# Automatic Course of Action Generation using Soft Data for Maritime Domain Awareness


Alex Plachkov[*]
School of Electrical Engineering & Computer Science, University of Ottawa, Canada
aplac099@uottawa.ca

Rami Abielmona
Research & Engineering, Larus Technologies Corporation, Ottawa, Canada
rami.abielmona@larus.com

Moufid Harb
Research & Engineering, Larus Technologies Corporation, Ottawa, Canada
moufid.harb@larus.com

Rafael Falcon
Research & Engineering, Larus Technologies Corporation, Ottawa, Canada
rafael.falcon@larus.com

Diana Inkpen
School of Electrical Engineering & Computer Science, University of Ottawa, Canada
diana@eecs.uottawa.ca

Voicu Groza
School of Electrical Engineering & Computer Science, University of Ottawa, Canada
groza@eecs.uottawa.ca



## ABSTRACT

Information Fusion (IF) systems have long exploited data provided by hard (physics-based) sensors with the aspiration of making sense of the environment they are monitoring. In recent times, the IF community has recognized the potential of utilizing data generated by people, also known as soft data. In this study, we demonstrate how course of action (CoA) generation, one of the key elements of Level 3 High-Level Information Fusion and a vital component for security and defense decision support systems, can be augmented using soft (human-derived) data for improved mission effectiveness. This conceptualization is validated through an elaborate experiment situated in the maritime world. To the best of the authors' knowledge, this is the first study to apply soft data to automatic CoA generation in the maritime domain.



## Keywords

course of action recommendation; decision support systems; multicriteria decision making; high-level information fusion; soft data


## 1. INTRODUCTION

Maritime Domain Awareness (MDA) can be understood as the situational knowledge of physical and environmental conditions that exist within or influence a maritime region. The intended scope of this awareness includes all behaviors that could, directly or indirectly, affect the security of the region, its economic activity or the local environment [1]. Maritime operators often rely on hard data sources (i.e., structured, quantitative, more objective, usually sensed data) generated by vessel traffic in order to identify suspicious events at sea. However, a wealth of relevant information can be extracted from soft data sources (i.e., unstructured/semi-structured, more subjective, qualitative data, such as textual reports on vessel sightings or marine incidents). As demonstrated in [2], Natural Language Processing (NLP) methods can draw meaningful information that is representative of human intuition, which is often not captured by hard data sources. These pieces of soft information can then supplement the existing hard information in order to provide a more comprehensive situational awareness.

In this paper, we delve into the generation of responses through the use of both hard and soft data and evolutionary multiobjective optimization (EMOO). Having used hard-soft data fusion in the past to augment the Information Fusion capabilities for situational awareness [3], this paper introduces a system that is able to process soft data and extract relevant features, which when fused with hard data, produces responses composed of assets that will carry out the specified mission, such as locating a vessel in distress (VID) or detecting a smuggling situation. The concept of introducing historical soft data into the response generation process should yield a higher mission-specific measure of performance as the embodiment of subject matter expertise (captured by the soft data) into the IF process provides the guidance and direction for the automated search techniques to arrive at near-optimal, yet entirely feasible and interpretable, solutions.

The rest of the paper is structured as follows. Section 2 briefly reviews relevant works. Section 3 presents some background information around maritime-based search patterns. Section 4 unveils the proposed response generation system and Section 5 illustrates its application within MDA, along with the associated experimental results. Section 6 concludes the paper.


[*]Corresponding author





*GECCO'16 Companion, July 20-24, 2016, Denver, CO, USA*





DOI: http://dx.doi.org/10.1145/2908961.2931678

## 2. RELATED WORK

Level 2 (L2) and Level 3 (L3) Fusion are respectively defined as *situation assessment* and *impact assessment* in the Joint Director of Laboratories (JDL)/Data Fusion Information Group (DFIG) models [4][5][6]. L2 Fusion aims to comprehend the current/unfolding situations via assessing relations between entities and their environment, as well as among the entities themselves. Once the situations have been characterized, L3 Fusion is responsible for generating viable Course of Action (CoA) recommendations and estimating their effects on the situations. Both levels have been successfully tasked with carrying out the typical MDA processes (i.e. anomaly detection, trajectory prediction, intent assessment, and threat assessment) [1][7]. The automatic generation of suitable CoAs in the maritime domain using EMOO has been addressed before in [8] and [9] but considering only hard data sources. Soft data has been employed in [2] and [3] at L2 Fusion to extract risk from reported maritime incidents and provide a hard-soft characterization of the maritime risk landscape, respectively. To the best of our knowledge, this is the first study that considers soft data as a main vehicle behind automatic CoA generation for the maritime world.

## 3. BACKGROUND

### 3.1 Maritime Response Basics – Search Patterns

In search and rescue operations, there are four popular types of search patterns with which response assets can be tasked to execute. These are the Track Crawl, Parallel Track Line, Outwards Expanding Square, and Inwards Expanding Square search patterns [10][11].

The *Track Crawl* is assigned to a vessel or an aircraft and tasks them to follow the track of the vessel at large (VaL). This search pattern is used when the VaL will most likely be close to its intended track.

The *Parallel Track Line* is a search pattern providing uniform area coverage. This pattern can be executed by one or more vessels and/or one or more aircraft. The vessels or aircraft follow (in an outward fashion) parallel tracks along the expected drift direction of the VaL. This search pattern is useful when the search area is large.

