Credal networks for military identification problems

Alessandro Antonucci
IDSIA
Galleria 2
CH-6928 Manno (Lugano)
Switzerland
alessandro@idsia.ch

Ralph Brühlmann
Armasuisse (W+T)
Feuerwerkerstrasse 39
CH-3600 Thun
Switzerland
ralph.bruehlmann@ar.admin.ch

Alberto Piatti
IDSIA
Galleria 2
CH-6928 Manno (Lugano)
Switzerland
alberto.piatti@idsia.ch

Marco Zaffalon
IDSIA
Galleria 2
CH-6928 Manno (Lugano)
Switzerland
zaffalon@idsia.ch

Abstract

Credal networks are imprecise probabilistic graphical models generalizing Bayesian networks to convex sets of probability mass functions. This makes credal networks particularly suited to capture and model expert knowledge under very general conditions, including states of qualitative and incomplete knowledge. In this paper, we present a credal network for risk evaluation in case of intrusion of civil aircrafts into a no-fly zone. The different factors relevant for this evaluation, together with an independence structure over them, are initially identified. These factors are observed by sensors, whose reliabilities can be affected by variable external factors, and even by the behavior of the intruder. A model of these observation mechanisms, and the necessary fusion scheme for the information returned by the sensors measuring the same factor, are both completely embedded into the structure of the credal network. A pool of experts, facilitated in their task by specific techniques to convert qualitative judgments into imprecise probabilistic assessments, has made possible the quantification of the network. We show the capabilities of the proposed network by means of some preliminary tests referred to simulated scenarios. Overall, we can regard this application as an useful tool to support military experts in their decision, but also as a quite general imprecise-probability paradigm for information fusion.

Keywords. Credal Networks, Information Fusion, Sensor Management, Tracking Systems.

1 Introduction

In the recent times, the establishment of a no-fly zone surveyed by the Air Force around important potential targets has become usual practice, also in neutral states like Switzerland, because of the potential danger of terror threats coming from the sky. In this paper we refer in particular to the Swiss case, where no-fly zones are usually established to protect international conferences, like the World Economic Forum in Davos, or to protect strategic buildings, like for example nuclear power plants and dams.

A no-fly zone for the protection of a single strategic object usually consists of a circular-shaped region with a radius of several kilometers around the target to defend. All the aircrafts flying in this region without the required permissions are considered *intruders*. The no-fly zone is usually divided in two concentric regions: the external no-fly zone is a large region, with many sensors, devoted to the identification of the intruder, while the internal no-fly zone is a small region, containing the object to protect, where fire is eventually released if the intruder is presumed to have bad aims.

But not all the intruders have the same intentions: there are intruders with bad aims (or *renegades*), intruders with provocative aims, and erroneous intruders. Since only renegades represent a danger for the protected object, the recognition of the intruder's aim plays a crucial role in the following decision, which, if it is wrong, is clearly critical. This is the recognition problem we address in this paper.

This problem is complex for many reasons: (i) the risk evaluation usually relies on qualitative expert judgments; (ii) it requires the fusion of the information coming from different sensors, and this information can be incomplete or partially contradictory; (iii) different sensors can have different levels of reliability, and the reliability of each sensor can be affected by exogenous factors, as geographical and meteorological conditions, and also by the behavior of the intruder. A short review of the problem and some detail about these difficulties is reported in Section 2.

Nowadays, the problem is faced by military experts without the support of any mathematical model. The reason is partly the difficulty of finding a suitable mathematical paradigm for this kind of problems.

In this paper, we propose *credal networks* (Section 3)

as a mathematical paradigm for the modeling of military identification problems. Credal networks are imprecise-probability graphical models representing expert knowledge by means of sets of probability mass functions, which are particularly suited for modeling and doing inference with qualitative, incomplete, and also conflicting information.

More specifically, we have developed a credal network for the considered identification problem. This is achieved by a number of sequential steps: determination of the factors relevant for the risk evaluation and identification of a causal structure between them (Section 4.1); quantification of this qualitative structure by imprecise probabilistic assessments (Section 5.1); determination of a qualitative model of the observation mechanism associated to each sensor, together with the necessary *fusion scheme* of the information collected by the different sensors (Section 4.2); quantification of this model by probability intervals (Section 5.2). An analysis of the main features of our imprecise-probability approach to information fusion is indeed reported in Section 6.

The credal network is finally used to evaluate the level of risk, which is simply the probability of the risk factor conditional on the information collected by the sensors in a given scenario. A description of the approximate procedure used to update the network, together with the results of a preliminary test, is reported in Section 7.

Summarizing, we can regard this model as a practical tool to support the military experts in their decisions for this particular problem. But, at the same time, this credal network can be regarded as a prototypical modeling framework for general identification problems requiring information fusion.

2 Military Aspects

This section is focused on the main military aspects of the identification problem. In particular, we explain: (i) what are the possible intentions of the intruder, (ii) what are the factors that are observed to determine the intention of the intruder, (iii) what are the sensors used to determine these factors.

