Demining Sensor Modeling and Feature-Level Fusion by Bayesian Networks

Silvia Ferrari, *Member, IEEE,* and Alberto Vaghi

***Abstract*—A method for obtaining the Bayesian network (BN) representation of a sensor's measurement process is developed so that the problems of sensor fusion and management can be approached from a unified point of view. Uncertainty, reliability, and causal information embedded in the sensor data are used to build the BN model of a sensor. The method is applied to model ground-penetrating radar, electromagnetic induction, and infrared sensors for humanitarian demining. Structural and parameter learning algorithms are employed to encode relationships among mine features, sensor measurements, and environmental conditions in the BN model. Inference is used to estimate target features in the presence of heterogeneous soil and varying environmental conditions. A multisensor fusion technique operating on BN models is developed to exploit the complementarity of the sensor measurements. Through the same approach, a BN classifier is obtained to estimate the target typology. The BN models and classifier also compute so-called confidence levels that quantify the uncertainty associated with the feature estimates and the classification decisions. The effectiveness of the approach is demonstrated by implementing these BN tools for the detection and classification of metal and plastic landmines that are characterized by different shape, size, depth, and metal content. Through BN fusion, the accuracy of the feature estimates is improved by up to 64% with respect to single-sensor measurements, and the number of objects that are both detected and classified is increased by up to 62%.**

***Index Terms*—Bayesian networks (BNs), classification, fusion, landmine detection, sensor modeling.**

I. Introduction

THE process of sensor fusion consists of combining the measurement data from multiple and, possibly, heterogeneous sensors in order to infer the characteristics or *features* of one or more targets. By taking into consideration the information obtained from different sensors the performance of detection, tracking, and identification algorithms can be significantly improved. The operational benefits of multisensor systems over single-sensor systems include extended spatial and temporal coverage, reduced uncertainty, and increased robustness that are due to the redundancy and complementarity of the sensor measurements [1, pp. 1–13]. The process is referred to as data-level, feature-level, or decision-level fusion, depending on the type of information that it combines. Several techniques, such as Bayes, linear discriminant, Dempster–Shafer, and voting methods, have been proposed for correlating and integrating this information, while taking into account the uncertainty associated with each sensor (see [1] and [2] for a comprehensive review). These techniques differ in the performance metric they optimize, and in how they incorporate the sensor uncertainty declarations measured by so-called *confidence* or *belief* functions.

In many applications, the sensors' measurements and uncertainty are greatly affected by their operational and environmental conditions. However, there are no systematic approaches for incorporating these conditions in the fusion process. Typically, confidence values (or levels) are chosen based on available sensor parameters that relate to the accuracy of the measurement process in a particular application. For instance, in the detection of landmines, the sampled and projected energy of ground-penetrating radars (GPR) are used as confidence values [3]. Other approaches obtain these values by comparing the predictions of sensor simulations to the actual sensor measurements [4], a procedure that also is used to extract features from the raw data. Although they account for *a priori* knowledge, these simulations require a detailed representation of the physical processes involved in a particular measurement process. This may require that a different simulation be developed for every sensor type by means of disparate formalisms, such as, ordinary or partial differential equations, and sensor data in the form of lookup tables. Thus, the respective results may only be coupled with fusion and processing techniques in an ad-hoc fashion. Also, due to the lack of a compact mathematical representation or *model* they can only be used in an open-loop fashion, i.e., by running a separate computer simulation for every set of conditions considered during tests or actual operation.

In this paper, a systematic approach for obtaining mathematical models of sensor measurements by means of Bayesian networks (BNs) is developed. These models are constructed based both on expert knowledge and on sensor data collected from tests or experiments *a priori*. The influence of operational and environmental variables on a particular sensor type is captured by its BN model, which can then be used for interpreting sensed information, for instance by means of sensor fusion and feature extraction. Unlike sensor simulations, the BN model of any sensor can be obtained by a common formalism, thereby allowing heterogeneous sensor measurements to be processed by means of unified techniques. Another advantage is that a BN model comprises a mathematical representation of the sensor that can be formally analyzed and used to derive closed-loop control and decision-making policies.