The *Outwards Expanding Square* is a search pattern which starts at the VaL's last known position (LKP) and expands outward in concentric squares. In addition to vessels, the search pattern can be used by an aircraft. Turns are 90 degrees. It is used when the VaL is thought to be within a small area.

The *Inwards Expanding Square* is a search pattern that is similar to the *Outwards Expanding Square*; however, instead of starting off at the last known VaL position, this location is visited last. It is used in the same situations as the outwards expanding square.

## 4. SOFT-DATA-DRIVEN RESPONSE GENERATION

### 4.1 System Description

The Soft-Data-Driven Response Generation (SDDRG) system is comprised of six modules, namely: (1) the *Anomaly Detection Module* (ADM), which is responsible for determining the confidence levels for different anomaly types (e.g. piracy, smuggling) for each of the assets being monitored; (2) the *Situation Assessment Module* (SAM), which determines the most pressing situation the system will tend to; (3) the *Response Requirements Determination Module* (RRDM), which uses historical incident data to infer the response requirements based on the type of unfolding situation and the manner in which similar, previous situations were dealt with; (4) the *Asset Selection Module* (ASM), which is responsible for selecting which assets will tend to the unfolding situation; (5) the *Asset Path Generation Module* (APGM), which generates tracks for all the assets, based on their designated search areas and assigned search patterns; and lastly, (6) the *Response Enactment Module* (REM), which is responsible for carrying out the response simulation. The architectural blueprint of the system is presented in Figure 1. This research focuses on the development of L3 modules, which are all shaded in grey within the blueprint.

### 4.2 Data Sources

The data sources used as part of this study are all simulated, and fall into three categories: (1) **Weather Data**, a hard or soft data source containing cloud density and precipitation information; (2) **Incident Response Data**, a soft data source containing semi-structured textual data providing detailed information about the responses that are enacted by coastal agencies; and lastly, (3) **Known Pirate Tracks**, a hard data source containing collections of known pirate tracks.

#### 4.2.1 The Role of Soft Response Data

The L3 information extracted from the Incident Response data source is used by the *Response Requirements Determination Module* of the system.

Although there exist some open L3 datasets, such as the International Maritime Organization[1] (IMO) and the Regional Cooperation Agreement on Combating Piracy and Armed Robbery against Ships in Asia[2] (ReCAAP), the response-related information provided by them was deemed to be vague and ultimately inadequate for the purposes of this study. Thus, the L3 information ingested by the system was synthetically generated.

The generated L3 soft data provides the following information: (1) the type of situation the incident is describing (e.g., piracy event, smuggling, vessel in distress); (2) the type of vessel involved in the incident; (3) the assets that took part in the mission, as well as their designated search patterns; and (4) a textual description of the physical landscape (viz., cloud coverage, rain intensity).

An example incident response report can be seen in Section 5.2.

### 4.3 System Module Inputs and Outputs

The inputs and outputs (IOs) of all the different modules of the SDDRG system are presented below:

- **Vessel location** – IO coming from the L2 Situation Assessment module. It specifies the LKP of the vessel which is involved in the most pressing situation that the system has to deal with.

[1] www.imo.org/en/OurWork/Security/PiracyArmedRobbery/Reports
[2] www.recaap.org/AlertsReports/IncidentReports