We consider only civil aircrafts; military aircrafts and flying weapons like rockets or cruise missiles are not taken into consideration. For the possible intentions of the intruder, four categories can therefore be considered: *renegade, agent provocateur, erroneous intruder* and *damaged intruder*. A *renegade* is an aircraft that has entered the no-fly zone with the purpose of attacking the protected object using itself as weapon; terrorists belong to this category. The purpose of an *agent provocateur* is the provocation of the protection structure for demonstrative purposes. An agent provocateur usually knows exactly what it is doing and does not want to die. An *erroneous intruder* has no particular purpose: it has entered the no-fly zone by mistake, because of bad preparation of the flight or due to a bad level of training of the pilot. Finally, a *damaged intruder* is an aircraft without bad aims that is incurring an emergency situation due to technical problems. A damaged intruder enters a no-fly zone because it cannot avoid it or because it is in a situation of panic. In our model, the intention of the intruder is modeled as a random variable, called the *risk factor*, whose possible values are the four cases described above.

The intruder is assumed to be observed for a sufficiently long time window, when it is flying in the external no-fly zone. The factors observed during this period to determine its intention can be divided into two categories: factors describing the *flight behavior* and factors describing the *reactions*. For this first category we consider: *height*, *changes in height*, *absolute speed*, *flight path* and *type of aircraft*. These factors are observed in a passive way, without any interaction with the intruder. The factors belonging to the second category are the *transponder (mode 3/A)*, the *reaction to radio communication with the civil Air Traffic Control (ATC)*, the *reaction to radio communication with the Air Defence Direction Center (ADDC)* and, finally, the *reaction to interception*. The common point of these factors is that they require an interaction (code emission, radio communication or visual contact) between the intruder and the civil or military control.

All these factors are regarded as random variables, taking only a finite number of possible values. Variables which are not intrinsically categorical, are discretized. For instance, regarding the height above the ground maintained by the intruder during the observation period, we are not interested in the precise elevation of the aircraft, but on its flight level. According to military practice, the airspace is divided in four levels: VERY LOW (0-150m), LOW (150-3'000), HIGH (3'000-7'000m) and VERY HIGH (above 7'000m).

Many sensors can be used to determine the factors described above. In our application the ADDC works as a centralized decision center receiving all the information collected by the sensors in order to evaluate the intention of the intruder. The network formed by the ADDC and all the sensors is called the *identification architecture*. The sensors in the identification architecture are divided in four main categories:

- *Signals intelligence*. Sensors belonging to this

category detect signals emitted by the intruder. In our application, the only sensor of this type is the secondary surveillance radar (SSR), that detects the Mode 3/A (identification code) and the Mode C (height) emitted by the intruder.

- *Radar intelligence*. Sensors belonging to this category are all the radars. In our application we have three types of radars: 3D radars, detecting the 3D position of the intruder in the airspace; 2D radars, detecting the 2D position but not the height of the intruder; and tracking radars, detecting the 3D position of the intruder but only at low heights and with a limited range.

- *Imagery intelligence*. Sensors belonging to this category record TV or infrared (IR) images of the intruder using cameras.

- *Human intelligence*. Sensors belonging to this category are sensors where the information is elaborated by humans before being transmitted to the ADDC. In our application there are two sensors of this type: ground-based observation units, where humans observe the intruder using optical instruments and communicate their observations to the ADDC, and interceptors, where the pilot observes directly the intruder and communicate the observations to the ADDC.

The identification architecture is a complicated non-homogeneous structure. In fact, not all the sensors are present at the same time in each point of the no-fly zone. The *presence* and the *reliability* of a sensor for observing a given factor of the intruder depend on the position of the intruder (in particular on its height), on the position of the sensors in the architecture and on the meteorological and geographical situation. In Section 4.2 we explain in detail how presence and reliability are modeled by our network.

3 Mathematical Aspects

In this section, we briefly recall the definitions of *credal set* and *credal network* [4], which are the mathematical objects we use to model expert knowledge and fuse the different kinds information in a single coherent framework.

3.1 Credal Sets

We use uppercase letters to denote random variables. Given a random variable X, we denote by Ω_X the possibility space of X, with x a generic element of Ω_X. Denote by $P(X)$ a mass function for X and by $P(x)$ the probability of x.

We denote by $K(X)$ a closed convex set of probability mass functions over X. $K(X)$ is said to be a *credal set* over X. For any $x \in \Omega_X$, the lower probability for x according to the credal set $K(X)$ is $\underline{P}(x) = \min_{P(X)\in K(X)} P(x)$. Similar definitions can be provided for upper probabilities, conditional credal sets, lower and upper expectations. Note that a set of mass functions, its convex hull, and its set of *vertices* (also called *extreme mass functions*) produce the same lower and upper expectations and probabilities.

Conditioning with credal sets is done by elements-wise application of Bayes rule. The posterior credal set is the union of all posterior mass functions. Denote by $K(X|Y = y)$ the set of conditional mass functions $P(X|Y = y)$, for generic variables X and Y. We say that two variables are *strongly independent*, when every vertex in $K(X, Y)$ satisfies stochastic independence of X and Y.