Manuscript received January 30, 2005; revised June 1, 2005. This work was supported in part by the Office of Naval Research Young Investigator Program (Code 321) and in part by the National Science Foundation under Grant ECS 0300236. The associate editor coordinating the review of this paper and approving it for publication was Dr. Giorgio Sberveglieri.

S. Ferrari is with the Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27708-0005 USA (e-mail: sferrari@duke.edu).

A. Vaghi is with the Politecnico di Milano, 20030 Senago (MI), Italy (e-mail: vaghi.alberto@tin.it).

Digital Object Identifier 10.1109/JSEN.2006.870162

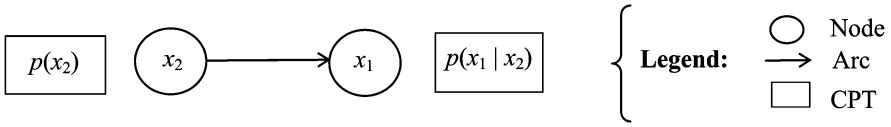

Fig. 1. Two-node BN, with attached CPTs.

The BN approach is demonstrated by modeling electromagnetic induction (EMI), GPR, and infrared (IR) sensors for the detection and classification of metal or plastic buried landmines. A BN model is obtained for each sensor type using a database of measurements collected from several known minefields with variable environmental conditions. Then, the three BN models are used to process the measurements obtained from a field with unknown buried objects and mines and known environmental conditions. For a single sensor, target features are estimated from the measurements through a BN *inference* algorithm (Section III-B). When multiple sensors collect measurements from the same minefield, feature-level fusion is carried out by combining the Dempster–Shafer (D-S) method [5], [6] with the BN approach (Section III-C). Finally, a BN classifier is developed to determine the probability that objects detected are either mines or false alarms (Section III-D).

II. Background on BNs

BNs organize the body of knowledge in a given system by mapping causal relationships among all relevant variables. They can be used to estimate unknown variables and make predictions by combining probabilistic data with heuristic arguments. A probabilistic model of a system is structured by building a graph comprised of a set of variables or *nodes*, and a set of directed links or *arcs* connecting the nodes. Each variable has a finite set of mutually exclusive state values, and each arc represents a dependency among the variables it connects. By attaching a table of conditional probabilities to each node, the network also represents the extent to which the variables are likely to affect each other. The nodes together with the arcs form a *directed acyclic graph* (DAG) [7].

Hence, a BN is defined by a triple $(\mathcal{G}, \Omega, \mathcal{P})$, where $\mathcal{G} = (\mathcal{X}, \mathcal{A})$ is a DAG with the set of nodes $\mathcal{X} = \{x_1, \ldots, x_n\}$, and the set of arcs $\mathcal{A} = \{(x_i, x_j) \mid x_i, x_j \in \mathcal{X},\ x_i \neq x_j\}$ representing dependencies among the node variables. Ω is the space of all possible state values or *instantiations* of the variables in $\mathcal{X}$, and $\mathcal{P}$ is a probability distribution over Ω realized by the *conditional probability tables* (CPTs) attached to the nodes. If there is an arc from x_i to x_j, x_i is said to be a *parent* of x_j. Given a node $x_j \in \mathcal{X}$, the notation $pa(x_j)$ is used to denote the set of parents of x_j in $\mathcal{X}$.

A simple example of BN representing the dependency between two variables x_1 and x_2 is shown in Fig. 1. Suppose x_2 is a variable with r_2 instantiations, $x_2^1, \ldots, x_2^{r_2}$, then $p(x_2)$ denotes its probability distribution over these values. $p\left(x_2 = x_2^\ell\right)$, or simply $p\left(x_2^\ell\right)$, represents the probability of the ℓ^{th} instantiation of x_2. The CPT attached to node x_1 is $p(x_1|x_2)$, which contains all probabilities $p\left(x_1^m|x_2^\ell\right)$, with $m = 1, \ldots, r_1$, and $\ell = 1, \ldots, r_2$. These conditional probabilities are referred to as BN parameters. If x_1 is instantiated (e.g., $x_1 = x_1^m$), the probability that x_2 assumes any one of its possible values is given by *Bayes' rule*

$$p\left(x_2|x_1^m\right) = \frac{p\left(x_1^m|x_2\right)p(x_2)}{p\left(x_1^m\right)} \tag{1}$$

