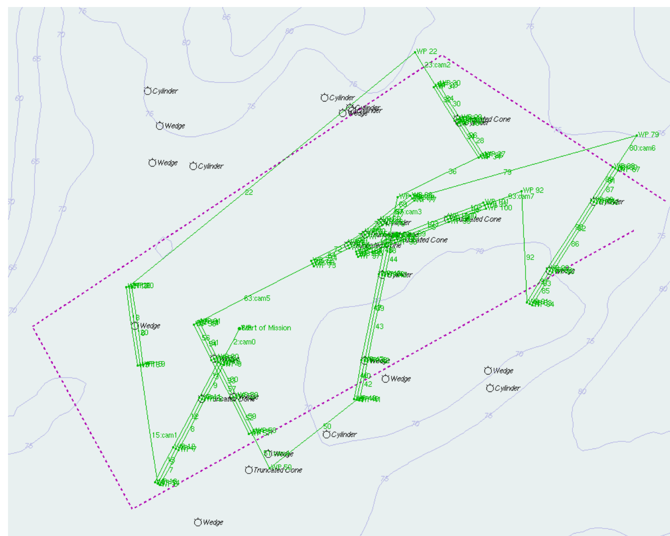

Figure 10: Mission plan for optical object inspection

The vehicle was able to cover all of the detected targets within the area of interest. An example of a target, an oil drum, is shown in Figure 11. As the image shows, the visibility was very low on the day of testing, which is often the case in this area during spring and summer. Still, it was possible to identify all targets as either practice mines or clutter.

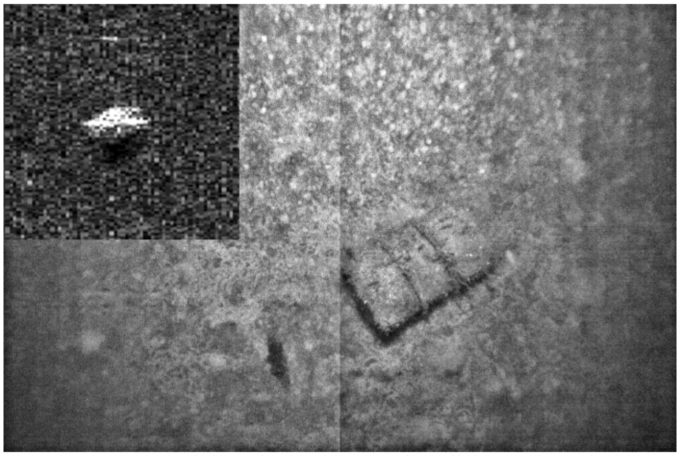

Figure 11: Sonar (top left) and optical image of an oil drum. The optical image was taken at an altitude of 5.5m. Courtesy of the RNoN.

VII. Conclusions and further work

This paper has presented a genetic algorithm for planning inspection missions for optical image gathering using an AUV. The software module takes a set of targets, and filters them to generate a set that fulfills criteria on type, location and classification confidence. The genetic algorithm then generates a mission plan that covers the remaining targets.

The algorithm was demonstrated on a sea trial with the RNoN's HUGIN 1000 V0 AUV. An in-mission ATR module was also demonstrated, and this module provided the target list for the planner. All the detected targets that fulfilled the filtering criteria was covered by the inspection mission plan, and the resulting photos made it possible to identify the objects as either practice mines or clutter.

Future work includes running the in-mission ATR algorithm on a vehicle with real time synthetic aperture sonar, which will increase the accuracy of this algorithm. It also includes extending the planning algorithm to consider the terrain, which will lead to more robust mission plans.

Acknowledgements

The authors would like to thank the Royal Norwegian Navy for deploying the targets, and for letting us use their HUGIN 1000 V0 AUV and the mine control vessel KNM Tyr, the surface vessel of the trials. We would also like to thank the crew of KNM Tyr for their assistance. Finally, we would like to thank the RNoN HUGIN operators Helge Telle and Arild Vangen for all their help and assistance during the trial.

References

[1] P. E. Hagen and N. Storkersen, "Military operations with the HUGIN AUVs," *Advances in Technology for Underwater Vehicles*, 2006.

[2] P. E. Hagen and J. Kristensen, "The HUGIN AUV 'plug and play' payload system," in *Proc. Oceans 2002*, 2002.

[3] P. E. Hagen and N. J. Størkersen, "The HUGIN AUV for Force Protection in the Littorals," in *System Concepts and Integration Symposium*, 2006.

[4] Ø. Midtgaard, P. E. Hagen, and N. Størkersen, "Imaging sensors for AUVs in military operations," in *Proc. NATO RTO SET-169*, 2011.

[5] T. R. Krogstad, M. Wiig, P. Cleophas, and Ø. Midtgaard, "Autonomous object identification in mine countermeasure missions," in *OCEANS 2011*, 2011.

[6] Ø. Midtgaard and P. E. Hagen, "Automatic Target Recognition for the HUGIN Mine Reconnaissance system," in *Proc. International Conference on Detection and Classification of Underwater Tragets*, 2007.

[7] H. Midelfart and Ø. Midtgaard, "Robust Template Matching for Object Classification," in *Proc. 4th Underwater Acoustic Measurements Conference*, 2011.

[8] M. Wiig, "Decisional Autonomy for the HUGIN Autonomous Underwater Vehicle," in *SCI-202 Symposium on Intelligent Uninhabited Vehicle Guidance Systems*, 2009.

[9] K. P. Valavanis, D. Gracanin, M. Matijasevic, R. Kolluru, and G. A. Demetriou, "Control architectures for autonomous underwater vehicles," *Control Systems, IEEE*, vol. 17, pp. 48–64, Aug. 1997.

[10] P. Ridao, J. Yuh, J. Batlle, and K. Sugihara, "On AUV control architecture," *Proceedings. 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2000)*, 2000.

[11] P. Hagen, T. Fossum, and R. Hansen, "HISAS 1030: The next generation mine hunting sonar for AUVs," in *UDT Pacific 2008 Conference Proceedings*, 2008, no. 3191.

[12] S. A. Synnes, R. E. Hansen, and T. O. Sæbø, "Assessment of shallow water performance using interferometric sonar coherence," in *Proc. of underwater Acoustic Measurements*, 2009.

[13] G. Laporte, "The traveling salesman problem: An overview of exact and approximate algorithms," *European Journal of Operational Research*, vol. 59, no. 2, pp. 231–247, 1992.

[14] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Professional, 1989, p. 432.

[15] B. A. Julstrom, “Very greedy crossover in a genetic algorithm for the traveling salesman problem,” in *Proceedings of the 1995 ACM symposium on Applied computing - SAC ’95*, 1995, pp. 324–328.

[16] W. Banzhaf, “The ‘molecular’ traveling salesman,” *Biological Cybernetics*, vol. 64, no. 1, pp. 7–14, Nov. 1990.