A set of *probability intervals* over Ω_X, say $\mathbb{I}_X = \{\mathbb{I}_x : \mathbb{I}_x = [l_x, u_x], , 0 \leq l_x \leq u_x \leq 1, x \in \Omega_X\}$, can be regarded as a specification of a credal set $K(X) = \{P(X) : P(x) \in \mathbb{I}_x, x \in \Omega_X, \sum_{x\in\Omega_X} P(x) = 1\}$. $\mathbb{I}_X$ is said to *avoid sure loss* if the corresponding credal set is not empty and to be *coherent* (or *reachable*) if $u_{x'} + \sum_{x\in\Omega_X, x\neq x'} l_x \leq 1 \leq l_{x'} + \sum_{x\in\Omega_X, x\neq x'} u_x$, for all $x \in \Omega_X$. $\mathbb{I}_X$ is coherent if and only if the intervals are tight, i.e., for each lower or upper bound in $\mathbb{I}_X$ there is a mass function in the credal set at which the bound is attained [12, 3].

3.2 Credal Networks

Let $\mathbf{X}$ be a vector of random variables and assume a one-to-one correspondence between the elements of $\mathbf{X}$ and the nodes of a *directed acyclic graph* $\mathcal{G}$. Accordingly, in the following we will use *node* and *variable* interchangeably. For each $X \in \mathbf{X}$, Π_X denotes the set of the *parents* of X, i.e., the random variables corresponding to the immediate predecessors of X according to $\mathcal{G}$.

The specification of a *credal network* over $\mathbf{X}$, given the graph $\mathcal{G}$, consists in the assessment of a conditional credal set $K(X_i|\pi_i)$ for each possible value $\pi_i \in \Omega_{\Pi_i}$ of the parents of X_i, for each variable $X_i \in \mathbf{X}$. The graph $\mathcal{G}$ is assumed to code strong dependencies among the variables in $\mathbf{X}$ by the so-called strong Markov condition: every variable is strongly independent of its nondescendant non-parents given its parents. Accordingly, it is therefore possible to regard a credal network as a specification of a credal set $K(\mathbf{X})$ over the joint variable $\mathbf{X}$, with $K(\mathbf{X})$ convex hull of the set of joint mass functions $P(\mathbf{X}) = P(X_1, ..., X_n)$ over the n variables of the net, that factorize according to $P(x_1, \ldots, x_n) = \prod_{i=1}^{n} P(x_i|\pi_i)$. Here π_i

is the assignment to the parents of X_i consistent with $(x_1, \ldots, x_n)$; and the conditional mass functions $P(X_i|\pi_i)$ are chosen in all the possible ways from the respective credal sets. $K(\mathbf{X})$ is called the *strong extension* of the credal network. Observe that the vertices of $K(\mathbf{X})$ are joint mass functions $P(\mathbf{X})$. Each of them can be identified with a Bayesian network [9], which is a precise probabilistic graphical model. In other words, a credal network is equivalent to a set of Bayesian networks.

3.3 Computing with Credal Networks

Credal networks can be naturally regarded as expert systems. We query a credal network to gather probabilistic information about a variable given evidence about some other variables. This task is called *updating* and consists in the computation, with respect to the network strong extension $K(\mathbf{X})$, of $\underline{P}(X|E = e)$ and $\overline{P}(X|E = e)$, where E is the vector of variables of the network in a known state e (the evidence), and X is the node we query. Credal network updating is an NP-hard task [5], for which a number of approximate algorithms have been proposed [8, 2].

4 Qualitative Assessment of the Network

We are now in the position to describe the credal network developed for our application. According to the discussion in the previous section, this task first requires the qualitative identification of the conditional dependencies between the different variables involved in the model, which can be coded by a corresponding directed acyclic graph.

As detailed in Section 2, the variables we consider in our approach are: (i) the *risk factor*, (ii) the nine variables used to assess the intention of the intruder, (iii) the variables representing the observations returned by the sensors, (iv) for each sensor two additional variables representing presence and reliability of the sensor. In the following, we refer to the variables in the categories (i) and (ii) as *core variables*.

4.1 Risk Evaluation

Figure 1 depicts the conditional dependencies between the core variables according to the military and technical considerations of the Expert.[1] As an example, the arcs connecting the nodes *type of aircraft*, *height*, and *risk factor* with the *speed*, correspond to the following Expert's remarks: *there is a strong relation between the height above the ground and the corresponding speed of an aircraft* (technical considerations); *a renegade is expected to fly as fast as possible* (military consideration); *an intruder flying with a light aircraft, because of the limited maximal speed of this type of aircrafts, would necessarily flight very slowly.* The specification of this part of the network has required a considerable amount of military and technical expertise that, due to space limitations, cannot be explained in more detail here.