$p\left(x_1^m\right)$ is obtained by *marginalization* of the CPTs, considering all r_2 instantiations of x_2, as follows:

$$p\left(x_1^m\right) = \sum_{\ell=1}^{r_2} p\left(x_1^m, x_2^\ell\right) = \sum_{\ell=1}^{r_2} p\left(x_1^m|x_2^\ell\right) p\left(x_2^\ell\right). \tag{2}$$

Then, x_2 is *inferred* from x_1^m using (1). Typically, x_2 is estimated by choosing the value with the highest probability, i.e., $\hat{x}_2 = x_2^\ell$ with $p\left(x_2^\ell|x_1^m\right) > p\left(x_2^k|x_1^m\right)$ for $\forall\ k \neq \ell$, and the probability $p\left(x_2^\ell|x_1^m\right)$ constitutes the *confidence level* of the estimate.

When processes and relationships between the variables of interest are unknown, a BN model of the system can be established by learning the structure $(\mathcal{A})$ and CPTs $(\mathcal{P})$ from a batch of cases. Typically, each case comprises the instantiations of all variables in $\mathcal{X}$ observed or measured from the real system. Learning techniques also have been devised to learn from incomplete cases, where some of the variables are unknown [8]. Two of these techniques are implemented and compared in Section V to build the BN models of GPR, EMI, and IR sensors. Once a BN model is obtained, it can be used to propagate knowledge or *evidence* from the observable variables to the unknown ones, similarly to (1)–(2). By using the separation properties of the network and by building its equivalent junction tree, marginalization can be simplified and evidence can be propagated rapidly even in large BNs [7]. These techniques are implemented in Section V for target feature estimation and classification.

The BN model of a system also can be used to automate and optimize system processes such as control and decision-making. In this case, the network can be expanded to include decision nodes and reward or outcome variables, obtaining a so-called *decision graph*. Then, optimal control techniques, such as dynamic programming [9], can be applied to the graph by exploiting the same separation properties used for inference. This topic is the subject of work in progress and will be addressed in a separate paper. The approach for modeling sensors by BNs is illustrated in the following section.

III. Approach

A. Sensor Modeling

A BN approach is developed for modeling key relationships in a sensor measurement process that is assumed to be static, using a system-theoretic perspective. Based on expert knowledge of the sensor's operating principles, the relevant variables

are identified and discretized to determine the nodes of the BN model. Then, the arcs connecting these nodes and the CPTs are determined from a database of sensor data, through batch learning algorithms. The predictive performance of the BN model is tested on a validation set of sensor data that is not used for training. If an important variable is missed, this performance will prove unsatisfactory and the modeling procedure may be repeated for a new set of nodes.

1) Node Selection: The nodes of a BN sensor model are selected by considering the following sets of variables, which typically influence the measurement process.

- Sensor Mode: Most sensors are equipped with adjustable parameters that affect the data collection process, the type of information that is measured, and its reliability. The set of parameters chosen to operate the sensor is referred to as mode, and is denoted by M.
- Sensor State: Sensor characteristics that may be observable but only can be adjusted indirectly through the mode's parameters constitute the sensor state, which is denoted by S.
- Environment: According to the physical phenomena underlying the measurement process, the environment surrounding the targets and the sensors may affect sensor performance. The set of environmental variables that influence the measurements is denoted by E.
- Measured Target Features: The set of measured features F is computed from the raw sensor measurements, and constitutes the information collected by the sensor to detect, classify, and, possibly, track a target.
- Actual Target Features: The set of actual features, T, contains the true values of F, which normally are unknown (except in tests and simulations) and must be inferred from the measurements.

