

Causal Bayesian Networks for Robust and Efficient Fusion of Information Obtained from Sensors and Humans

G. Pavlin^{1,2}, M. Maris^{2,3}, F. Groen³

¹Thales research and technology,
Delftechpark 24, Delft, The Netherlands
E-mail: pavlin@decis.nl

²TNO Defence, Security and Safety
Oude Waalsdorperweg 63, Den Haag, The Netherlands
E-mail: marinus.maris@tno.nl

³University of Amsterdam
Kruislaan 403, Amsterdam, The Netherlands
E-mail: groen@science.uva.nl

Abstract – This paper discusses techniques for fusion in contemporary situation assessment applications. Such applications often require reasoning about phenomena that cannot be observed directly, but information about their effects (i.e. symptoms) can be accessed through the existing sensory and communication infrastructure. Reasoning about hidden phenomena requires interpretation of relevant observations. Observations can be of heterogeneous types and can originate from humans as well as various sensory systems. Interpretation in such settings can be very challenging, as there might exist complex dependences between different phenomena. In addition, we are often confronted with significant modeling and observation uncertainties. Particularly challenging is the fact that a large portion of such information often originates from humans. Consequently, it can be very difficult to obtain perception models that precisely describe the distributions of hidden phenomena and human reports. In this paper we show that Bayesian networks (BNs) are suitable for the development of fusion systems in such settings, because they can efficiently describe the monitoring domains. Moreover, BNs support construction of efficient and robust distributed fusion systems.

Keywords – Information fusion, Bayesian networks, heterogeneous information sources.

I. INTRODUCTION

Contemporary decision making processes require accurate and quick situation assessment. Particularly challenging and increasingly relevant is situation assessment in crisis management, such as detection of toxic gases, disease outbreaks, fires, etc. In such settings, critical hidden events must be inferred through interpretation (i.e. fusion) of large quantities of uncertain and very heterogeneous information. The information can be accessed via static sensors or ad-hoc sensor networks formed at runtime as sensors are delivered to the area of interest via mobile platforms (e.g. unmanned aerial vehicles). In

addition, in crisis situations we might be able to obtain valuable information from humans in the field by using the existing communication and storage infrastructure, such as fixed and mobile phone networks, WWW, databases, etc. In fact, the information originating from humans might be prevalent due to insufficient coverage and density of sensor networks.

Interpretation of different types of information requires adequate domain models which provide a mapping between heterogeneous observations and hypotheses of interest. However, situation assessment in crisis management settings introduces several substantial challenges:

- Domains are often complex, which means that models are abstractions associated with significant uncertainties.
- Information sources are heterogeneous and noisy. The heterogeneity of observations implies complex domain models.
- Constellations of information sources are often not known prior to the operation and they change at runtime.

In this paper we show that probabilistic causal models facilitate design of fusion systems which can cope with these challenges. Namely, monitoring processes can often be viewed as causal stochastic processes, where hidden events cause observations with certain probability. Such processes can be modeled with the help of Causal Bayesian networks (BN) [1], which provide a theoretically rigorous and compact mapping between hidden events of interest and observable events. BNs support construction of powerful distributed fusion systems. In addition, information fusion based on inference in Bayesian networks can be very robust.

Note also, that this paper is intended to give an overview of the most important properties of causal BNs and design approaches to robust and efficient fusion systems. Proofs of the key properties and thorough discussion can be found in the listed references.

II. STATE ESTIMATION

Adequate decision making requires the knowledge of critical situations. We assume that such situations can be associated with states of a discrete variable H with states h_1, \dots, h_n . Hence, from a decision making point of view, we must know the state of a hidden variable H that materialized in a particular situation and remained constant during the fusion process. We call such a state the *true state* and denote it by h^* .