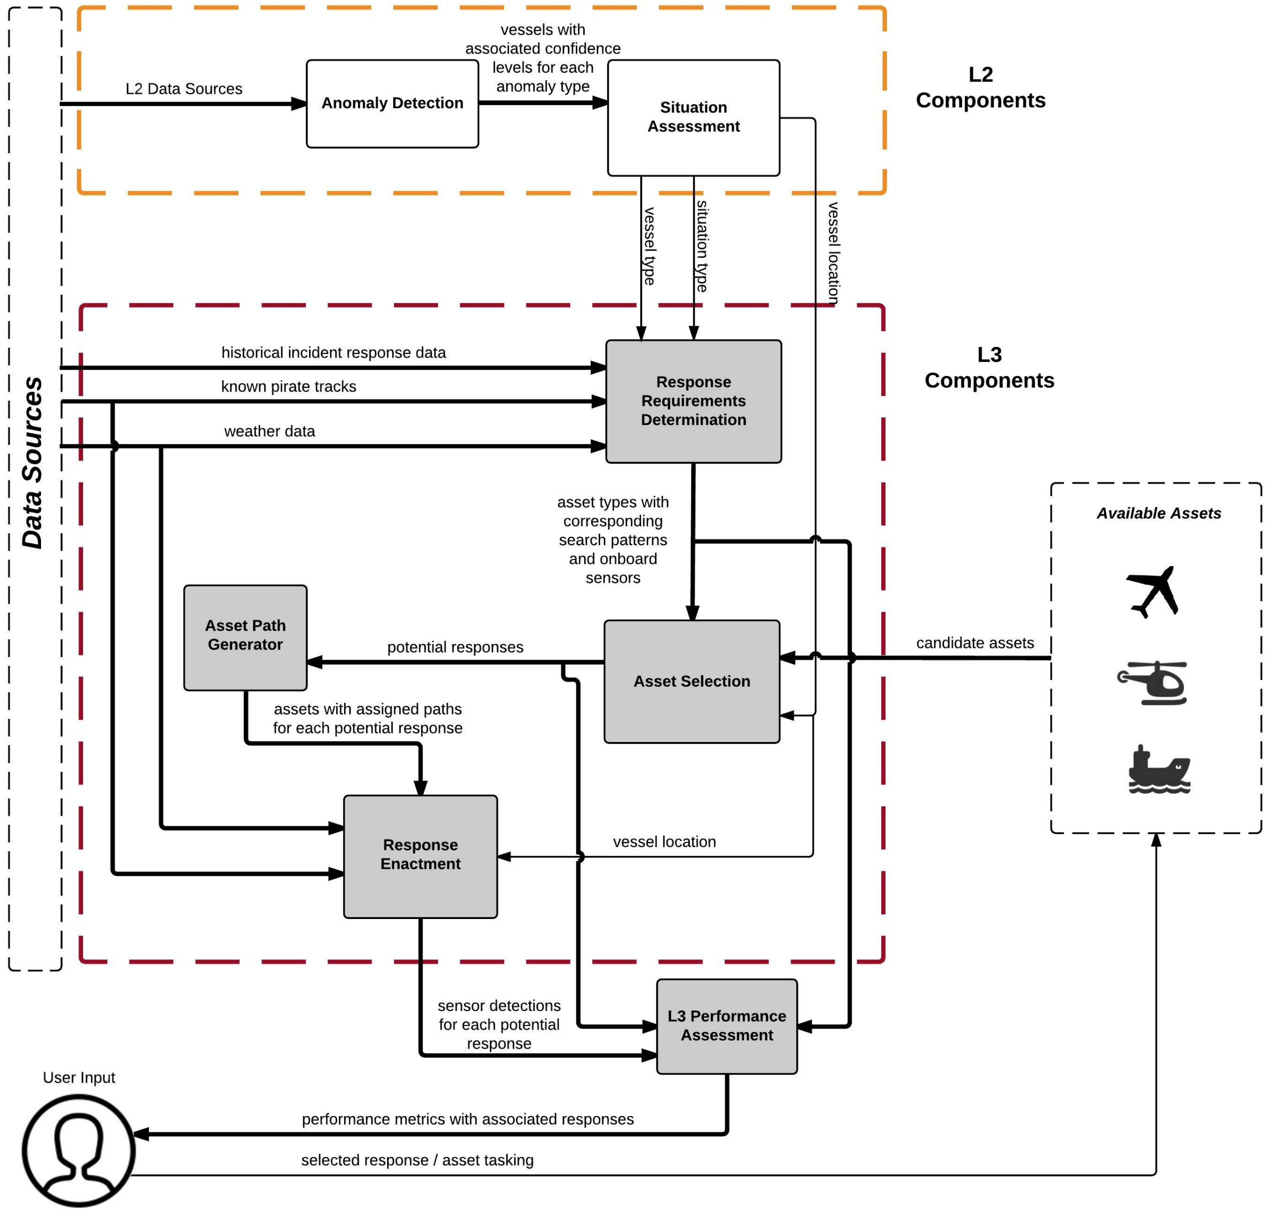


Figure 1: Soft-Data-Driven Response Generation System

- **Vessel type** – IO coming from the L2 Situation Assessment module. It specifies the type of vessel which is involved in the most pressing situation that the system has to deal with.

- **Situation type** – IO coming from the L2 Situation Assessment module. It specifies what the type of the situation which the system shall deal with is.

- **Asset types with corresponding search patterns and onboard sensors** – IO determined based on the historical response data. The most similar scenario is located in the response data, based on: (1) the current situation type, (2) the vessel involved, and (3) the current weather conditions. The RRDM module then takes that scenario, and outputs the asset types, assigned search patterns, and onboard sensors which partook in the historical mission.

- **Candidate response assets** – Assets which are candidates to be used within a response. This includes coast guard assets, as well as opportunistic platforms (i.e. ships or aircraft which are in the vicinity and able to participate).

- **Potential responses** – A collection of potential responses, which can be enacted to tend to the current situation. Each response has a collection of assets which have been selected to participate. Each asset has a designated search area to cover, with a specific search pattern to be executed within the search area. Each asset also has an assigned sensor with a particular quality level. More information on the sensors used as part of this study is presented in Section 5.1.

- **Assets with assigned paths for each potential response** – A search path (a collection of sequential latitude, longitude, and altitude values) is generated for each asset participating in a response.

- **Sensor detections for each potential response** – A statistic provided for each response. It quantifies the likelihood of the assets detecting the VaL.

## 4.4 L3 System Modules

This section unveils the different modules comprising the SDDRG system.

### 4.4.1 Response Requirements Determination Module

The RRDM uses characteristics of the current situation and analyzes the soft, historical incident response data to determine the optimal asset requirements for the mission. For a more detailed description on this module's IOs, please refer to Figure 1 and Section 4.3.

The RRDM contains two submodules – the *Soft Information Extraction* (SIE) submodule and the *Case-Based Reasoning* (CBR) submodule.

The SIE Submodule is responsible for processing the raw textual incident and response reports, and extracting information pertaining to the different categories of interest. More formally, the SIE submodule processes a textual response report, and produces the set, R, containing relevant response information extracted from the report:

$$R = \{situationtype, vesseltype, W, A\} \quad (1)$$

where: $A$ is the superset containing the sets of tuples of assets along with their assigned search patterns, sensor types, and sensor qualities: $A = \{\{asset_1,\ searchpattern_1,\ sensortype_1,\ sensorquality_1\},\ ...,\ \{asset_n,\ searchpattern_n,\ sensortype_n,\ sensorquality_n\}\}$; and $W$ is the set containing the weather-related information encircling the response, or more formally: $W = \{cloudDensity,\ rainDensity\}$. The elements of the set $W$ are all numerical values between 0 and 1, inclusively. The intensity and density values selected to quantify the qualitative descriptions are configurable by the human operator. Table 1 presents one possible configuration.