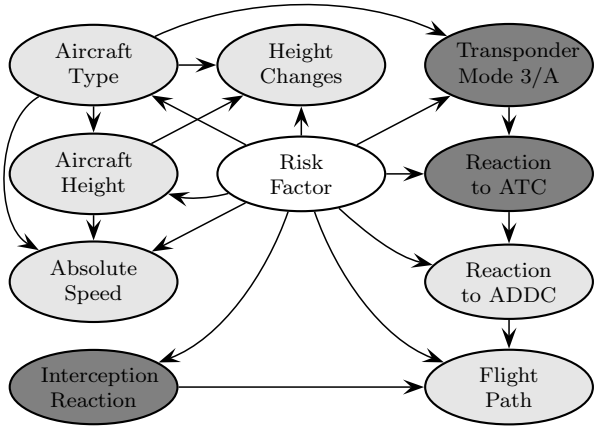

Figure 1: The core of the network. Dark gray nodes are observed by single sensors, while light gray nodes are observed by set of sensors for which an information fusion scheme (see Section 4.2) is required.

4.2 Observation and Fusion Mechanism

We distinguish between *latent variables*, that are assumed to be unobservable, and *manifest variables*, which are actually observed. The *core variables* in Figure 1 are regarded as latent variables that, to be determined, usually require the fusion of information coming from different sensors, with different levels of reliability. Nevertheless, in the case of the identification code emitted by the intruder (*Transponder Mode 3/A*), the *reaction to interception* observed by the pilot, and the reaction to civil air traffic control (ATC) observed by the controllers through SSR, the observation mechanism is immediate; thus we simply identify the latent with the corresponding manifest variable, adding the value MISSING, as possible value of the variable. This value can have particular meanings (eg., a missing Mode 3/A probably means a switched off transponder) and will be also added to the possibility space of the other manifest variables.

Clearly, if the risk factor was the only latent variable, the network in Figure 1 would be the complete network needed to model the risk evaluation. But, because we are dealing with latent variables observed by many sensors, a model of the observation and a fusion

[1] In this paper we briefly call *Expert* a pool of military experts from the Swiss Air Force, we have consulted during the development of the model.

mechanism has to be added to the current structure.

Observation Mechanism We begin by considering observations by single sensors, and then we explain the fusion scheme for several sensors. Consider the following example: suppose that an intruder is flying at low height and is observed by ground-based observation units in order to evaluate its *flight path*. For this evaluation, the intruder should be observed by many units. If our identification architecture is characterized by too a low number of observation units, it is probable that the observation of the flight path would be incomplete or even absent, although the meteorological and geographical conditions are optimal. In this case, the low quality of the observation is due to the scarce presence of the sensor. Suppose now that the architecture is characterized by a very large number of observation units but the weather is characterized by a complete cloud cover with low clouds, then the quality of the observation is very low although the presence of units is optimal. In this case the low quality of the observation is due to the low reliability of the sensor under this meteorological condition. This example motivates our choice to distinguish between *reliability* and *presence* of the sensors in the network.

Figure 2 illustrates, in general, how the evidence provided by a sensor about a latent variable is assessed. The manifest variable depends on the relative *latent variable*, on the *presence* of the sensor and on its *reliability*. Both *reliability* and *presence* are categorical variables with three possible values, HIGH, MEDIUM and LOW for the *reliability*, and PRESENT, PARTIALLY PRESENT and ABSENT for the *presence*.

The *reliability* of a sensor depends on the meteorological and geographical situation and on the height, while the *presence* of a sensor depends only on the identification architecture and on the height of the intruder. The dependence on the latent variable *height* can be explained considering the technical limits of the sensors. There are sensors that are specific of the low and very low heights, like tracking radars and TV or IR cameras. There are other sensors, like the 3D radars of the fixed military radar stations, that are always present at high and very high heights, but are not always present at low and very low heights.

The *meteorological and geographical conditions* do not affect the presence of a sensor, but only its reliability. It is important to point out that these conditions are always observed and we will not display them explicitly as variables in the network, being already considered by the Expert during his quantification of the reliability.

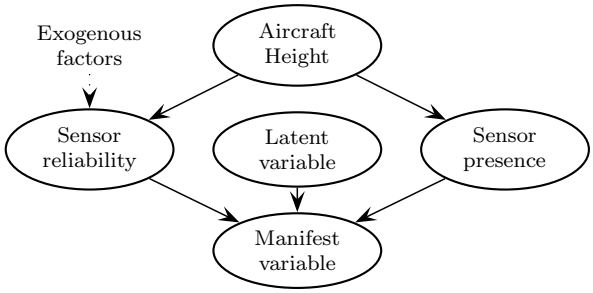

Figure 2: Observation mechanism for a single sensor. The *latent variable* is the variable to be observed by the sensor, while the *manifest variable* is the value returned by the sensor itself.

Sensors Fusion We can finally explain how the information collected by the different observations of a single latent variable returned by different sensors can be fused together. Consider, for example, the determination of the latent variable *type of aircraft* depicted in Figure 3. The *type of aircraft* can be observed by four types of sensors: TV cameras, IR cameras, ground-based observation units and air-based interceptors. For each possible sensor, we model the observation using a structure like the network in Figure 2: there is a node representing the *presence* of the sensor and a node representing the *reliability* of the sensor, while the variable *height* influences all these nodes. This structure permits the fusion of the evidence about the latent variables coming from the different sensors, taking into account the reliability of the different observations in a very natural way and without the need of any external specification of explicit fusion procedures. Similar approaches have already been used for Bayesian networks [6].