Typically, we cannot observe the states of H directly, but often we can interpret a set of observations \mathcal{E} in order to estimate probability distribution¹ $\hat{P}(H|\mathcal{E})$ over the possible states of H and \mathcal{E} . On the other hand, we are reasoning about phenomena where we know that one of the possible states of variable H has materialized. Consequently, the state estimation is equivalent to classification of possible states of interest. Given a discrete hypothesis variable H and a set of observations \mathcal{E} , the estimation corresponds to the determination of the hypothesis \hat{h} for which the estimated posterior probability $\hat{P}(H|\mathcal{E})$ or, equivalently, the joint probability $\hat{P}(H, \mathcal{E})$ is maximum:

$$\hat{h} = \underset{h_i}{\operatorname{argmax}} (\hat{P}(h_i, \mathcal{E})) \quad (1)$$

Thus fusion corresponds to the computation of $\hat{P}(h_i, \mathcal{E})$ over the states h_i and \mathcal{E} for every possible h_i .

Moreover, we assume that each state of a random variable H is associated with an optimal decision $d_i \in \{d_1, \dots, d_n\}$. If the decision maker knew that state h_i of H materialized, she would make the decision d_i corresponding to that state. Since the decision maker cannot directly observe the ground truth h^* , she makes use of a posterior probability distribution $\hat{P}(h_i|\mathcal{E})$ over the possible states of H . Moreover, for each possible state h_i we define a threshold θ_{h_i} in such a way that only one of the possible thresholds can be exceeded at a time. If the estimated $\hat{P}(h_i|\mathcal{E}) > \theta_{h_i}$ then decision d_i is made as if the true state would be h_i . This decision rule is optimal if $\hat{P}(h_i|\mathcal{E})$ is the true distribution over variables H and \mathcal{E} [2].

In this context, fusion is accurate in a particular situation if the decision based on $\hat{P}(h_i, \mathcal{E})$ is identical to the decision that would have been made if the true state were known. In the next section we show that causal Bayesian networks support efficient computation of $\hat{P}(h_i, \mathcal{E})$ in complex domains.

III. PROBABILISTIC INFERENCE WITH CAUSAL BAYESIAN NETWORKS

In a significant class of domains we can view sequences of observations as outcomes of stochastic causal processes, which in turn can be described through probabilistic causal models (see example in figure 1). Such models basically capture dependencies, such as the fact that materialization of certain states makes realization of other states more/less probable.

¹ In this paper modeling parameters and estimated probabilities are denoted by $\hat{P}(\cdot)$ while true probabilities are denoted by $P(\cdot)$.

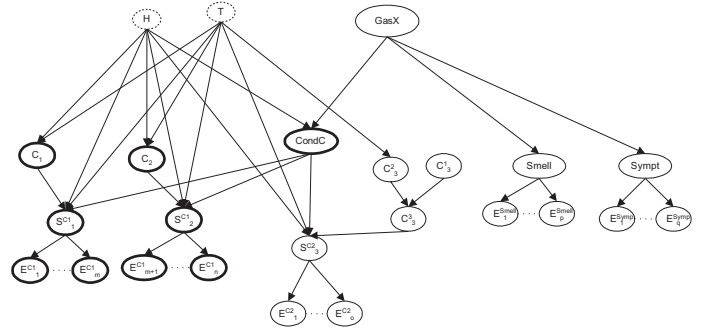


Fig. 1. A causal model capturing relations between the possible states, such as presence/absence of GasX represented by node $GasX$ and different types of sensor observations. Highlighted portion of the graph captures causal relations between $CondC$, sensor propensities S_i^{C1} and sensor reports E_j^{C1} .

By using causal models we can formulate information fusion as a mapping from different observations to possible hidden causes. Particularly useful in this context are Bayesian networks (BNs), probabilistic models which can describe stochastic causal processes in a mathematically rigorous and compact way [1]. BNs support efficient modeling of causal processes involving very heterogeneous phenomena and uncertainties. In such models every hidden or observable phenomenon is represented by a node. In addition, BNs support theoretically rigorous belief propagation which can be used for the computation of $\hat{P}(H|\mathcal{E})$.