Depending on what type of sensor is being modeled, one or more of these sets may be empty. For instance, in GPR, EMI, and IR sensors, S is empty because the sensor state is uniquely determined by the sensor mode. After the relevant nodes have been selected, all of their possible instantiations must be identified such that they are countable and mutually exclusive. If one or more variables are continuous, they can be discretized or treated as continuous nodes with Gaussian distributions, as explained in [10]. For simplicity, in this paper, all continuous variables (e.g., the mine's depth) are discretized. Subsequently, the chosen sensor variables constitute the nodes of the BN sensor model, $\mathcal{X} = \{M, S, E, F, T\}$, that each have r_i instantiations. If, after the learning phase, the predictive performance of the BN sensor model is unsatisfactory, the set of nodes and the corresponding instantiations may be modified and learning may be repeated. For example, additional target features or environmental variables may be considered and r_i may be increased to improve the granularity of variables that have been discretized.

2) Learning: The final BN arcs and CPTs are determined by a batch learning algorithm. First, an initial structure $\mathcal{A}_0$ is proposed by the designer based on expert knowledge of the sensor working principles. The architecture in Fig. 2 can be used as a guide in building $\mathcal{A}_0$, since it represents the typical relationships between the variable sets or node *clusters* listed above. This

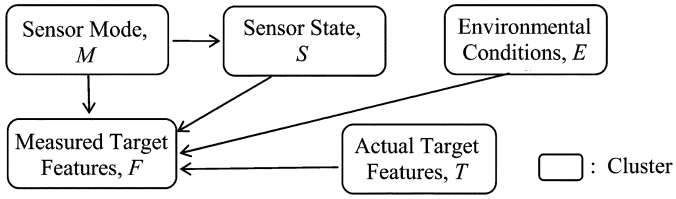

Fig. 2. Typical relationships between the variable sets in sensor models.

step allows to incorporate relationships that are known *a priori*, such as the influence of selected environmental variables and sensor modes on the measured features. For example, the mine's depth measured by a GPR is known to be influenced by its mode and by the soil uniformity. Therefore, arcs initially are placed from the mode and soil-uniformity variables to the measured-mine-depth variable. Then, a database $\mathcal{D} = \{\mathcal{C}_1, \ldots, \mathcal{C}_q\}$ of q independent cases is used to learn the BN structure that best captures the sensor measurement process.

Ideally, the BN would be obtained by considering the weighted average over the inferences of every possible network containing $\mathcal{X}$. Since this approach is too computationally expensive, a BN structure that maximizes a so-called scoring metric is sought instead. Let $\mathcal{B}_\mathcal{S} = (\mathcal{A}, \Theta)$ be a BN structure with parameters $\Theta = \{\theta_1, \ldots, \theta_n\}$ where, θ_i is the CPT attached to node x_i, containing all conditional probabilities $p(x_i \mid pa(x_i))$. A scoring metric is chosen such that the probability of a network structure for a given database, $p(\mathcal{B}_S|\mathcal{D})$, is maximized. Since $p(\mathcal{D})$ is independent of $\mathcal{B}_S$, the joint probability $p(\mathcal{B}_S, \mathcal{D})$, which is more easily computed, is maximized in place of $p(\mathcal{B}_S|\mathcal{D})$.

The scoring metrics proposed in the literature (e.g., [11]–[13]) reduce computation by factorizing $p(\mathcal{B}_S, \mathcal{D})$ based on the following assumptions: all n nodes are discrete; all cases in $\mathcal{D}$ are independent; all possible CPTs and structures are equally probable before observation of the database $\mathcal{D}$; and, finally, an expert-based ordering of nodes is given such that if a variable x_2 precedes x_1 (e.g., Fig. 1), then no causal arc is allowed from x_1 to x_2. Consequently, the scoring metric $p(\mathcal{B}_S, \mathcal{D})$ can be factorized as follows:

$$p(\mathcal{B}_s, \mathcal{D}) = p(\mathcal{B}_s) \cdot \prod_{i=1}^{n} \prod_{j=1}^{r_{\pi_i}} \frac{(r_i - 1)!}{(q_{ij} + r_i - 1)!} \prod_{k=1}^{r_i} q_{ijk}! \qquad (3)$$

where, if $\pi_i = pa(x_i)$, $r_{\pi_i} = \prod_{x_j \in \pi_i} r_j$ is the number of all possible instantiations of π_i. Then, for a given database $\mathcal{D}$, q_{ijk} is the number of cases where $x_i = x_i^k$ and $\pi_i = \pi_i^j$, and $q_{ij} = \sum_k q_{ijk}$ is the number of cases where $\pi_i = \pi_i^j$ regardless of the value of x_i.

The structure of each BN sensor model is learned from a database of measurements obtained by the corresponding sensor type from a set of known targets. Each case in the database is an instantiation of $\mathcal{X}$, i.e.,: $\mathcal{C}_k = \{M^k, S^k, E^k, F^k, T^k\}$. Thus, it consists of the sensor mode, state, and measurements, of the known target features, and of any known environmental conditions. Often, signal processing techniques are necessary in order to extract F from the raw sensor measurements (e.g., [24]–[26]). Then, the so-called K2 algorithm determines the BN structure that maximizes (3) [8]. This structure can be expected to differ

from $\mathcal{A}_0$ because the learning algorithm discovers nodal relationships that are not always known *a priori* from the sensor data. In this paper, the K2 algorithm is implemented by the function *learn_struct_K2* in the MATLAB Bayesian Network Toolbox [17]. This learning procedure is used in Section IV to determine the GPR, EMI, and IR sensor models for landmine detection and classification.

B. Inference

Once the BN model of a sensor has been determined through the procedure outlined in Section III-A2, it can be used to estimate unknown variables based on knowledge of the observable ones. During actual operation, a sensor collects measurements from foreign targets under environmental conditions that may or may not be known. In this paper, it is assumed that all environmental variables are observable. However, the same approach can be applied when E is not fully observable. After processing the raw sensor measurements, the sensor operational conditions M and S, the measured features F, and all variables in E are known. Then, the actual target features T are inferred from the BN model. Information about an observable variable is referred to as *evidence*. It may consist of an instantiation in Ω, e.g., $x = x^\ell$, or of a statement about the impossibility of some of its values, e.g., $x \neq x^\ell$. For example, when the sensor mode is known, the corresponding variable in S is instantiated. The probability distribution of the unknown variables is inferred from the available evidence, similarly to the simple example in (1)–(2).

For a BN with more than two nodes, finding the joint probability $p(\mathcal{X}) = p(x_1, \ldots, x_n)$, which is required for marginalization, can be computationally expensive. In recent years, efficient algorithms have been developed to exploit the following *recursive factorization*

$$p(\mathcal{X}) = \prod_i p(x_i \mid pa(x_i)) \tag{4}$$

as well as BN separation properties. The directed Markov property, for example, states that a variable is conditionally independent of its nondescendents given its parents. The factorization in (4) can be significantly simplified by identifying all conditional independencies among the variables in $\mathcal{X}$, given the evidence. This is achieved by a series of graphical manipulations that construct an acyclic moral graph from the BN, and build a so-called *junction tree* that captures most of the original conditional independencies [18, pp. 9–26]. In this paper, inference is realized through the MATLAB commands *jtree_inf_engine*, *enter_evidence*, and *marginal_nodes* [17].

The BN approach utilizes prior knowledge from tests and field experiments $(\mathcal{D})$ and evidence about the operational sensor to estimate the features of an unknown target. Given the measurements and the sensor operational and environmental conditions, $e = \{M, S, E, F\}$, the actual target features, T, are inferred from the BN sensor model. Inference also provides a measure of reliability of the estimated features in the form of a probability distribution, $p(T|e)$, over all of their possible values.