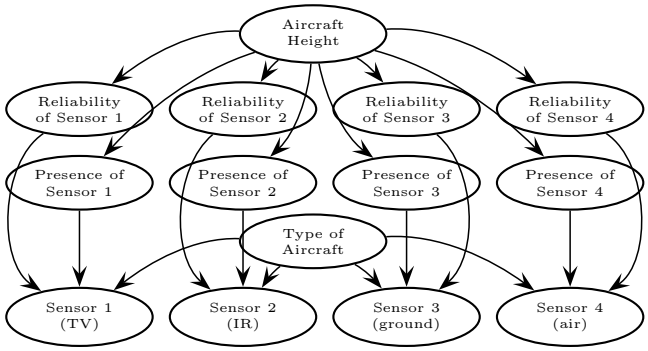

Figure 3: The determination of the latent variable *type of aircraft* by four sensors.

We similarly proceed for all the latent variables requiring the fusion of information from many sensors. This practically means that we add a subnetwork similar to the one reported in Figure 3 to each light gray

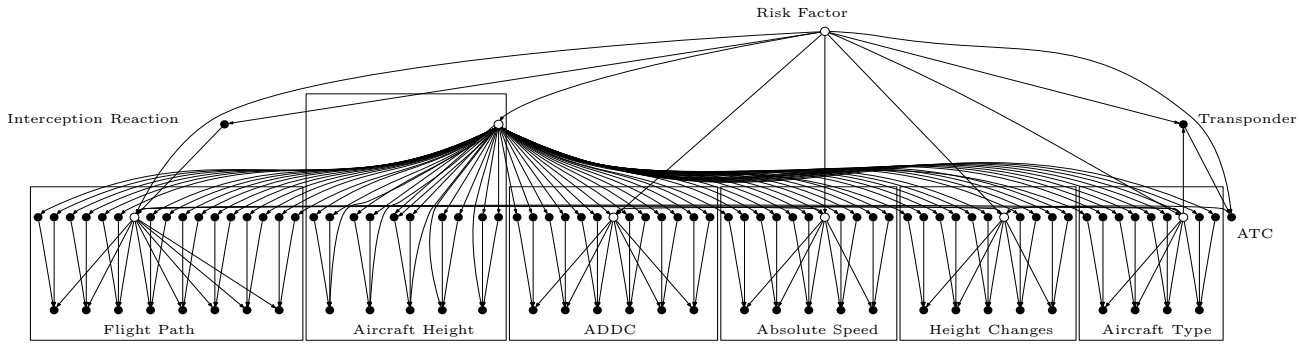

Figure 4: The complete structure of the credal network. Black nodes denote manifest variables, while latent variables correspond to white nodes. Boxes are used to highlight the different subnetworks modeling the observations of the latent variable as in Figure 3.

node of the core network in Figure 1. The resulting directed graph, which is still acyclic, is shown in Figure 4.

5 Quantitative Assessment of the Network

As outlined in Section 3, the specification of a credal network over the variables associated to the directed acyclic graph in Figure 4 requires the specification of a conditional credal set for each variable and each possible configuration of its parents.

For the core variables, these credal sets have been obtained by means of probability intervals explicitly provided by the Expert (Section 5.1), while, regarding observations, presence and reliability, a quantification procedure to automatically transform Expert's qualitative judgments in conditional credal sets specifications has been developed (Section 5.2).

5.1 Quantification of the Network Core

Because of the scarcity of historical cases, the quantification of the conditional credal sets for the core variables in Figure 1 is mainly based upon military and technical considerations. Together with the Expert we have isolated a number of principles, later translated into probability intervals and hence into conditional credal sets. We point the reader to [10] for a detailed description of this quantification task. Here, we cite as an example only some of the principles used to quantify this part of the network: *a renegade is not expected to use balloons or gliders; the light aircraft is the type of aircraft more probable to be used by a terrorist; erroneous intruders are usually light aircrafts and we do not expect a business jet or an airliner to be an erroneous intruder; balloons and gliders are subject to defects due to the meteorological conditions.*

In some situations, the Expert was also able to identify logical constraint among the variables. As an example, the fact that *balloons cannot maintain high levels of height* represents a constraint between the possible values of the variables *type of aircraft* and *height*. These kinds of constraints have been embedded in the structure of the network by means of zero probability assessments.

5.2 Observations, Presence and Reliability

To complete the quantification of our credal network, we should discuss, for each sensor, the quantification of the variables associated to the observation, the reliability and the presence.

We begin by explaining how presence and reliability are specified. Consider the network in Figure 2. The Expert should quantify, for each of the four possible values of the variable *height*, a credal set for the reliability and a credal set for the presence of the sensor. In practice, the Expert is simply required to suggest a value for the presence and a value for the reliability. To assess the value of the presence, he should take into consideration only the structure of the identification architecture; while to assess the value for the reliability level, also the actual meteorological and geographical situation should be considered.