A Bayesian network is defined as a tuple $\langle \mathcal{D}, P \rangle$, where $\mathcal{D} = \langle \mathcal{V}, E \rangle$ is a directed a-cyclic graph (DAG) defining a domain $\mathcal{V} = \{V_1, \dots, V_n\}$ and a set of directed edges $\langle V_i, V_j \rangle \in E$ over the domain. The joint probability distribution over the domain \mathcal{V} is defined as

$$\hat{P}(\mathcal{V}) = \prod_{V_i \in \mathcal{V}} \hat{P}(V_i | \pi(V_i)),$$

where $\hat{P}(V_i | \pi(V_i))$ is the conditional probability table (CPT) for node V_i given its parents $\pi(V_i)$ in the graph. In this paper, we assume that each node represents a discrete variable.

It turns out that BNs are relevant for a significant class of information fusion systems, where we can identify causal relationships between hidden and observable phenomena. We illustrate this with the help of a simple gas detection system that makes use of two different types of sensors as well as reports from humans (see the causal model in Figure 1).

The existence of $GasX^2$ is represented by binary variable $GasX$, where the instantiations $GasX = true$ and $GasX = false$ correspond to the presence and the absence of $GasX$, respectively. In addition, we assume that the system uses two types of chemical sensors in conjunction with

² For the sake of simplicity we say that a toxic gas is present if its concentration exceeds some critical value. Otherwise we say that the gas is absent.

thermometers and hygrometers. Both types of chemical sensors measure the conductivity (i.e. electrical current) in a semiconductor exposed to the contaminated air; $GasX$ reacts with the semiconductor which influences the conductivity. But each sensor type evaluates the signal in a different way. States of the binary variable $CondC$ correspond to the situations where electrical current under ideal circumstances would either exceed some detection threshold (i.e. $CondC = true$) or be below that threshold (i.e. $CondC = false$). The current will exceed the detection threshold with the probability $P(CondC = true|GasX = true)$ if the gas concentration exceeded a critical value, i.e. $GasX = true$. Also, with probability $P(CondC = false|GasX = true)$ the current might not exceed the detection threshold despite $GasX = true$. This might be due to the presence of another gaseous component which would inhibit the desired chemical reaction between $GasX$ and the semiconductor.

Moreover, we assume that $GasX = true$ within a finite time interval, i.e. time slice. Within such a time slice electronic circuitry of a sensor evaluates the current and generates a stream of sensor reports. Let the j th report from a sensor of type i be represented by a random variable E_j^i . A sensor report claiming the presence or absence of some hypothetical phenomenon is expressed by $E_j^i = true$ or $E_j^i = false$, respectively. Given the existence of a critical gas concentration, the i th sensor is considered to be working correctly if $P(E_j^i = true|GasX = true)$, the probability that a report from this sensor asserts the presence of the gas, is greater than $P(E_j^i = false|GasX = true)$, the probability that a report asserts the absence of the gas. Similarly, given the absence of a critical gas concentration, the relation $P(E_j^i = true|GasX = false) < P(E_j^i = false|GasX = false)$ should be satisfied when a sensor is considered to be working correctly. In other words, if a sensor works correctly the majority of the sensor reports should assert the presence or absence of $GasX$ given $GasX = true$ or $GasX = false$, respectively.

However, probability distribution $P(E_j^i|GasX)$ does not depend only on the presence of the phenomenon that we are trying to infer. For example, the sensor electronics might fail due to proximity to a source of electromagnetic noise, the sensor could be covered by ice, etc. In other words, the true distribution over reports depends on many factors which are often not well known. We avoid detailed modeling of causal processes producing sensor reports by introducing the *sensor propensity* concept that represents two types of situations, each corresponding to certain combinations of the states of the electronic components and the states in the domain (e.g. conductivity of the chemical element) [3]. We represent the *sensor propensity* by a binary variable S_i . Here, $S_i = true$ denotes the class of state combinations which influence the sensing process, such that the probability of obtaining a sensor report affirming the existence of some phenomenon is greater than 0.5; i.e. $P(E_j^i = true|S_i = true) > P(E_j^i = false|S_i = true)$. Situations corresponding to $S_i = false$ would influence the

distribution over types of sensor reports, such that $P(E_j^i = true|S_i = false) < P(E_j^i = false|S_i = false)$.