Table 1: Weather Conditions and Their Associated Intensity and Density Values

| Weather Category | Keywords | Associated Intensity/Density Value |
|---|---|---|
| Rain condition | Extreme rain | 1.00 |
| Rain condition | Heavy rain | 0.75 |
| Rain condition | Moderate rain | 0.50 |
| Rain condition | Light rain | 0.25 |
| Rain condition | No rain, clear skies | 0.00 |
| Cloud category | Extremely dense clouds | 1.00 |
| Cloud category | Dense clouds | 0.75 |
| Cloud category | Moderate clouds | 0.50 |
| Cloud category | Light clouds | 0.25 |
| Cloud category | No clouds, clear skies | 0.00 |

The SIE submodule makes use of the NLP technique called *Named-Entity Recognition* (NER) in order to extract information from the textual incident report and construct the set $R$. There is a lexicon constructed for each of the elements in this set (the *vesseltype* element's lexicon is presented in *Table III: Vessel types and their corresponding category* located in [3]). There is also a lexicon constructed for each of the elements (rain condition and cloud category) found in sets $W$ and $A$, as presented in Table 1.

The RRDM employs a CBR technique, *k-Nearest Neighbor* (kNN) with $k = 1$, in order to locate the most similar scenario in the historical response data. The dissimilarity between two response reports, $X$ and $Y$, is calculated using a modified Euclidean Distance function that takes categorical data into account:

$$D(X,Y) = \sqrt{\delta_S^2(X,Y) + \delta_V^2(X,Y) + \sum_{w \in W}(w_X - w_Y)} \quad (2)$$

where: $W$ remains as previously defined, $\delta_S^2$ calculates the dissimilarity between two situations, and $\delta_V^2$ calculates the dissimilarity between two vessel types. More formally:

$$\delta_S(X,Y) = 1 - \Omega_S(X,Y) \quad (3)$$

$$\delta_V(X,Y) = 1 - \Omega_V(X,Y) \quad (4)$$

where: $\Omega_S$ calculates the similarity between two different situations, and is derived from Table 2; and $\Omega_V$ calculates the similarity between two different vessel types, and is derived from *Table II: Similarity matrix for vessel categories* and *Table III: Vessel types and their corresponding categories*, both of which are located in [3].

Table 2: Similarity matrix for situations

| | Piracy | Smuggling | Vessel in Distress |
|---|---|---|---|
| **Piracy** | 1 | 0.5 | 0 |
| **Smuggling** | 0.5 | 1 | 0 |
| **Vessel in Distress** | 0 | 0 | 1 |

### 4.4.2 Asset Selection Module

The ASM is responsible for selecting which response assets will be carrying out the mission, based on the mission requirements. The module provides a designated search area for each asset. Asset search area designation is optimized with the popular Non-dominated Sorting Genetic Algorithm II (NSGA-II) [12]. This algorithm is used for EMOO to produce a set of spread non-dominated candidate solutions with varying degrees of latency, cost, and response area gap sizes. For a more detailed description on this module's IOs, please refer to Figure 1 and Section 4.3.

The response grid shall be broken down into a square grid. Each cell is a square, which is entirely enclosed within the sweep area of the smallest-sweeping sensor of any of the selected assets.

Responses will be encoded as three-layer chromosomes. Each gene represents an asset, which is available to engage in the response. The first layer will codify an inclusion/exclusion bit on the gene. The second layer encodes the designated response subgrid for that asset (by encoding the row and column indices of the top left corner location of the response grid, as well as the length and width of its designated subgrid). The third layer encodes the type of search pattern that the asset will have to execute. The types of search patterns that an asset can execute are derived from the historical incident response data. Assets have specific sensors (specific types and qualities) which are already mounted on them; thus, the sensor details are not part of the encoding, but are instead used within the *Mission Requirements* objective function. A candidate response encoding is presented in Table 3.

Table 3: Chromosome Encoding of a Candidate Response

| Asset | $A_1$ | $A_2$ | $A_3$ | ... | $A_N$ |
|---|---|---|---|---|---|
| Inclusion | Include | Exclude | Exclude | ... | Include |
| Designated Subgrid | <0,0,2,3> | <1,4,2,2> | <0,7,5,3> | ... | <7,6,2,3> |
| Designated Search Pattern | Parallel Track Line | Track Crawl | Outwards Square | ... | Track Crawl |

In the case that no soft data is available, the designated search pattern defaults to an *adhoc* pattern. This pattern is generated according to Algorithm 1. The *adhoc* search pattern starts off by selecting the upper left corner cell of the response subgrid, and proceeds by selecting a random neighbouring cell. The selection process will give precedence to unexplored neighbouring cells; however, if no unexplored ones exist, it will choose one which has already been visited.

