Automatic threat evaluation for border security and surveillance

Bert van den Broek [*], Jos van der Velde, Michiel van den Baar, Loek Nijsten, Rob van Heijster

TNO, Oude Waalsdorperweg 63, 2597 AK The Hague, The Netherlands

ABSTRACT

We present a study of border surveillance systems for automatic threat estimation. The surveillance systems should allow border control operators to be triggered in time so that adequate responses are possible. Examples of threats are smuggling, possibly by using small vessels, cars or drones, and threats caused by unwanted persons (e.g. terrorists) crossing the border. These threats are revealed by indicators which are often not exact and evidence for these indicators incorporates significant amounts of uncertainty. This study is linked to the European Horizon 2020 project ALFA, which focuses on the detection and threat evaluation of low flying objects near the strait of Gibraltar.

Several methods are discussed to fuse the indicators while taking the uncertainty into account, including Fuzzy Reasoning, Bayesian Reasoning, and Dempster-Shafer Theory. In particular the Dempster-Shafer Theory is elaborated since this approach incorporates evaluation of unknown information next to uncertainty. The method is based on belief functions representing the indicators. These functions show a gradual increase or decrease of the suspiciousness depending on input parameters such as object speed, size etc. The fusion methods give two output values for each track: a suspect probability and an uncertainty value. The complete dynamic risk assessment of detected flying objects is evaluated by the automatic system and targets with probabilities exceeding a certain threshold and appropriate uncertainty values are presented to the border control operators.

Keywords: treat evaluation, automation, uncertainty, border security.

1. INTRODUCTION

When a person or object is approaching the border, automatic systems will use data from sensors, network sources and databases to calculate indicators that refer to possible threats. These indicators are not exact because they are based on judgments from operators and on experiences from previous events that may not always have sufficient predictive value for threats in the future, since the ways of smuggling and migrating are changing continuously. Therefore uncertainty and incompleteness concerning the knowledge about the indicators need to play a role in the calculation of the threat level.

A combination of indicators is often necessary since a single indicator is not specific enough to reveal a threat. For example, abnormal sized luggage may indicate smuggling, but since many people carry oversized luggage this indicator alone will not suffice for an accurate threat estimation and other independent indicators are needed for extra information. The need for multiple indicators implies the application of fusion methods that can handle uncertain and incomplete information. The output of a border surveillance system should therefore not only present an estimate of the threat level, but also an assessment of the uncertainty.

In this paper we present the use case of the project ALFA: Advanced Low Flying Aircraft Detection and Tracking. The project exploits a new sensor suite which overcomes the limitations of conventional sensors to detect low flying aircrafts and drones. Near the strait of Gibraltar, smuggling of drugs with boats, small airplanes, and drones from Morocco to Spain and Portugal is an evident border control issue, which requires timely responses. Therefore the project also aims at automatic processing of the sensor data and threat indicators so that threat alerts can be produced in time.

The paper is organized as follows: in section 2 we focus on the requirements for the indicators and in section 3 we discuss several methods for fusing indicators and uncertainty. In section 4 we present the ALFA use case. The supporting

[*] bert.vandenbroek@tno.nl; phone +31 888 66 4075; http://www.tno.nl

evidence for the indicators in the form of belief masses are introduced in section 5. Section 6 is dedicated to threat evaluation. Finally in section 7 we give a summary and conclusions.

2. REQUIREMENTS FOR INDICATORS REVEALING THREATS

For threat analysis the goal is to state a belief whether an approaching object is suspect, non-suspect, or unknown. To that end, we define indicators. An indicator judges to what extend a single attribute, such as speed, object size, and location gives evidence that the flying object is suspect or non-suspect. The verdicts of all indicators should then be fused into a single threat level. Five requirements have guided the design of the threat analysis component:

1) A single indicator verdict should be computed for every flying object. This overall verdict should be a logical consequence of all indicators, whereby the uncertainty of each indicator is taken into account.
2) For each situation, each indicator should state a belief whether the situation is either suspect or non-suspect, and the amount of uncertainty given this situation.
3) Uncertainty is referring to the relative amount of evidence for a situation being suspect versus non-suspect as well as to the to lack of evidence. Concerning indicators uncertainty implies that:
    a. The higher the uncertainty of an indicator, the lower its evidence should be weighted when combined with other indicators.
    b. Conflicting evidence between the various indicators should lead to a higher uncertainty in the total threat level.
4) The configuration of the beliefs of an indicator should be expert-based, because we do not have much historical data. Configuration of an indicator should therefore be straightforward for the expert operator.
5) All indicators should be either dependent on one variable, or in addition also be dependent on the object type (e.g. whether a flight speed is suspicious might be very different for a drone than for a fixed wing aircraft). This makes configuration easier, and is in line with the indicators that are designed by the end users.

Configuration of the threat levels per indicator will be performed by experts and should be straightforward and user-friendly. To simplify and to keep the configuration transparent we assume that each indicator only supports either suspect or non-suspect and not both. This means we have sets of indicators which focus on suspect and sets of indicators which focus on non-suspect. Suspect is primarily associated with abnormal behavior and anomalous events, whereas non-suspect is usually often associated with normal patterns. Since normal behavior is also exhibited by suspects, it is usually easier to find and to asses indicators for threatening events than for normal phenomena (see also [1]). For instance, the indicator in illegal area gives evidence for a flying object being suspect if the distance to the area is low, but if the distance is high, this is no sign of the flying object being non-suspect, it is just a lack of evidence of it being either suspect or non-suspect, so the indicator has a high uncertainty.

Ideally, both types of evidence (suspect and non-suspect) should be equally incorporated in the system. When only using one type of indicators. i.e. that focus on suspect behavior, the system will be biased, and likely to give high suspect values to all encountered situations. Because it is not easy to find useful indicators for non-suspect, except for the indicator that states if a flight is registered and therefore non-suspect, we will counter this bias by adding a default normal indicator, that gives a constant non-suspect value for each situation. This makes sure that the total non-suspect level is higher than the total suspect level for situations in which only a small amount of suspect evidence is available. We therefore consider only a set of indicators for suspect balanced by a general prior for non-suspect to prevent biases.

In addition we assume that indicators are independent. By this simplification, we ignore correlations between indicators. To understand the consequences of this simplification, take for example the correlated indicators for altitude and speed. Suppose that a high speed is suspicious, as is a high altitude. Moreover, suppose that objects with a high speed are more likely to have a high altitude. Then an object with a high altitude might be suspicious regardless of the speed, whereas the speed is very relevant for a low-flying object. Ignoring this correlation will lead to relatively higher suspiciousness when both indicators find a situation suspicious and vice versa. In our model we solve this issue by combining highly correlated attributes into a single parameter and use this for the indicator.

In the following paragraphs we discuss methods which combine evidence for these indicators and which handle uncertain information.

3. FUSION METHODS SUITABLE FOR HANDLING UNCERTAINTY

Given a list of indicators, each having incomplete and uncertain evidence of the threat level of a situation, we want to fuse this evidence into a total threat level. To identify the best data fusion methodology given our requirements, we will now touch upon three methods of data fusion (see [2] for an overview of data fusion methods) which are capable of handling uncertainty:

1) Rule Based Fusion with fuzzy reasoning
2) Probabilistic fusion (Bayes)
3) Evidential belief fusion (Dempster-Shafer)

This selection of methods contains often used methods of data fusion, i.e. the first two of this list. The list is completed by a less used method that seems to fit well with our requirements. After applying each of these data fusion methods to our problem and analyzing their compatibility with our requirements, we will summarize our choices of the data fusion framework in a conclusion.

3.1 Rule based fusion with fuzzy reasoning

Fuzzy logic is an extension of Boolean logic to handle partial truths. To apply fuzzy logic to threat analysis, the evidence of each indicator needs to be mapped to a set. Next, rules should be designed mapping a value of the set to a certain threat level.

