Rationale Visualization for Decision Support

Roeland Scheepens, Steffen Michels, Huub van de Wetering, Jarke J. van Wijk

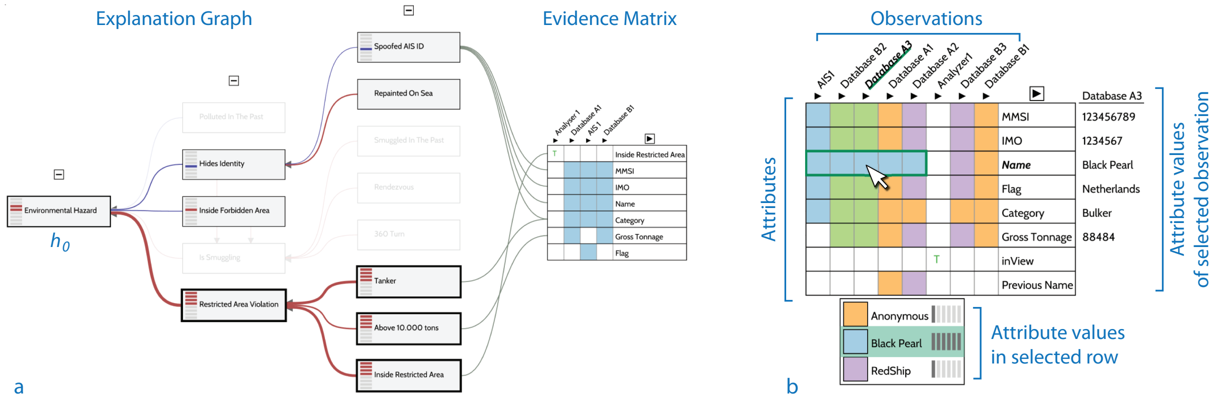

Fig. 1. (a) A visualization to explain why a reasoning engine has concluded that a vessel poses an environmental hazard: The vessel is in an area where no tankers with a weight above 10,000 tons are allowed to be. This is supported by a set of non-conflicting observations in an evidence matrix. (b) An evidence matrix where each column is an observation and each row is an attribute. The cells are colored such that the user can easily see on which attributes the observations agree and on which they conflict. In this case most observations agree on the name of the vessel, but conflict on other attribute values.

Abstract—We present a method to visually explain the rationale of a reasoning engine that raises an alarm if a certain situation is reached. As a case study we look at the maritime safety and security domain. Based on evidence and a reasoning structure, the engine concludes with a certain probability that, e.g., the vessel is an environmental hazard. This engine is part of a larger safety and security system manned by an operator who makes decisions based on the output of the reasoning engine. To support decision making we visualize the rationale, an abstraction of the reasoning structure that allows users to understand why the conclusion has been reached, and display the evidence in a color-coded matrix that easily reveals if and where observations contradict.

Index Terms—Reasoning, Rationale Visualization, Decision Support

1 Introduction

In a decision support system in which objects of interest (OOI) are monitored, a reasoning engine may reason on these objects using information gathered from multiple, heterogeneous, and possibly unreliable sources. The system will raise an alarm whenever the reasoning engine decides certain conditions have been met, i.e., the probability that some task-based hypothesis is true is above a set threshold. As data becomes increasingly available and complex, the need for such automated methods, and especially, automated reasoning rises. Automated reasoning methods, however, are often monolithic black boxes, where data goes in and a hypothesis with a probability for its validity, comes out without an end-user understanding the rationale behind the reasoning. Especially where these conclusions are required for high-cost decision-making, this poses a problem. Trust in the system, an understanding of the situation, and also acceptance of the results are essential to make such decisions.Therefore, we propose to visualize the rationale of a reasoning engine. The end user can then use our visualization to understand why the alarm was raised and confidently take appropriate action.

Our reasoning engine is based on a *first-order probabilistic logic* model, which represents relations between attributes of objects in the domain of interest as joint probability distributions. The model can be used to compute probabilities of arbitrary statements given some evidence. Such a reasoning approach is similar to widely used *Bayesian Networks*,but more general due to the first-order nature of the used language. A more detailed description of this model as used in the maritime safety and security domain is given by Michels *et al.* [1]. To explain our method, we use a case study from that domain, where the OOIs are vessels.