**Algorithm 1: Adhoc Search Pattern**

```
cell = designatedAssetSubgrid.upperLeftCornerCell
for i = 1 to asset.numCellsInSubArea
  pattern.add(cell)
  newCell = getRndUnexploredNeighbour(cell)
  if newCell = null
    newCell = getRndNeighbour(cell)
  end if
  cell = newCell
end for
return pattern
```

The study considers custom evolving operators (mutation and cross-over). The crossover probability which was used was 0.9 and the mutation probability was 0.3. No constraints were imposed on chromosomes, the epsilon value for the Epsilon-Box Dominance Archive was 0.01, and the selection operator used was a binary tournament [12].

The NSGA-II optimizer will make use of the following objective functions:

- Minimize mission time, $MT$:

$$MT = \sum_{a \in A} (T_{a.loc,a.subgrid.start} + T_{a.subgrid.start,a.subgrid.end}) \quad (5)$$

- Minimize mission expenses[3], $ME$:

$$ME = \sum_{a \in A} (E_{a.loc,a.subgrid.start} + E_{a.subgrid.start,a.subgrid.end}) \quad (6)$$

- Minimize unexplored search area, $USA$:

$$USA = 1 - \frac{|\cup_{a \in A} a.gridCells|}{SearchGrid.numGridCells} \quad (7)$$

- Maximize the level to which mission requirements are met by comparing the number and type of required assets to the number and type of ones that have been selected to participate in the response, $MR$:

$$MR = \sum_{a \in A} \beta(\alpha) \quad (8)$$

[3]For the sake of brevity, details on specific asset movement costs are not presented herein.

where $A$ is the set of selected response assets; $E_{a.loc,a.subgrid.start}$ and $T_{a.loc,a.subgrid.start}$ represent respectively the expenses that would accumulate and the time it would take asset $a$ to traverse from its current location to the starting point of its assigned response subgrid; $E_{a.subgrid.start,a.subgrid.end}$ and $T_{a.subgrid.start,a.subgrid.end}$ represent respectively the expenses that would accumulate and the time it would take asset $a$ to traverse from the starting location to the ending location of its subgrid; $a.gridCells$ is the set of grid cells belonging to the search grid that asset $a$ is responsible for visiting, $\cup$ is the union of such sets (for all selected assets); $SearchGrid.numGridCells$ returns the total number of grid cells that the search grid contains; and $\beta(\alpha)$ is a function which returns 1 when the selected asset $a$ is satisfying a selection requirement, and 0 otherwise; satisfying a selection requirement entails: (1) asset $a$ being of a required asset type (e.g. aircraft), and (2) asset $a$ having an onboard sensor of equivalent or higher quality than required quality (the requirement is extracted from the soft data).

If the system is running with historical response reports, all of the above objective functions will be used by the genetic optimizer. However, if the system is running without any historical response reports, the last objective function, $D_A$, is ignored by the genetic optimizer.

#### *4.4.3 Asset Path Generation Module*

The APGM is a minute module whose sole responsibility is to create actual search paths (collections of waypoints) for each asset, based on the collection of grid cells that each response asset has to visit. For a more detailed description on this module's IOs, please refer to Figure 1 and Section 4.3.

#### *4.4.4 Response Enactment Module*

The REM is responsible for carrying out a simulation of each potential response, and providing high-fidelity sensor detection information from the sensors of each of the assets participating in the response. In order to achieve this functionality, the REM's underpinning simulation engine of choice is Larus Technologies' propriety Intelligence, Surveillance, and Reconnaissance (ISR) Tool. The ISR Tool contains models for Modulation Sideband Technology for Absolute Ranging (MSTAR) and Synthetic Aperture Radar (SAR) sensors. The tool accepts weather information, which influences the probability of detection of the different sensors. The qualities of the sensors that are aboard the assets also affect the probability of detection of targets. For a more detailed description on this module's IOs, please refer to Figure 1 and Section 4.3.

### 4.5 L3 Performance Assessment

There are six metrics used to assess the performance of each potential response:

1. **Potential Contact Detections per Response Asset** (PCDRA) – Quantifies the number of potential VaL contacts that are detected by each response asset during the simulations. The higher this value is, the greater the probability of detecting the VaL during the execution of the real-life mission. More formally, the PCDRA is defined as:

$$P = \frac{1}{|A|} \sum_{a \in A} \sum_{v \in V} \lambda_a(v.C) \quad (9)$$

where: $A$ is the set of response assets; $V$ is the set of potential paths that the VaL could have taken; $v.C$ is the set containing the sorted collection of contacts in potential path $v$; and $\lambda_a()$ is a function that takes in a set of potential contacts, and returns the number of those that have been detected by asset $a$.

2. **Response Generation Latency** (RGL) – Represents how long (in milliseconds) it took for the system to generate the set of potential courses of action. More formally ($S$ stands for system):
$$RGL = S_{endtime} - S_{starttime} \tag{10}$$

3. **Unexplored Search Area** (USA) – Quantifies the extent to which the search area (SA) contains gaps. Remains as previously defined in Equation 7.

4. **Mission Time** (MT) – The estimated time of carrying out the mission. Remains as previously defined in Equation 5.

5. **Mission Expenses** (ME) – The estimated expenses of the mission. Remains as previously defined in Equation 6.

6. **Mission Requirements** (MR) – Quantifies the extent to which the MRs were met based on the selected assets to carry out the mission. It is only calculated if soft response data is added to the system. Remains as previously defined in Equation 8.

The human operator is presented with the list of responses, along with their associated performance metrics, and proceeds to select which response, if any, should be carried out, given their training, expertise and intuition.