Let us for instance consider the altitude of a flying object, which may be mapped to the set {low,medium,high}. To make the reasoning fuzzy, we allow for the possibility of belonging to multiple values within this set, by assigning a degree membership to all values. An object flying at a certain altitude may for instance have a degree membership of 0.3 of having altitude medium and 0.7 of altitude high.

Next, fuzzy inference rules are designed. Here follow some examples:

$$altitude = high \rightarrow suspect = medium$$
$$altitude = medium \rightarrow suspect = low$$

Each rule outputs a fuzzy assignment to the *suspect* set $\{low, medium, high\}$. Multiple approaches are possible to aggregate these assignments into a single total suspect level, for instance taking the maximum or the minimum of all the assignments depending on the fusion approach.

For our requirements, these rules exhibit the problem that it does not output a suspect level value and a separate uncertainty value for each indicator. When the rules state that the suspect level is low, we can interpret it either as representing evidence that the situation is non-suspect, or that this particular indicator has no interesting evidence for suspect in this situation.

To distinguish the evidence of being non-suspect and the uncertainty of the indicator, we might add a *non-suspect* set $\{low, medium, high\}$, and interpret the membership of the suspect and non- suspect set as representing the amount of evidence for the situation being threatening and non-threatening:

$$altitude = high \rightarrow suspect = medium, \quad non\text{-}suspect = low$$
$$altitude = medium \rightarrow suspect = low, \quad non\text{-}suspect = high$$
$$in\ illegal\ area = true \rightarrow suspect = high, \quad non\text{-}suspect = low$$
$$in\ illegal\ area = false \rightarrow suspect = low, \quad non\text{-}suspect = low$$

Here we can distinguish between the altitude indicator stating that a medium altitude is probably non-suspect, with the in illegal area indicator stating that it has no evidence at all about the threat level if the subject is not flying over that area.

Aggregating all assignments of this suspect and non-suspect set is now less straightforward. But, since the framework does not restrict us to implement our own aggregation rules, we could combine the assignments by designing a rule that takes the amount of conflict into account. We could, for instance, take the maximum suspect membership, and lower it with the maximum non-suspect membership. See Figure 1 for the processing flow.

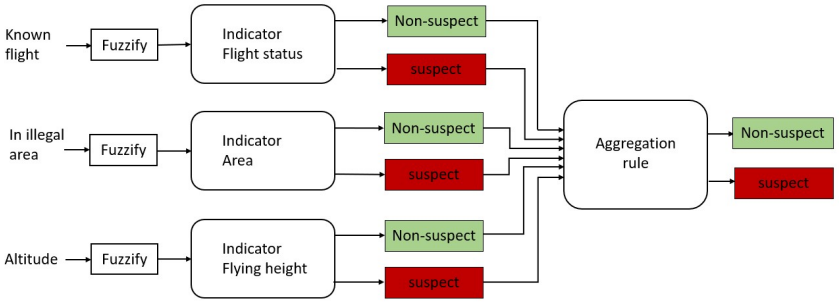

Figure 1: Processing flow for rule based fusion with fuzzy reasoning.

The described procedure would fit our requirements, but the separate suspect and non-suspect sets might seem to be a rather contrived design, and we might question what we obtained from using fuzzy reasoning in this way. We have to create a fuzzy mapping between the evidence and a set of possible values for this indicator, and next map these values to suspect and non-suspect levels. These two steps are intertwined, because the desired suspect levels need to be taken into account in both steps: in the first step the definition of, for example, altitude = high and altitude = low needs to be tuned in order for the second step, the rules, to output the right suspect level. Fuzzy reasoning adds an extra configuration step for our use case, and this extra step does not make the configuration straightforward and user-friendly as was required in section 2.