1.1 Data

The actual reasoning structure is a complicated structure of interdependent hypotheses and too complex for an operator to understand. Therefore, we extract an abstraction to make it comprehensible. This abstraction is a directed acyclic graph (DAG), which we call the *explanation graph*. We assume that working with an abstraction is okay, since the operator is expected to be a domain expert. Each node h_i in the graph represents a hypothesis which can be supported by evidence and/or child hypotheses. Each hypothesis is formulated to either represent a situation that requires attention or to support another hypothesis that does. The nodes are connected by directed edges e_{ij} from node h_i (child) to node h_j (parent), where the hypothesis of node h_i is understood to support that of node h_j. The explanation graph has a single root node h_0, which is the main hypothesis, e.g., the vessel is involved in smuggling operations, or as in the example in Figure 1, the vessel poses an environmental hazard.

- *Roeland Scheepens, Huub van de Wetering, and Jarke J. van Wijk are with Department of Mathematics and Computer Science, Eindhoven University of Technology, The Netherlands. E-mail: {r.j.scheepens, h.v.d.wetering, j.j.v.wijk}@tue.nl*
- *Steffen Michels is with Institute for Computing and Information Sciences, Radboud University Nijmegen, The Netherlands. E-mail: s.michels@science.ru.nl.*

Manuscript received 31 March 2013; accepted 1 August 2013; posted online 13 October 2013; mailed on 4 October 2013.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

For each node h_i, the probability p_i^e given the evidence currently available is computed. Along with that the prior probability p_i^{prior} is given, which is the probability of the statement without any evidence. Furthermore, each node has a short descriptive label. Each edge e_{ij} has a dynamic weight $w_{ij} \in [-1,1]$ that signifies the influence, based on the dependencies, a hypothesis has on its parent hypothesis.

The evidence is a set of observations. An observation is a tuple consisting of a number of attributes, which are used to support hypotheses. Multiple observations can contain the same attribute, but their actual values may differ or be absent. Attributes may have completely different domains, e.g. from strings to categorical values to continuous numerical values, and are considered independent. Each attribute value has a probability that it is the actual value.

2 Visualization

We aim to visualize the rationale by showing the explanation graph, the relations between the hypotheses of the graph and the evidence, and the relations between observations in the evidence.

We visualize the explanation graph by drawing boxes for nodes and curved lines with arrowheads for the directed edges. To layout the graph, we use a variation of the Sugiyama layout algorithm [2]. In most instantiations of the explanation graph, i.e., an explanation of the main hypothesis for a given object of interest, not all child hypotheses are required to explain it. To make the graph easier to understand, we show only the sub hypotheses that are required to explain the main hypothesis. We say a hypothesis h_k is not required to explain the main hypothesis h_0 if all paths from h_k to h_0 contain at least one edge e_{ij} with $|w_{ij}| < \varepsilon$, where ε is some threshold. To preserve the mental map of the user, we use a fixed layout for the whole graph and for readability fade out irrelevant nodes to the background–see Figure 1a.

We assume that an operator has domain knowledge and therefore restrict ourselves to showing the deviation from normal situations by using the difference of probabilities $\Delta p_i = p_i^{prior} - p_i^e$. This allows us to more easily show anomalies, i.e., situations that are different from normal and may require attention. The hypotheses are formulated such that $\Delta p_i > 0$ means that more attention is required, while $\Delta p_i < 0$ means no special attention is required in the context of the main hypothesis. Each node i is visualized as a box containing its label. Its probability Δp_i is visualized in the side of the box using an indicator with red boxes above the center for $\Delta p_i > 0$ and blue boxes below center for $\Delta p_i < 0$, as shown in Figure 1a. We use a color scheme red (hot) and blue (cold) to signify elements requiring attention (a dangerous or anomalous situation) versus elements that do not. As a double encoding, hypotheses that require more attention are visualized with a thicker border.