Besides the semiconductor sensors we assume that there are humans in the area who have olfactory reactions to $GasX$ and who submit reports of what they smell via a call service. The states of binary variable $Smell$ represent situations in which people familiar with a typical smell of $GasX$ are either able or unable to recognize the smell under ideal circumstances, respectively. The probabilities of having either of the possible two states given $GasX$ are captured by the conditional probability distribution $P(Smell|GasX)$. Moreover, each individual report is represented by a node E_i^{Smell} . Since the smelling capabilities are not perfect, we use a probabilistic "sensor" model $P(E_i^{Smell}|Smell)$. Similarly, aid workers might be able to report about health symptoms which are typical results of exposure to $GasX$. A situation in which observable symptoms take place is denoted by variable $Sympt$ and the reports are denoted by variables E_i^{Sympt} . A BN describing such a domain is depicted in Figure 1.

Also, given a hypothesis variable H we can define a *conditionally independent network fragment* as follows:

Definition 1: Given a BN and a classification variable H , the i^{th} *conditionally independent network fragment* \mathcal{F}_i^H is a portion of the BN corresponding to a set of nodes that include node H and are conditionally independent of other parts of the BN given H .

For example, in Figure 1 we can identify fragment \mathcal{F}_2^{GasX} consisting of nodes $GasX$, $Smell$ and several E_m^{Smell} . Moreover, conditional independence between the sets of variables corresponds to a specific factorization of $\hat{P}(H, \mathcal{E})$. Namely, $\bigcap_i \mathcal{F}_i^H = \{H\}$, which means that each network fragment \mathcal{F}_i^H is associated with a factor $\phi_i(H)$ resulting from a marginalization of all variables from this fragment except H and the instantiated evidence nodes. For example, the network depicted in Figure 1, corresponds to the following factorization of the joint probability distribution over a hypothesis variable $GasX$ and a set of observations \mathcal{E} :

$$P(GasX, \mathcal{E}) = \phi_0(GasX)\phi_1(GasX)\phi_2(GasX)\phi_3(GasX),$$

where $\phi_0(GasX) = P(GasX)$, $\phi_2(GasX) = \sum_{Smell} P(Smell|GasX) \prod_{\mathcal{E}_4} P(E_m^{Smell}|Smell)$, $\phi_3(GasX) = \sum_{Sympt} P(Sympt|GasX) \prod_{\mathcal{E}_5} P(E_n^{Sympt}|Sympt)$. Factor ϕ_1 corresponds to a rather complex marginalization involving variables representing states of sensor components $C_1, C_2, C_3^1, C_3^2, C_3^3$, etc. $\mathcal{E}_1, \dots, \mathcal{E}_5$ denote sets of observations³ corresponding to instantiations in different parts of the BN.

Note that with each sensor we introduce an independent partial causal process which is initiated through some hidden phenomenon. For example, by introducing a new gas sensor the presence of gas will initiate different processes in the sensor's circuitry which will eventually produce sensor reports. This

³ reports from sensors and humans

is reflected in a BN by introducing a new propensity node S_m along with variables representing states of sensor components and sensor report nodes E_j^m (see Figure 1).

IV. DISTRIBUTED FUSION

While BNs can cope with modeling and observation uncertainties, in contemporary fusion problems we are often confronted with additional challenges. Namely, constellations of information sources change at runtime. Furthermore, every information source must be captured in a BN. Thus for each constellation of information sources we need a specific BN. In addition, large quantities of heterogeneous information accessed through the existing communication and sensing infrastructure often require large BNs which in turn require significant processing and communication resources.

These problems can be tackled by introducing modular approach to modeling and processing. By using basic fusion building blocks, each specialized for a limited fusion task, adequate domain models (i.e. BNs) can be assembled at runtime as the information sources are discovered. In addition, fusion in such an assembled network can easily be distributed throughout several machines thus avoiding processing and communication bottlenecks. Beside sound and efficient fusion algorithms, basic fusion modules must support also efficient communication and cooperation protocols. In addition, a distributed fusion system should be able to adapt to the current situation autonomously. Therefore, modules should form fusion systems consisting of relevant modules autonomously and they should be able to reason about resource allocation with respect to sensing and processing capacity. In order to be able to cope with such complex functionality in a systematic way, we make use of the multi agent systems paradigm [4].