# 5. CASE STUDY: VESSEL IN DISTRESS

The VID scenarios use a simplistic vessel drift model. The drift course for the VID for each simulation is determined based on a uniform distribution probability value in the range of [0, 360). There are 12 potential responses which have been generated by running the system with the historical incident response data, and 12 potential responses which have been generated by running the system without this soft data. These responses are a subset of the Pareto Archive Set (PAS), and the details of how they have been chosen are unveiled in Section 5.4. There are 25 simulations run for each of the responses.

The overall simulation flow is as follows:

1. Select the next available response to simulate from the collection of viable responses.
2. Get a random number between 0 and 359, inclusively, and generate a VID path (collection of waypoints representing a potential track) that follows this course from the VID starting location to the end of the response area.
3. Run the REM module with the generated VID path, and compute the response's associated performance metrics.
4. Repeat *Steps 2 and 3* 25 times.
5. Go to *Step 1* if there are more available responses to simulate.

## 5.1 Available Response Assets

There are two types of assets that can partake in a response, namely, coast guard assets (CGAs) and opportunistic response assets (ORAs). CGAs can be either grounded (docked) at the coast guard site or traveling somewhere in the surrounding area (not grounded). ORAs are ones which are not owned by the coast guard, but are in the vicinity, and are able to provide assistance. In this scenario, most of the CGAs are grounded at St John's, which is located approximately at latitude of 47.555, and longitude of -52.7067. More specifically, 22 of the CGAs are grounded, and the remaining two CGAs are traversing in the surrounding region. There are also five ORAs in the vicinity, which can partake in the response. Vessel platforms are mounted with MSTAR sensors, whereas air platforms (aircraft, uavs, helicopters, etc.) have SAR sensors. Each of the sensor types has four associated quality levels, as presented in Table 4; the power level values used in this study were experimentally determined. The higher the quality of the sensor mounted on a particular platform, the higher the chance of that platform detecting a VaL is (the PCDRA value would be higher); however, higher quality sensors are mounted on more costly platforms (e.g. a speedboat with a high quality MSTAR sensor costs more to move than a speedboat mounted with a low quality MSTAR, but would have a better chance of detecting a VaL).

Table 4: Sensor Qualities and Associated Power Levels

| Sensor Type | Sensor Quality | Sensor Power Level (Watts) |
|---|---|---|
| MSTAR | Low | 1.200 |
| MSTAR | Medium | 12.00 |
| MSTAR | High | 120.0 |
| MSTAR | Very high | 1200 |
| SAR | Low | 0.008 |
| SAR | Medium | 0.080 |
| SAR | High | 0.800 |
| SAR | Very high | 8.000 |

The grounded CGAs include: (1) **Speedboats:** six equipped with low quality sensors, and two equipped with very high quality sensors; (2) **Tugboats:** five equipped with low quality sensors, one equipped with a medium quality sensor, and two equipped with very high quality sensors; (3) **General cargo boats:** one equipped with a low quality sensor; (4) **Slow UAVs:** one equipped with a low quality sensor, and one equipped with a high quality sensor; (5) **Fast UAVs:** one equipped with a medium quality sensor; and (6) **Helicopters:** one equipped with a low quality sensor, and one equipped with a high quality sensor.

The non-grounded CGAs include: (1) **Fast UAVs:** one equipped with a very high quality sensor, located at <47.0, -50.0>; and (2) **Helicopters:** one equipped with a very high quality sensor, located at <44.5, -46.853>.

The ORAs include five **aircraft** equipped with low quality sensors, located at <50.0, -49.0>, <48.0, -48.0>, <49.0, -47.0>, <47.0, -46.0>, and <48.0, -45.0>.

CGA and ORA locations are not randomized, as the set of experiments to be conducted aim to study the effect including soft data has on the CoA generation for the same (identical) situation.

## 5.2 Scenario Description

This scenario comprises a VID which needs to be located in the worst of weather conditions (dense clouds with extreme rain). A visualization of the response region is presented in Figure 2; the VID's LKP is in the centre of the the response area. The historical incident response report which was deemed by the system to be the most pertinent to the current situation is presented in the following box.

*While underway a bulk carrier seized transmitting AIS information at approximately 100 nautical miles south-east of St John's, Canada. After an unsuccessful attempt to contact the vessel crew, a response mission was launched. Helicopter-3 was assigned a square in search pattern. FastUAV-2 was assigned a parallel track line search pattern. SlowUAV-3 was assigned a square in search pattern. Aircraft-3 was assigned a parallel track line search pattern. Tugboat-3-A was assigned a square out pattern. Tugboat-3-B was assigned a square out pattern. Tugboat-3-C was assigned a square out pattern. Lastly, Speedboat-3 was assigned a parallel track line search pattern. At the time of the incident, moderate rain present with extremely dense clouds were in the response region.*

This incident response report describes a VID situation taking place during bad weather. As a result, the coastal agency decided to dispatch assets with higher quality sensors (all of the assets had *very high* quality sensors, with the exception of FastUAV-2, which only had a *high* quality onboard sensor). This report tells the system to explore solutions comprising of assets with high quality sensors, via the *MR* objective function used within the NSGA-II.