The edges are visualized as colored, curved lines. Here we use the same color scheme as before, where red signifies a positive influence, and blue signifies a negative influence. The thickness of the edge is determined by $|w_{ij}|$. An edge from a faded hypothesis is also faded–see Figure 1a.

Since the evidence may contain many attribute values, simply displaying the values as text labels in a table will not enable the user to quickly see agreement or contradiction. Also, since each attribute represents an independent and possibly completely different domain, we cannot map a set color map to attribute domains. Therefore, we visualize the evidence in a colored evidence matrix.

The evidence matrix–see Figure 1b–is a table in which each column represents an observation and each row represents a unique attribute. The cells in each row are colored such that each attribute value receives a unique coloring. This allows the user to quickly see on which attributes observations agree and for which attributes there is a conflict. The cells are colored according to the following requirements:

- The total number of different colors k should be minimized.
- For each row, two cells have the same color if and only if they have the same value.
- Cells with no value should have no color.
- For each column, the number of colors should be minimized.

We minimize the number of different colors in a column to reduce visual clutter and to make it easier for the user to see color differences in the horizontal direction. To make clear that the colors are related in the horizontal direction and not in the vertical direction, we separate the rows using black lines and separate the columns using gray lines. From the requirements it follows that k is exactly equal to the maximum number of different attributes per row.

Attributes are connected to their hypotheses using curved, gray lines. To avoid visual clutter, these lines are bundled and routed in between nodes where needed.

2.1 Interaction

To investigate further, we allow the user to interact with the visualization. When hovering over an hypothesis h_i, all paths from to h_i to h_0, as well as all paths to attributes in the evidence are highlighted, allowing the user to explore the structure of the graph and evidence. Additionally, when hovering over a cell in the evidence matrix, attribute values of the observation are shown to the right of the matrix and all values for the attribute are shown below the matrix, as shown in Figure 1b. We visualize the probabilities of the attribute values using an indicator with boxes in a similar style as before.

3 Use Case

In Figure 1a we show an instance of an explanation graph where hypothesis h_0 states that the OOI is an environmental hazard. The probability increase Δp_0 is high enough to raise an alarm such that the operator chooses to investigate. From the visualization it immediately becomes clear that this is due to the hypothesis labeled "restricted area violation". The operator knows this restricted area violation means a rule violation that forbids vessels of a certain type and size (tankers above 10,000 tons) to be near nature preserves. The operator can also immediately see that all sources agree on the relevant attributes of the vessel, and can therefore confidently conclude that the vessel is a potential environmental hazard due to a restricted area violation. In this use case, all observations in the evidence matrix are in agreement, whereas in Figure 1b we show an evidence matrix with contradiction.

4 Conclusions & Future Work

We have presented a method to visualize the reasoning of a decision support system that allows the user to quickly follow the rationale behind the reasoning and confidently take appropriate action. Despite their widely varying heterogeneous attribute values, we show the observations in the evidence in a compact and quick to read matrix.

While the visualization of the rationale is intended for end-users, it has already proven its value in developing the reasoning engine: Several errors in the reasoning engine's model have been detected, which before remained undetected.

The visualization has been designed in close cooperation with both automated reasoning experts and maritime experts. We would like to evaluate the visualization using operational experts, and improve where needed. Additionally, we would like to visualize how the rationale changes over time while more observations become available. Lastly, we would like to try our method in more domains.

Acknowledgments

This publication was supported by the Dutch national program COMMIT. The research work was carried out as part of the Metis project under the responsibility of Embedded Systems Innovation by TNO with Thales Nederland B.V. as the carrying industrial partner.

References

[1] S. Michels, M. Velikova, A. Hommersom, and P. J. Lucas. A decision support model for uncertainty reasoning in safety and security tasks. In *Systems, Man, and Cybernetics (SMC), 2013 IEEE International Conference on*, pages 663–668. IEEE, 2013.

[2] K. Sugiyama, S. Tagawa, and M. Toda. Methods for visual understanding of hierarchical system structures. *Systems, Man and Cybernetics, IEEE Transactions on*, 11(2):109–125, Feb 1981.