Definition 2 (Fusion Agent) **Fusion agent** A_i is a processing unit, a module, which can compute probability distributions over variables in its local BN.

- A local BN is defined through a local DAG $\mathcal{G}_i = \langle \mathcal{V}_i, E_i \rangle$ and a set of conditional probabilities which encode factorization of a joint probability distribution $P(\mathcal{V}_i)$ over local variables.
- Each agent A_i maintains a set of **service variables** $\mathcal{R}_i \subset \mathcal{V}_i$ and a set of **input variables** $\mathcal{L}_i \subset \mathcal{V}_i$.
- Each agent A_i can compute marginal posterior probabilities over the local service variables.

Two agents can exchange their local estimates of marginal (posterior) distributions for any variable contained in local BNs of both agents.

A. Fusion Organization

Each fusion process depends on the constellation of cooperating agents, which corresponds to a particular problem/task decomposition. In this context we use the concept of a *Fusion Organization* [3]. A particular organization Ω is a function of a given query concept (also hypothesis) and a set of available agents.




Fig. 2. An organization of fusion agents implementing a distributed causal model. Each solid rectangle corresponds to an agent with a local modeling fragment represented by dashed rectangles. Dashed arrows show the flow of the inter-agent fusion messages.

A global task of a fusion organization Ω is computation of probability distribution $P(H|\mathcal{E})$ over some hypothesis variable $H \in \mathcal{R}_i$ of agent A_i which correctly reflects the entire evidence set \mathcal{E} . This is only the case if $P(H|\mathcal{E}) = P'(H|\mathcal{E})$, where $P'(H|\mathcal{E})$ is computed through propagation of the entire evidence \mathcal{E} in a monolithic BN which correctly captures all dependencies and independences that exist between variables in and between the agent's local models in Ω . This requires well defined cooperation of fusion agents since evidence \mathcal{E} corresponds to instantiations of variables in different agents in a fusion organization Ω .

B. Model Decomposition

BNs support construction and combination of local BNs of different agents, such that an agent organization Ω supports correct computation of $P_i(H|\mathcal{E})$ for some hypothesis variable H in agent A_i without any centralized configuration and coordination of local fusion processes. This is achieved by using simple rules for the generation of local BNs from a monolithic BN and agent organization [3].

The *design rules* guarantee that all dependencies between different local graphs obtained through partitioning of an arbitrary monolithic BN are preserved. The *assembly rules* prescribe how agents with different local BNs should assemble fusion organizations which support exact belief propagation without any centralized control of local fusion processes⁴.

C. Distributed Perception Networks

The introduced theory has been implemented in a SW toolbox called *Distributed Perception Networks (DPN)* [3], which supports design of complex agent based fusion systems that can cope with changing constellations of information sources.

⁴ i.e. central processes which determine messaging sequences between agents.

DPNs are in essence distributed self-organizing classifiers, which consist of agents with very heterogeneous domain expertise. Agents wrap information sources and provide uniform communication and fusion protocols. In contrast to other approaches to distributed inference in BNs [5], the DPN framework supports theoretically correct fusion which does not require any compilation of secondary fusion structures (e.g. junction trees [5]) that span several modules. As new information sources wrapped by DPN agents enter the scene, the domain models of fusion systems are adapted on the fly, without any centralized control. DPN agents wrapping mobile sensor suites supply local BNs, arbitrarily sophisticated sensor models, which are plugged into the overall system as the agents join the fusion organization. In other words, each agent contributes a local model which relates agent’s observations with the rest of the distributed causal model.

In addition, DPN agents exploit causal models to generate meaningful queries which are communicated to the people via SMS or WWW and can be replied by simple yes/no (e.g. ”Do you smell Ammonia?”). In this way we obtain only information which can be processed by the fusion system and interpretation of natural language is avoided. In this way DPN agents facilitate efficient acquisition of very heterogeneous information from humans in the field. This is very relevant for crisis management applications where humans can be viewed as omnipresent and often the only available information sources. Clearly, sound fusion of such information is very challenging, since we cannot obtain precise sensor models for human reports. However, we can show that this is possible if we use BNs (see the following section).