3.2 Probabilistic fusion

Using probability theory, we can define the probability that the situation is suspect S (or, conversely, the probability that the situation is non-suspect N) given the total evidence E. This probability can be denoted as $P(S|E)$, and can be computed using Bayes' rule:

$$P(S|E) = \frac{P(E|S)P(S)}{P(E)} \propto P(E|S)P(S) = P(E_1, E_2, \dots E_n|S)P(S)$$

Here $E_1, E_2, \dots E_n$ denote the evidence of the n indicators. The non-suspect case $P(N|E)$ is analogous where $P(S|E) + P(N|E) = 1$. To simplify, let us use the assumption of our framework that states that all indicators are independent. We then arrive at the following equations:

$$P(S|E) \propto P(S)\prod_{i=0}^{n} P(E_i|S), \quad P(N|E) \propto P(N)\prod_{i=0}^{n} P(E_i|N)$$

where $P(E_i|S)$ and $P(E_i|N)$ are conditional priors for each indicator. In Figure 2 we present the processing flow for probabilistic fusion.

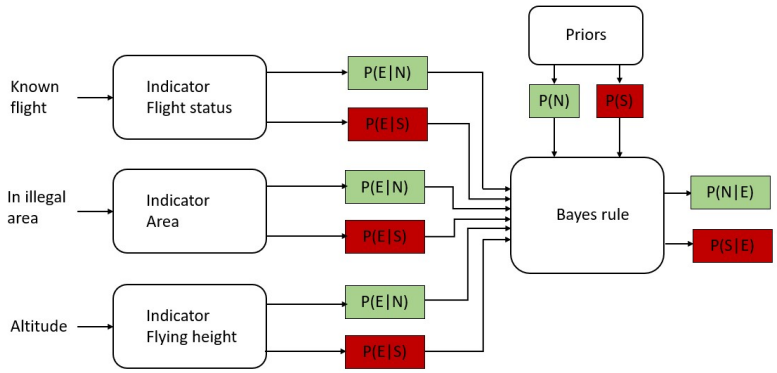

Figure 2: Processing flow for probabilistic fusion.

These conditional priors already give a conflict with our requirements: we cannot distinguish between a lack of evidence and suspect or non-suspect evidence, since the conditional priors for each indicator do not concern lack of the evidence. In other words, the probabilistic framework does not provide a notion of uncertainty. This makes it impossible for an indicator to state "given this situation, I have no clue, so do not take my judgement into account." The inability to handle uncertainty has been the main motivation to establish the theory of belief functions (see next section).

Also it not possible to have indicators which are only associated with suspect since each indicator has to be evaluated for the suspect and non-suspect conditional priors which in practice is hard to do.

3.3 Evidential belief fusion

The theory of belief functions [3], also known as Dempster-Shafer theory, is a generalization of probability theory. It naturally handles uncertainty by assigning a belief mass m not only to the possibility of being suspect $m_E(S)$ and non-suspect $m_E(N)$, but also to the possibility of being either one of them $m_E(SN)$. This belief mass can be interpreted as the proportion of available evidence for either being suspect, non-suspect, and the absence of evidence (i.e. unknown), so that we denote $m_E(SN) = m_E(U)$, and where $m_E(S) + m_E(N) + m_E(U) = 1$.

The framework leaves room for different rules of combination to be applied. Dempster's rule of combination [4] is the most common. Given the belief mass m_1 of the first indicator and m_2 of the second indicator, and a normalization factor C, the joint mass will be:

$$
\begin{aligned}
m_{1,2}(S) &= C\big(m_1(S)m_2(S) + m_1(S)m_2(U) + m_1(U)m_2(S)\big) \\
m_{1,2}(N) &= C\big(m_1(N)m_2(N) + m_1(N)m_2(U) + m_1(U)m_2(N)\big) \\
m_{1,2}(U) &= C\, m_1(U)m_2(U)
\end{aligned}
$$

The rule can be repeated to allow combining more indicators:

$$m_{1,2,3}(\cdot) = m_{(1,2),3}(\cdot)$$

To use these rules, experts first have to design three mass functions for each indicator that return the belief masses of $m(S)$, $m(N)$ and $m(U)$ given the current evidence. Following the requirements we can restrict ourselves to indicators that are independent from each other so that the belief mass is simply a function of the indicators evidence.