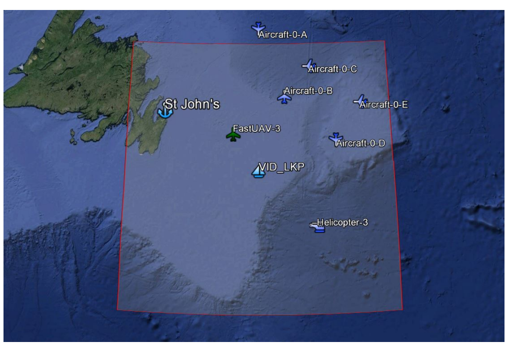


Figure 2: VID Scenario – 22 CGAs grounded at St John's (not visualized), 2 non-grounded CGAs (1 helicopter and 1 fast UAV), and 5 ORAs (aircraft)

## 5.3 Expected Results

The coastal agency in this scenario has assets exhibiting a correlation between sensor qualities and operating costs (the higher the sensor quality on an asset, the more expensive it is to move that asset). Most of the assets in the scenario are grounded, and therefore roughly at the same location. When the system is run without the historical incident response data, it is not aware of how to effectively respond to the current (bad weather) situation; it is expected that the NSGA-II will gravitate towards solutions which include assets with lower quality sensors, as they are located in the same physical location as their more expensive (higher-sensor-quality-equipped) counterparts, yet have the same maximum speeds (i.e. the *overall mission time* objective will be unaffected), yet are cheaper to operate (i.e. *overall mission expense* objective will be better satisfied). However, when the system is run with the historical incident response data, it is expected that it will present solutions, which include assets with higher quality sensors. It is expected that the solutions produced will have a correlation between the degree to which the *MR* objective is met, and the calculated *mission-level probability of detection* (the better the MR objective is satisfied, the higher the probability of detection should be). There is also an expected correlation between the *MR* objective, and the *ME* objective, as higher-quality sensors are on more expensive platforms.

## 5.4 Experimental Results

There were two types of experiments run through the SD-DRG system. The first type, *Experiment Type 1*, was concerned with running the system with the historical incident response data enabled, whereas the second type, *Experiment Type 2*, was concerned with running the system without any historical data. In both cases, the NSGA-II was set to run for 10,000 generations. When the historical data was enabled, *RGL* was calculated to be 127,000 ms for generating the list of viable responses, whereas when it was run without this data, *RGL* was calculated to be 69,000 ms.

### 5.4.1 Experiment Type 1 – With Response Data

This section presents the experimental results gathered by running the system with the historical incident response data enabled. The performance metrics gathered from this experiment are laid out in Table 5.

Table 5: Experimental Results with Historical Incident Response Data

| Response # | ME($) | MR | MT(min) | USA(%) | PCDRA |
|---|---|---|---|---|---|
| 0 | 1285485.98 | 4 | 7716.27653 | 0.0000000000 | 38.29788739 |
| 1 | 1129650.55 | 4 | 8037.72841 | 0.0000000000 | 55.95162705 |
| 2 | 1207408.25 | 4 | 7809.22727 | 0.0000000000 | 38.85187839 |
| 3 | 1190460.33 | 4 | 7841.85429 | 0.0000000000 | 27.95111062 |
| 4 | 3023906.55 | 4 | 5549.42594 | 0.0000000000 | 39.25647836 |
| 5 | 963590.19 | 3 | 7020.96539 | 0.0000000000 | 23.97133441 |
| 6 | 33795.41 | 0 | 298.19482 | 94.8096885800 | 0.00000000 |
| 7 | 72400.29 | 1 | 845.35252 | 98.9619377200 | 54.77510338 |
| 8 | 78438.53 | 1 | 278.57823 | 95.5017301000 | 0.00000000 |
| 9 | 80671.99 | 0 | 258.34687 | 98.9619377200 | 0.00000000 |
| 10 | 206503.21 | 1 | 537.51785 | 93.4256055400 | 1.907509158 |
| 11 | 173736.18 | 2 | 553.06215 | 97.2318339100 | 21.98264251 |

A human operator looking at this solution set, who is primarily concerned with overall response time, may consider viable *Response 11*, which has a low response time, yet a decently high chance the VID (as judged by the response's *PCDRA* value). If this operator, however, is willing to allocate a bit of extra time (for a 52.84% longer mission), he or she could opt in to use viable *Response 7* and drastically bring down the overall mission cost (by about 58%), whilst greatly increasing the overall *PCDRA* level (roughly 2.5 times higher), and hence increasing the probability of detecting the VID.

An interesting observation one can make about the above solution set is that there is a good correlation between the level to which the *MRs* are met and the *PCDRA* (as was previously discussed and expected). Responses which met the *MR* objective with a level of 4 had an average of 2.73 times higher *PCDRA* versus the remainder of the responses (meeting the *MR* with levels ranging from 0 to 3). Responses which met the *MRs* with a level of 4 were also on average 6.82 times more expensive than the remainder of the responses, due the use of higher-quality, but more expensive platforms (another expected trend in the results).

### 5.4.2 *Experiment Type 2 – No Response Data*

This section presents the experimental results gathered by running the system without the historical incident response data. The performance metrics gathered from this experiment are laid out in Table 6.