Also, human information sources have limited capacities; a person can reply to a limited number of SMS queries in a certain time interval. Therefore, a DPN system should first generate queries that have the greatest impact on its belief about the critical events. Fortunately, rigorous causal models support theoretically sound and efficient approaches to the ranking of evidence types with respect to their relevance [6].

V. ROBUST INFORMATION FUSION

If parameters in a BN are identical to the true distributions over the modeled variables, then a Bayesian classifier is optimal (see [2]). However, in many real world domains, it is very difficult or even impossible to obtain probabilistic models that precisely capture the true probability distribution over the phenomena in the observed domains. Namely, training data sets are finite and human experts cannot precisely specify the domain models.

Fortunately, it has been shown that for a significant class of models we can achieve very accurate estimation even if we use imprecise modeling parameters. The recently proposed theory of Inference Meta Models (IM) describes information fusion with BNs from a coarse runtime perspective and can be used to characterize inherently robust fusion systems [7]. IM is based

on very coarse assumptions and exposes properties of BNs that are relevant for the construction of fusion systems which can process information obtained from humans in a robust way.

With the help of IM we show that inference in BNs featuring many conditionally independent network fragments (i.e. factors) can be very insensitive to the values of the modeling parameters and can have asymptotic properties with respect to the estimation accuracy.

We can achieve very robust classification by taking into account updating tendencies of belief propagation in a BN. Recall that each instantiation of a network fragment that is separated from the other parts of the network by H corresponds to a factor in the expression describing the distribution over classification variable H (see definition 1). For each such conditionally independent network fragment we can observe that if we multiply the conditional equation with the corresponding factor and normalize over all states of H , the posterior probability of one state will increase the most. One could say that given a certain observation one state ‘wins’. Obviously, the state that wins sufficiently often will end up with the highest posterior probability. This suggests that it is not the exact factor values, but the relations between them that matter most with respect to the classification accuracy. Therefore, for each factor $\phi_i(H)$ we introduce a factor reinforcement:

Definition 3 (Factor Reinforcement) Assume a classification variable H and a fragment \mathcal{F}_i^H . Given some instantiation of the evidence variables within \mathcal{F}_i^H , we can compute a factor $\phi_i(h_j)$ for each state h_j of variable H and determine the corresponding factor reinforcement r_i^H as follows:

$$r_i^H = \arg \max_{h_j} \phi_i(h_j). \quad (2)$$

In other words, reinforcement r_i^H is a function that returns the state h_j of variable H , whose probability is increased the most (i.e. reinforced) by instantiating nodes of the fragment \mathcal{F}_i^H corresponding to factor $\phi_i(H)$.

Moreover, we can define an accurate reinforcement:

Definition 4 (Accurate Reinforcement) Let H be a classification variable and let h^* be its hidden true value. A reinforcement r_i^H contributed by factor ϕ_i is accurate iff:

$$h^* = r_i^H. \quad (3)$$

Thus, the true state of H is reinforced. It turns out that in a given situation, an inference step results in an accurate factor reinforcement for very different parameters as long as the parameters satisfy very simple relations.

By considering this we can design very robust classifiers based on the *Reinforcement propagation algorithm* [7], a simple yet very robust approach to inference in BNs. The *Reinforcement propagation algorithm* reduces belief updating in BNs to counting of state updates (i.e. reinforcements) at different nodes. The algorithm basically counts how often the instantiation of observable nodes increased or decreased the probability of different states of hidden variables.

Obviously, by ignoring the magnitude of belief updates the same evidence results in the same updating tendencies for very different CPT parameters. On the other hand, however, certain information encoded in the model is ignored. Fortunately, we can still obtain very accurate classifiers with asymptotic properties if BNs feature many conditionally independent fragments given the hypothesis variable H .

This is true if each factor $\phi_i(H)$ corresponding to a fragment \mathcal{F}_i^H in BN reinforces correct state of H with the probability $p_\phi > 0.5$ in a particular situation. This seems to be a rather strong assumption. However, we can show that for BNs with tree topologies $p_\phi > 0.5$ if (i) belief updating is based on the *Reinforcement propagation algorithm* and (ii) the true conditional probabilities for any pair of dependent variables from a fragment satisfy simple greater-than/smaller-than relations that are also captured by the corresponding CPT parameters. A thorough rationale defending this assumption can be found in [7]. Obviously, it is plausible to assume that designers or learning algorithms can easily identify such relations.