Belief theory fulfils most requirements of our use-case in combination with Dempster's rule, but one issue remains: contradictory belief will not lead to a lower certainty in the total threat level: only the amount of uncertainty of each indicator is taken into account while computing the joint mass of SN.

Yager's rule of combination [5] fixes this issue by adding contradictory belief into the joint mass of the unknown. The joint masses $m_{1,2}(S)$ and $m_{1,2}(N)$ are thereby not changed, although a bit simplified since the normalization factor equals to one:

$$
\begin{aligned}
m_{1,2}(S) &= m_1(S)m_2(S) + m_1(S)m_2(U) + m_1(U)m_2(S) \\
m_{1,2}(N) &= m_1(N)m_2(N) + m_1(N)m_2(U) + m_1(U)m_2(N) \\
m_{1,2}(U) &= m_1(U)m_2(U) + m_1(S)m_2(N) + m_1(N)m_2(S)
\end{aligned}
$$

Combining more than two indicators is less straightforward though, since Yager's rule is not associative: $m_{(1,2),3}(S) \neq m_{1,(2,3)}(S)$. This can be overcome by first computing the joint masses $m_{total}(S)$, $m_{total}(N)$ using Dempster's (associative) rule of combination without the normalization factor, and thereafter computing $m_{total}(U) = 1 - m_{total}(S) - m_{total}(N)$. In Figure 3 we present the processing flow for evidential belief fusion.

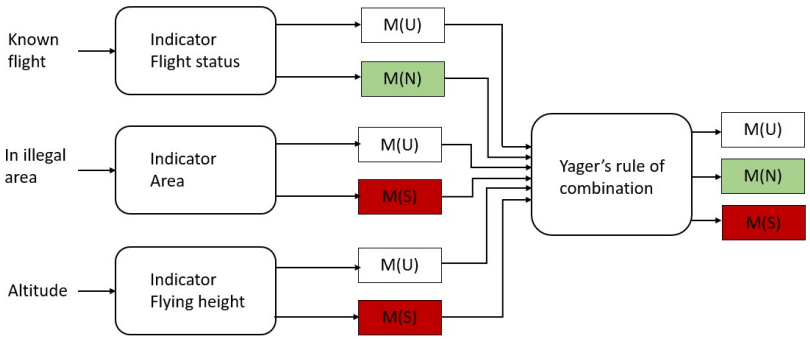

Figure 3: Processing flow for evidential belief fusion.

The main criticism of belief functions [6] focusses on its counterintuitive results in certain cases. These counterintuitive results would arise if we had another possibility besides suspect S and non-suspect N, for instance military action M, giving three possibilities: $\{N, S, M\}$. When combining two indicators, one of which states that it has 0.99 evidence for the situation being S and 0.01 for N, another 0.99 for M and 0.01 for N, the combination will lead to the most likely possibility of the situation being non-suspect, N. This is counterintuitive, because both indicators belief that this is not a likely option. For our use case, this scenario is impossible, because we only have two possibilities: $\{S, N\}$.

3.4 Discussion of methods

Fuzzy reasoning would add an extra configuration step for the end users, to make it to match with our requirements for uncertainty: first, designing the membership functions, and second, designing the fuzzy inference rules. Because this extra step enhances the complexity of the configuration it is not straightforward and user-friendly. We refrain from using fuzzy reasoning.

Probability theory fails to incorporate a notion of unknown, and is therefore not able to distinguish between evidence of a non-suspect situation and the absence of evidence of a suspect situation.

The theory of belief functions (also known as the Dempster-Shafer theory), in combination with Yager's rule for data fusion, forms a good match with our requirements. It naturally couples the lack of information or unknown to uncertainty. The configuration of belief functions and the inference rules can be created straightaway. The theory of belief functions is therefore a suitable candidate for our data fusion framework and is therefore applied in the project.