For each specified level of presence or reliability, the Expert should also decide whether or not he is uncertain about this value. His judgments are then translated into coherent probability intervals, from which we can compute the corresponding credal sets reflecting his beliefs. To this purpose, we have defined, together with the Expert, a set of fixed credal sets that are used to model the different combinations of values and uncertainty values. This procedure sub-

stantially simplifies the quantification of the network, while maintaining a large flexibility in the specification of presence and reliability.

Regarding the observations, a conditional credal set for each possible value of the corresponding latent variable and for each possible level of reliability and presence has been assessed. The Expert has answered questions like, *what is the probability (interval) that the ground-based observers have medium reliability in observing the type of aircraft of an intruder that is flying at low height, if the meteorological condition is characterized by dense low clouds and we are in the plateau?*

Clearly, it can be extremely difficult to answer dozens of questions of this kind in a coherent and realistic way. It is much easier to answer questions like the following, *what is the reliability level that you expect from ground-based observers observing the type of aircraft of an intruder that is flying at low height, if the meteorological condition is characterized by dense low clouds and we are in the plateau?* The latter question is much simpler than the former, because one is required to specify something more qualitative than probabilities. This is exactly the type of question that we asked the Expert to quantify the necessary probabilities in our network. In the following we explain, in order, our quantification of presence and reliability of sensors, and the observation mechanism.

Let X be a latent variable, and O the manifest variable corresponding to the observation of X as returned by a given sensor. For each combination of *reliability* and *presence*, we should assess, for each $x \in \Omega_X$ and $o \in \Omega_O$, the bounds $\underline{P}(X = x|O = o)$ and $\overline{P}(X = x|O = o)$.

This quantification step can be simplified by defining a symmetric non-transitive relation of *similarity* among the elements of Ω_X. The *similarities* between the possible values of a latent variable according to a specific sensor can be naturally represented by an undirected graph as in the example of Figure 5. In general, given a latent variable X, for each possible outcome $x \in \Omega_X$, there are outcomes of X that are similar to x and outcomes that are not similar to x.

Having defined, for each latent variable and each corresponding sensor, the similarities between its possible outcomes, we can then divide the possible observations in four categories: (i) observing the correct value of X; (ii) confounding the real value of X with a similar one; (iii) confounding the true value of X with a value that is not similar; (iv) the observation is MISSING. The idea is to quantify, instead of a probability interval for $P(X = x|O = o)$ for each $x \in \Omega_X$ and each $o \in \Omega_O$, only four probability intervals, corresponding to the four categories of observations described above.

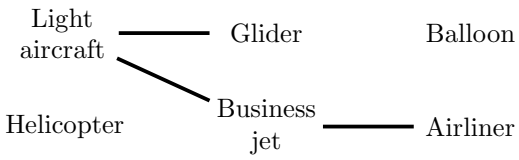

Figure 5: An undirected graph depicting *similarity* relations about the possible values of the variable *types of aircraft* according to the observation of a TV camera. Edges connect similar states. The sensor can mix up a light aircraft with a glider or a business jet, but not with a balloon or a helicopter.

Let us finally explain how the four probability intervals are quantified in our network for each combination of *reliability* and *presence* and for each sensor. The probability interval assigned to the case where the observation is missing depends uniquely on the *presence*. In particular, if the sensor is ABSENT, then the probability of having a MISSING observation is set equal to one and therefore the probability assigned to all the other cases are equal to zero. It follows that we have only seven combinations of *reliability* and *presence* to quantify. To this extent, we use constraints based on the concept of *interval dominance* to characterize the different combinations.[2] In order of accuracy of the observation, the combinations are the following:

1. HIGH, PRESENT: the correct observation dominates (clearly) the similar observations. The probability for not similar observations is zero and is therefore dominated by all the other categories.

2. HIGH,PARTIALLY PRESENT: the correct observation dominates the similar observations and dominates (clearly) the not similar observations. The similar observations dominates the not similar observations.

3. MEDIUM,PRESENT: the correct observation dominates the similar observations and dominates the not similar observations. The similar observations dominates the not similar observations.

4. MEDIUM, PARTIALLY PRESENT: the correct observation does not dominate the similar observations but dominates the not similar observations.

[2]Given a credal set $K(X)$ over a random variable X, and two possible values $x, x' \in \Omega_X$, we say that the x *dominates* x' if $P(X = x') < P(X = x)$ for each $P \in K(X)$. It is easy to show that that interval dominance, i.e., $\overline{P}(X = x') < \underline{P}(X = x)$, is a sufficient condition for dominance.

5. LOW,PRESENT: no dominance at all.

6. LOW,PARTIALLY PRESENT: no dominance at all.

7. ABSENT: the probability of a MISSING observation is equal to one, this value dominates all the other values.

6 Information Fusion by Imprecise Probabilities

The procedure described in Sections 4.2 and 5.2 to fuse the observations gathered by the sensors, can be regarded as a possible imprecise-probability approach to the general *information fusion* problem. In this section, we take a short detour from the military aspects to illustrate some key features of such an approach by a simple example.