Table 6: Experimental Results without Historical Incident Response Data

| Response # | ME($) | MR | MT(min) | USA(%) | PCDRA |
|---|---|---|---|---|---|
| 0 | 6209923.14 | N/A | 44954.87537 | 0.0000000000 | 19.72071006 |
| 1 | 6701877.41 | N/A | 35368.15994 | 0.0000000000 | 26.72473344 |
| 2 | 7106243.64 | N/A | 38450.66588 | 0.0000000000 | 2.465088757 |
| 3 | 4842652.20 | N/A | 38929.73464 | 0.0000000000 | 12.79773646 |
| 4 | 148505.91 | N/A | 1172.78313 | 84.7750865052 | 0.00000000 |
| 5 | 281949.99 | N/A | 2698.34101 | 85.1211072664 | 2.873764084 |
| 6 | 291217.74 | N/A | 2707.27653 | 79.9307958478 | 5.950011902 |
| 7 | 336652.87 | N/A | 2239.92386 | 79.9307958478 | 3.775260535 |
| 8 | 375789.14 | N/A | 1293.78110 | 84.0830449827 | 3.686337301 |
| 9 | 674632.95 | N/A | 1612.65306 | 71.9723183391 | 0.00000000 |
| 10 | 522222.34 | N/A | 2003.35055 | 80.2768166090 | 3.813138069 |
| 11 | 1027456.31 | N/A | 2238.13831 | 67.1280276817 | 7.935238095 |

A human operator, who does not have access to historical incident response data, would find themselves with the results presented in Table 6. He/she would conduct a similar analysis, as in Section 5.4.1, and observe trade-offs in the different responses in terms of the different performance metrics. For instance, if the operator wants to maximize the chance of detecting the VID, but is not concerned with the overall time nor cost of the mission, he or she could opt in for *Response 1*. If the operator is more sensitive to cost, but wishes to maintain the probability of detection reasonably high, he or she could select *Response 3* and save approximately 27.74% in mission expenses whilst roughly halving the number of potential VID contacts that each response asset can detect (the *PCDRA*).

### 5.4.3 *Value of Response Data*

The *PCDRA* obtained by running the SDDRG was calculated to be roughly 3.38 times higher when historical incident response data was available. This significant increase represents a tangibly higher chance of a VID detection in the real world, and can be attributed to the fact that the historical data pointed the system towards the use of higher-quality sensors (due to the prevailing weather conditions) at the expense of producing costlier solutions. And indeed, the average mission cost was calculated to be 3 times higher for the responses presented in *Experiment Type 1* vs those found in *Experiment Type 2*.

## 6. CONCLUSIONS

In this study, we have proposed a novel L3 fusion methodology in the presence of both hard and soft data for the purpose of automatic CoA generation. We have demonstrated the tangible benefit of including soft data through an intricate maritime experiment, where the chance of detecting a VaL was significantly increased when soft data was added in the system. Future work will entail extending the system's capabilities through the support of a greater variety of situation types that can be dealt with. To the best of our knowledge, this is the first study that considers soft data as a pivotal element behind automatic CoA generation for the maritime world.

## 7. ADDITIONAL AUTHORS

Additional authors: Emil Petriu (School of Electrical Engineering & Computer Science, University of Ottawa, Canada, email: `petriu@uottawa.ca`).

## 8. REFERENCES


[1] Rami Abielmona. Tackling big data in maritime domain awareness. *Vanguard*, pages 42–43, August-September 2013.

[2] Amir H. Razavi, Diana Inkpen, Rafael Falcon, and Rami Abielmona. Textual Risk Mining for Maritime Situational Awareness. In *2014 IEEE International Inter-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, pages 167–173, San Antonio, TX, USA, March 2014.

[3] Rafael Falcon, Rami Abielmona, Sean Billings, Alex Plachkov, and Hussein Abbass. Risk management with hard-soft data fusion in maritime domain awareness. In *Computational Intelligence for Security and Defense Applications (CISDA), 2014 Seventh IEEE Symposium on*, pages 1–8. IEEE, 2014.

[4] Erik Blasch and Susan Plano. Dfig level 5 (user refinement) issues supporting situational assessment reasoning. In *Information Fusion, 2005 8th International Conference on*, volume 1, pages xxxv–xliii. IEEE, 2005.

[5] Erik Blasch, Ivan Kadar, John Salerno, Mieczyslaw M Kokar, Subrata Das, Gerald M Powell, Daniel D Corkill, and Enrique H Ruspini. Issues and challenges of knowledge representation and reasoning methods in situation assessment (level 2 fusion). In *Defense and Security Symposium*, pages 623510–623510. International Society for Optics and Photonics, 2006.

[6] Erik Blasch, Élio Bossé, and Dale A Lambert. *High-level information fusion management and systems design*. Artech House, 2012.

[7] Rafael Falcon, Rami Abielmona, and Erik Blasch. Behavioral learning of vessel types with fuzzy-rough decision trees. In *Information Fusion (FUSION), 2014 17th International Conference on*, pages 1–8. IEEE, 2014.

[8] Rafael Falcon and Rami Abielmona. A Response-Aware Risk Management Framework for Search-and-Rescue Operations. In *2012 IEEE Congress on Evolutionary Computation (CEC)*, pages 1540–1547, Brisbane, Australia, June 2012.

[9] Rafael Falcon, Rami Abielmona, and Sean Billings. Risk-Driven Intent Assessment and Response Generation in Maritime Surveillance Operations. In *2015 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, Orlando, FL, March 2015.

[10] Australian Maritime Safety Authority. *National Search & Rescue Manual*, 1 edition, 6 2014.

[11] Dianne Timmins. *National Search and Rescue Manual*. Canadian Coast Guard, 6 2000.

[12] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *Evolutionary Computation, IEEE Transactions on*, 6(2):182–197, 2002.