4. THE ALFA USE CASE AND INFORMATION COLLECTION

The ALFA maritime border surveillance system for the detection of low-flying objects (aircraft, helicopters, and drones) is based on a set of heterogeneous sensor systems covering radar, cameras (EO), and passive radio frequency (RF). The sensor stations that will be set-up can be fixed or mobile. In the sensor station the different representations of information from the radar (tracks), camera (images), and passive RF (plots) will be combined into one information stream (track) per detected object. These tracks will be enriched by a classifier, after which the tracks will be sent to the ALFA control station (ALFA core) for further processing by applying indicators and by fusing them into levels of suspiciousness. The part of the ALFA system that extracts information for the indicators from the sensor data and processes them, consists of four components: sensor fusion, object classification, behavior analysis, and threat analysis.

The sensor fusion component fuses the information of all sensors in the sensor station, and fuses the tracks from multiple sensor stations in the control station. The classification module forms a further processing step, in which the known information is enriched with likelihood values for each of the possible object types. The behavior analysis delivers additional behavioral attributes that will be used in the threat analysis using a complex event processor. These attributes represent basic characteristics of the detected low-flying objects, such as type (aircraft, helicopter, drone), size, and features representing behavior, such as speed and height. Additional data from geographical information systems (providing the relation to valleys, nature reserves, urban areas, vessels, etc.), weather forecast systems and accessible intelligence systems (historical data and air traffic control data) are produced to enhance the analysis.

In the threat analysis component all the attribute information provided by the previous steps will be used to calculate and fuse the indicators into a single level of suspiciousness for each detected low-flying object. Threat indicators can be divided into three types: (1) target attributes that represent the basic characteristics of the detected low-flying objects such as type (aircraft, helicopter, drone), size, and elevation; (2) behavior attributes based on trajectory, maneuvers, and the relation to land or sea objects (valleys, nature reserves, urban areas, vessels); (3) indicators from other information sources such as weather, historic data, and air traffic control. Several approaches for fusion have been considered (see previous section) and the Dempster-Shafer approach is considered here. The advantage of this approach is that it provides us with a straightforward and user-friendly methodology to convey end-user knowledge to the ALFA system by means of belief mass functions, whereby uncertainty and incompleteness are explicitly defined.

The threat analysis output of the control station will be presented to the ALFA mobile sensor-station operators and the border guards and other law enforcement agents in the field, by means of a handheld mobile device. This way, the system will provide decision support, enabling the border guards to take proper actions.

5. BELIEF (MASS) FUNCTIONS

For each indicator a mass-function needs to be constructed by the expert. By creating mass functions evidence for an indicator is directly mapped to beliefs for being suspect, non-suspect or unknown. These functions depend on observed attribute parameters such as speed, height, distance etc. Some mass functions are dependent on multiple attributes. E.g. the indicator for *altitude* depends on the *object type*. In this case each tuple *(altitude, object type)* is mapped to a separate mass. We call these mass functions conditional mass functions. Conditional mass functions should allow for uncertainty and we assume that this uncertainty is expressed in probabilities for different object types, for instance a probability of 70% that it is a helicopter, and 30% that it is a drone. The combined mass function thus becomes a marginal mass function:

$$m_A(S) = \sum_{T \in O} m_{A|T}(S)\ p_T(S)$$

where $m_A(S)$ is the belief mass of being suspect given *altitude* A, O is the set of all *object types*, $m_{A|T}(S)$ is the conditional mass of being suspect given *altitude* A and *object type* T, and $p_T(S)$ is the probability of being suspect given *object type* T.

The configuration of mass functions can be made intuitive by manually changing a two dimensional graph for each indicator. As displayed in Figure 4, the points of each graph can be dragged to alter the mass function.

In a similar way, the mass functions of a dependent indicator can be configured, by creating a mass-function graph for each of the object types.