Let us first formulate the general problem. Given a latent variable X, and the manifest variables $O_1, \ldots, O_n$ corresponding to the observations of X returned by n sensors, we want to update our beliefs about X, given the values $o_1, \ldots, o_n$ returned by the sensors.

The most common approach to this problem is to assess a (precise) probabilistic model over these variables and compute the conditional mass function $P(X|o_1, \ldots, o_n)$. That may be suited to model situations of *consensus* among the different sensors. The precise models tend to assign higher probabilities to the values of X returned by the majority of the sensors, which may be a suitable mathematical description of these scenarios.

The problem is more complex in case of *disagreement* among the different sensors. In these situations, precise models assign similar posterior probabilities to the different values of X. But a flat posterior probability mass function models *indifference*, while sensors disagreement seems to reflect instead a condition of *ignorance* about X.

Imprecise-probability models are more suited for these situations. Posterior ignorance about X can be represented by the impossibility of a precise specification of the conditional mass function $P(X|o_1, \ldots, o_n)$. The more disagreement we observe among the sensors, the wider we expect the posterior intervals to be, for the different values of X.

The case where the size of the posterior probability intervals results to be increased by conditioning is known in literature as *dilation* [11], and is relatively common with coherent imprecise probabilities.

The following simple example, despite its simplicity, is sufficient to outline how these particular features are obtained by our approach.

Example 1 *Consider a credal network over a latent variable X, and two manifest variables O_1 and O_2 denoting the observations of X returned by two identical sensors. Assume to be given the strong independencies coded by the graph in Figure 6. Let all the variables be Boolean. Assume $P(X)$ to be uniform and both $P(O_i = \mathrm{T}|X = \mathrm{T})$ and $P(O_i = \mathrm{F}|X = \mathrm{F})$ to take values in the interval $[1-\epsilon, 1]$, for each i=1,2, where $\epsilon > \frac{1}{2}$ models a (small) error in the observation mechanism. Since the network in Figure 6 can be regarded as a* naive credal classifier *[13], where the latent variable X plays the role of the class node and the observations correspond to the class attributes, we can exploit the algorithm presented in [13, Section 3.1] to compute the following posterior interval:*

$$P(X = \mathrm{T}|O_1 = \mathrm{T}, O_2 = \mathrm{T}) \in [\frac{(1-\epsilon)^2}{1-2\epsilon(1-\epsilon)}, 1].$$

It follows that, in case of consensus between the two sensors, the corresponding probability for the latent variable increases, given that the lower extreme is larger than $\frac{1}{2}$. In the case of disagreement, instead, we obtain that $P(X = \mathrm{T}|O_1 = \mathrm{F}, O_2 = \mathrm{T}) \in [0, 1]$, which means that our ignorance about X dilates, leading to a completely uninformative posterior interval.

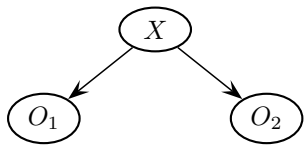

Figure 6: The credal network for Example 1.

Remarkably, assuming fixed levels of height, reliability and presence, Figure 3 reproduces the same structure of the prototypical example in Figure 6, with four sensors instead of two. The same holds for any sub-network modeling the relations between a latent variables and the relative manifest variables in our network.

7 Algorithmic Issues and Simulations

The discussion in Section 4 and Section 5 led us to the specification of a credal network, associated to the graph in Figure 4, over the whole set of random variables we consider, i.e., core variables, observations collected by the different sensors, reliability and presence levels.

At this point, we can evaluate the risk associated to an intrusion, by simply updating the probabilities for the four possible values of the *risk factor*, conditional on the values of the observations returned by the sensors

and on the levels of reliability and presence observed by the Expert.

As a preliminary test of the model, we have considered a simulated scenario of a single object in the Swiss Alps, like for example a dam, surveyed by an identification architecture that is characterized by the absence of interceptors and by a relatively good coverage of all the other sensors. We assumed as meteorological conditions discontinuous low clouds and daylight. The simulated scenario reproduces a situation where an agent provocateur is flying very low with a helicopter and without emitting any identification code. The decision maker is assumed to have uniform prior beliefs about the four classes of risk.

The size of the network suggests the opportunity of an approximate approach to this updating problem. In our approach, we have first reformulated our model as a *locally specified* credal network, according to the procedure developed in [1]. Then, we have transformed each non-binary variable of the credal network into a set of binary variables, according to the *binarization* algorithm, reported in [2]. The resulting credal net has been finally updated by the *loopy* version of the 2U algorithm (L2U) [7]. The overall procedure, which can be proved to be approximate only because of the L2U algorithm, can be implemented in polynomial time. In our case, the credal network has been updated in few seconds on a 2.8 GHz Pentium 4 machine, and convergence of L2U has been observed after seven iterations.

Figure 7.a depicts the posterior probability intervals for this simulated scenario. The upper probability for the outcome *renegade* is zero, and we can therefore exclude a terrorist attack. Similarly, the lower probability for the outcomes *agent provocateur* and *damaged intruder* are strictly greater than the upper probability for the state *erroneous*, and we can reject also this latter value because of interval dominance. Both these results are reasonable estimates for this simulated scenario.