A. Fusion of Human Reports

The properties illustrated in the previous section suggest that we could exploit huge quantities of relevant information accessible through the existing communication infrastructure, such as GSM networks, Internet, etc. Such information usually originates from humans and it might be very difficult to obtain modeling parameters that precisely describe the true probability of obtaining certain types of reports (e.g. $P(E_m^{Smell}|Smell)$ and $P(Smell|GasX)$). However, we can assume that the simple relations between the true probabilities can be identified reliably. Moreover, we could obtain reports about many different symptoms corresponding to many independent network fragments relative to the classification variable. In other words, we might achieve very accurate classification, despite significant discrepancies between the true and the modeling distributions.

VI. APPLICATIONS

A typical application example is the detection and tracking of gases in real world applications. This is because there are no devices that would reliably detect different types of gases. State of the art gas sensors are usually noisy and they can fail due to different reasons, such as physical damage or unusual operating conditions. Clearly, noisy and faulty information sources often produce erroneous reports, which can have a negative impact on the assessment quality. In order to be able to cope with these problems, the gas detection must be based on the interpretation of large quantities of different types of information.

As the currently available sensor networks are typically not dense enough to support robust detection and tracking, situation assessment is often strongly supported by information obtained from humans, who can be considered as omnipresent information sources that can provide rich and useful information.

By exploiting the existing communication technology, such as mobile phones, we can access humans in the affected area who can provide valuable information on the situation, as for example a specific smell and effects on the health of humans (e.g. haziness, irritation, etc).

In general, applications of our method are found in complex automated information processing systems that need to fuse information from many different sources. In particular, operators who need to make an assessment about a recently emerged (potentially threatening) situation, can be strongly supported by our approach. Currently, in cooperation with DCMR Milieu-Dienst Rijnmond, an organization monitoring air quality in the Rotterdam harbor, we are investigating how the DPN architecture and IM theory could be used in their gas detection system.

VII. CONCLUSION

The DPN architecture in conjunction with the IM theory provides a solid basis for the development of efficient and reliable fusion systems, which support situation assessment in crisis management. Robust assessment requires fusion of large quantities of heterogeneous information. In addition, we have to deal with ad-hoc constellations of information sources; i.e. it is difficult to provide adequate domain models prior to the operation. To tackle these challenges, distributed causal BNs are proposed; adequate probabilistic models, mapping observations to hypotheses, are assembled during runtime, creating a system for robust fusion.

VIII. ACKNOWLEDGEMENTS

The authors would like to thank the Decis laboratory for providing the facilities and the research environment for carrying out this research. We thank our colleagues from the University of Amsterdam for their academic support and in building components for the DPN system.

IX. REFERENCES

- [1] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, 1988.
- [2] R. O. Duda and P. E. Hart, *Pattern classification and scene analysis*, John Wiley & Sons, New York, 1973.
- [3] G. Pavlin, P. de Oude, M. Maris, J. Nunnink, and T. Hood, "A distributed approach to information fusion based on causal probabilistic models," Tech. Rep., 2007.
- [4] N. Jennings, K. Sycara, and M. Wooldridge, "A roadmap of agent research and development," in *Autonomous Agents and Multi-Agent Systems*, vol. 1, pp. 7 – 38. July 1998.
- [5] Y. Xiang, *Probabilistic Reasoning in Multiagent Systems: A Graphical Models Approach*, Cambridge University Press, 2002.
- [6] J. Nunnink and G. Pavlin, "A probabilistic approach to resource allocation in distributed fusion systems," in *Proc. Int. Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS-05)*, Utrecht, Netherlands, 2005, pp. 846–852.
- [7] G. Pavlin and J. Nunnink, "Inference meta models: Towards robust information fusion with bayesian networks," in *The Ninth International Conference on Information Fusion*, Florence, Italy, 2006.