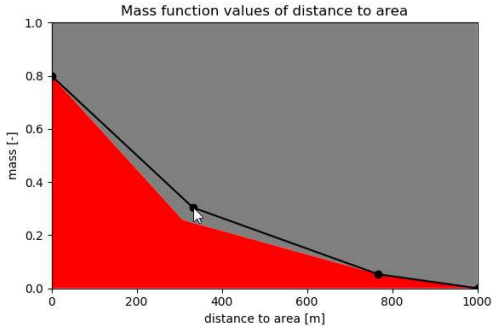

Figure 4: Configuring the mass function of distance to a restricted area. The red area is suspect, while the grey area is uncertain. Dragging the points of the graph changes the mass function.

6. RESULTS FOR THREAT ANALYSYS

To visualize the threat level, we make use of bar plots such as shown in Figure 5. Hereby the information of the suspect and non-suspect evidence are crushed into a single color, ranging from green being completely non-suspect to red being completely suspect, with colors light green, yellow and orange for intermediate cases. The width of the colored bar plot signifies the amount of certainty. The empty part of the box signifies the uncertainty

Figure 5: The threat indicator concept, where the color signifies the ratio of suspect to non-suspect, and the width of the colored bar signifies the amount of certainty.

Figure 6 visualizes six example situations (use cases) including the judgements of all indicators and the total threat level. Note that all indicators are either supporting suspect or non-suspect and never both, implying that there is no uncertainty in the purpose of the indicator. The strength of the support by an indicator (colored part of the bar) signals the uncertainty and is directly related to the lack of information (i.e. unknown, non-colored part of the bar). After fusion of the contributions from the various indicators, the total gives both the ratio between suspect/non-suspect and the unknown: the two parameters for uncertainty in this context. In Figure 6 the individual indicators are visualized on top and the total threat level at the bottom of the diagrams.

Three situations for normality are taken (high, unknown, low), all combined with weak or strong support of the indicators for suspect in case of an detected object. In case the normality is high, the expectation is low that suspicious objects are present and the alert status is low and vice versa.

Indicator	Visualisation	Suspect ratio	Unknown
Default normal		0	25
Area		100	86
Altitude		100	81
Indicator c		100	90
Indicator d		100	71
Indicator e		100	90
Total		33	55

Indicator	Visualisation	Suspect ratio	Unknown
Default normal		0	25
Area		100	21
Altitude		100	81
Indicator c		100	53
Indicator d		100	79
Indicator e		100	90
Total		83	72

Figure 6a: Normality high (alert level low), Case A: weak suspect support (left) , Case B: strong suspect support (right)

Indicator	Visualisation	Suspect ratio	Unknown
Default normal		0	50
Area		100	86
Altitude		100	81
Indicator c		100	90
Indicator d		100	71
Indicator e		100	90
Total		60	50

Indicator	Visualisation	Suspect ratio	Unknown
Default normal		0	50
Area		100	21
Altitude		100	81
Indicator c		100	53
Indicator d		100	79
Indicator e		100	90
Total		94	50

Figure 6b: Normality unknown (alert level medium), Case C: weak suspect support (left) , Case D: strong suspect support (right)

Indicator	Visualisation	Suspect ratio	Unknown
Default normal		0	75
Area		100	86
Altitude		100	81
Indicator c		100	90
Indicator d		100	71
Indicator e		100	90
Total		82	45

Indicator	Visualisation	Suspect ratio	Unknown
Default normal		0	75
Area		100	21
Altitude		100	81
Indicator c		100	53
Indicator d		100	79
Indicator e		100	90
Total		98	28

Figure 6c: Normality low (alert level high), Case E: weak suspect support (left) , Case F: strong suspect support (right)

Figure 6a shows that if the expectation for a normal situation is high, indications for suspect enhances primarily the unknown in the total. In this case the conflict between normal and suspect is dominant, causing the larger unknown.

In case the expectations for a normal situation are low (Figure 6c) we see that indications for suspect are dominant with a relatively low unknown in the total.