Remarkably, the indecision between *agent provocateur* and *damaged intruder* disappears as we assume higher levels of reliability and presence for the sensors devoted to the observation of the *height*. The results, reported in Figure 7.b, state that the intruder is an *agent provocateur*, as we have assumed in the design of this simulation.

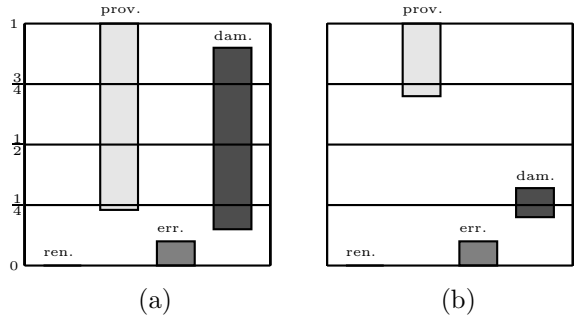

Figure 7: Posterior probability intervals for the risk factor, corresponding to a simulated scenario reproducing a helicopter entering the no-fly zone because of technical difficulties. The histogram bounds denote lower and upper probabilities. The sensors observing the *aircraft height* are assumed more reliable in (b) than in (a).

8 Conclusions and Future Work

A model for determining the risk of intrusion of a civil aircraft into no-fly zone has been presented. The model embeds in a single coherent mathematical framework human expertise expressed by imprecise-probability assessments, and a structure reproducing complex observation mechanisms and corresponding information fusion schemes.

The risk evaluation corresponds to the updating of the probabilities for the risk factor conditional on the observations of the sensors and the estimated levels of presence and reliability. Preliminary tests considered for a simulated scenario are consistent with the judgments of an expert domain for the same situation.

As future work we intend to test the model for other historical cases and simulated scenarios. The approximate updating procedure considered in the present work, as well as other algorithmic approaches will be considered, in order to determine the most suited for this specific problem.

In any case, it seems already possible to offer a practical support to the military experts in their evaluations. They can use the network to decide the risk level corresponding to a real scenario, but it is also possible to simulate situations and verify the effectiveness of the different sensors in order to design an optimal identification architecture.

Finally, we regard our approach to the fusion of the information collected by the different sensors as a sound and flexible approach to this kind of problems, able to work also in situations of contrasting observations between the sensors.

Acknowledgments

This research was partially supported by the Swiss NSF grants 200020-109295/1 and 200021-

113820/1. The Java implementation of the L2U algorithm included in the software tool *2UBayes* (http://www.pmr.poli.usp.br/ltd/) has been used to update the (binary) credal networks. The software tool *lrs* (http://cgm.cs.mcgill.ca/∼avis/C/lrs.htm) has been used to compute the extreme mass function of the conditional credal sets corresponding to the probability intervals provided by the Expert. The authors of this public software tool are gratefully acknowledged.

References

[1] A. Antonucci and M. Zaffalon. Locally specified credal networks. In *Proceedings of the third European Workshop on Probabilistic Graphical Models (PGM-2006)*, pages 25–34, Prague, 2006.

[2] A. Antonucci, M. Zaffalon, J.S. Ide, and F.G. Cozman. Binarization algorithms for approximate updating in credal nets. In IOS Press, editor, *STAIRS'06: Proceedings of the third European Starting AI Researcher Symposium*, pages 120–131, Amsterdam, 2006.

[3] L. Campos, J. Huete, and S. Moral. Probability intervals: a tool for uncertain reasoning. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 2(2):167–196, 1994.

[4] F. G. Cozman. Credal networks. *Artificial Intelligence*, 120:199–233, 2000.

[5] C. P. de Campos and F. G. Cozman. The inferential complexity of Bayesian and credal networks. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1313–1318, Edinburgh, 2005.

[6] E. Demircioglu and L. Osadciw. A Bayesian network sensors manager for heterogeneous radar suites. In *IEEE Radar Conference*, Verona, NY, 2006.

[7] J. S. Ide and F. G. Cozman. IPE and L2U: Approximate algorithms for credal networks. In *Proceedings of the Second Starting AI Researcher Symposium*, pages 118–127, Amsterdam, 2004. IOS Press.

[8] J. S. Ide and F.G. Cozman. Approximate inference in credal networks by variational mean field methods. In F. G. Cozman, R. Nau, and T. Seidenfeld, editors, *ISIPTA '05: Proceedings of the Fourth International Symposium on Imprecise Probabilities and Their Applications*, pages 203–212. SIPTA, 2005.

[9] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, 1988.

[10] A. Piatti, M. Zaffalon, and A. Antonucci. Auswahl und provisorische spezifizierung der kredalen struktur. Technical report, Armasuisse, 2006.

[11] T. Seidenfeld and L. Wasserman. Dilation for sets of probability. *Annals of Statistics*, 21(3):1139–1154, 1993.

[12] B. Tessem. Interval probability propagation. *International Journal of Approximate Reasoning*, 7(3):95–120, 1992.

[13] M. Zaffalon. The naive credal classifier. *Journal of Statistical Planning and Inference*, 105(1):5–21, 2002.