In Figure 7 we present the results of the six use cases in an uncertainty diagram with the two uncertainty parameters on the axes. The parameter on the vertical axis represents the amount of knowledge, i.e. complete knowledge minus lack of knowledge (100%-unknown) and the parameter on the horizontal is the ratio suspect/non-suspect. The green symbols represent the total results for low suspect support and red for strong suspect support in cases of an observed (detected) object.

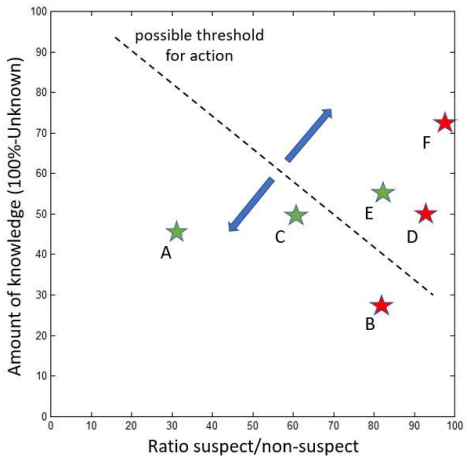

Figure 7: Uncertainty diagram for the six use cases. The arrows indicate the direction of the decision threshold for action by the operator

On basis of this information the operator has to decide whether action should be taken, for example by applying a threshold line in the uncertainty diagram. This threshold depends on the impact of an actual threat for the society (also depending on the socio-political context) and on the available resources for the operation. In the top-right of the diagram the operator usually decides for action, while in the bottom-left he will refrain from action. In between he has to evaluate the impact given the resources and shift the threshold accordingly (arrows in Figure 7). To the bottom-left when the threat is more dangerous (i.e. has more impact) and/or plenty of resources are available and to the top-right when there are limited resources and/or when the threat is considered not so important. The operational benefit of this way of handling uncertainty in automatic threat evaluation has to be proven in practice.

7. SUMMARY AND CONCLUSIONS

We have studied threat analysis for automatic border surveillance, based on the ALFA EU Horizon 2020 project, which aims at improving the detecting of low flying objects by exploiting an improved sensor suite. In the analysis we explicitly used threat indicators which allow uncertain and incomplete information. Several indicators need to be fused in the threat evaluation so that all the information provided by the sensors and other sources are processed into a single level of suspiciousness for each detected low-flying object.

We studied several approaches that can handle uncertainty and the evidential belief fusion method (Dempster-Shafer) was selected for implementation. The advantage of this approach is that we do not need to collect a large body of statistics (which is simply not available) but that it provides us the methodology to convey end-user knowledge to the ALFA system by means of belief functions.

We presented and discussed several use cases, where different levels of suspiciousness in case of an detected object were evaluated, taking into account the uncertainty parameters, lack of information (unknown), and the ratio between suspect/non-suspect.

ACKNOWLEDGEMENTS

This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 700002).

REFERENCES

[1] Broek, B. van den, Smith, A., Breejen, E. den, Voorde, I. van de, "Inference of vessel intent and behaviour for maritime security operations", Proceedings of SPIE Vol. 9248, 2014.

[2] Bahador Khaleghi, Alaa M. Khamis, Fakhri Karray, Saiedeh N. Razavi, "Multisensor data fusion: A review of the state-of-the-art", Information fusion 14.1, pp. 28-44, 2013.

[3] Shafer, G., "A mathematical theory of evidence", Vol. 42. Princeton university press, 1976.

[4] Dempster, Arthur P., "A generalization of Bayesian inference", Journal of the Royal Statistical Society: Series B (Methodological) 30.2, pp. 205-232, 1968.

[5] Yager, R. R., "On the Dempster-Shafer framework and new combination rules", Information Sciences, 41(2), pp. 93-137, https://doi.org/10.1016/0020-0255(87)90007-7, 1987.

[6] Zadeh LA., "On the validity of Dempster's rule of combination of evidence", Electronics Research Laboratory, College of Engineering, University of California, Berkeley, 1979