C. Sensor Fusion

When multiple sensors collect measurements from the same target the BN modeling approach can be combined with a feature-level fusion technique to improve feature estimation. By employing the BN models, a systematic fusion approach is developed that accounts for prior expert knowledge and data, and for real-time operational and environmental conditions. The D-S rule of evidence combination [5], [6] is used to develop a fused probability interval for each state of the target's features. The D-S method is considered to be a generalization of Bayesian inference, because it assigns probability not only to individual variables but also to Boolean combinations of variables. It is well suited for sensor fusion because it considers evidence both in support and in negation of the inferred features. When measurements from different sensors are in disagreement, they are treated as contrasting expert opinions about the target features. In the approach presented here, each expert's confidence level corresponds to the probability distribution obtained through inference from the BN model.

The probability that a variable is in one of its states, say $x = x^\ell$, expresses the certainty of its knowledge and, thus, it is considered to be *in support* of the state x^ℓ. Ignorance of the variable is quantified by the probability of its negation, $p(x \neq x^\ell) \equiv p(\bar{x}^\ell)$, which is a measure of the evidence that refutes x^ℓ. The so-called *plausability* accounts for the evidence that does not rule out a state, i.e., $pl(x^\ell) \equiv 1 - p(\bar{x}^\ell)$, and represents the probability that can be moved in support of that state (x^ℓ). Suppose z is a Boolean combination of the states $\{x^1, \ldots, x^r\}$, whose probability is assessed from evidence collected by two different sources, a and b. Then, the fused probability is obtained from the following D-S rule of combination

$$p(z) = \frac{\sum\limits_{\ell \in \mathcal{S}} p_a(x^\ell) p_b(x^\ell)}{1 - \sum\limits_{m \in \mathcal{N}} p_a(x^m) p_b(x^m)} \quad \ell, m = 1, \ldots, r \tag{5}$$

where $\mathcal{S}$ is the set of variables in support of z and $\mathcal{N}$ is the set of variables that negate z, i.e.,: $\left\{\ell \in \mathcal{S} : (x^\ell)_a \cap (x^\ell)_b = z\right\}$ and $\left\{\ell \in \mathcal{N} : (x^\ell)_a \cap (x^\ell)_b = \emptyset\right\}$. $p_s(x^\ell)$ denotes the probability that $x = x^\ell$ according to source s, and $\cap$ is the intersection of two states.

When two or more sensors collect measurements from the same target, the respective BN models are used to obtain the probability distribution of each feature, $x_i \in T$, over its possible r_i states. Let $e_s = \{M^s, S^s, E^s, F^s\}$ be the evidence available from the s^{th} sensor measurements and operational conditions. Then, its BN model is used to infer the target features and the distribution $p(T|e_s)$, which includes $p(x_i|e_s)$ or, simply, $p_s(x_i)$. Thus, if a feature x_i is measured by two sensors, a and b, the fused probability of each of its states can be obtained from (5), which simplifies to

$$p\left(x_i^\ell\right) = \frac{p_a\left(x_i^\ell\right) p_b\left(x_i^\ell\right)}{1 - \sum\limits_{\ell,m=1}^{r_i} (1 - \delta_{\ell m}) p_a\left(x_i^\ell\right) p_b\left(x_i^m\right)} \tag{6}$$

$\delta_{\ell m}$ is the Kronecker delta, which equals one if $\ell = m$, and equals zero if $\ell \neq m$. With the chosen modeling approach, the

features have mutually exclusive states. Therefore, the probability of a feature variable being in a particular state is in support of that state, e.g., x_i^ℓ. The probabilities associated with all the other states (i.e., x_i^m with $m \neq \ell$) negate or refute x_i^ℓ.

According to the approach illustrated above, the probability distribution for target features $p(T|e)$ is determined through inference for a single sensor, or through inference and fusion for multiple sensors. From this distribution, the value of each feature can be estimated by selecting the instantiation with the highest probability, i.e., $\hat{x}_i = x_i^\ell$ with $p\left(x_i^\ell|e\right) > p(x_i^m|e)$, $\forall\, m \neq \ell$, and $x_i \in T$. In many applications, knowing the probability or confidence level, $p\left(x_i^\ell|e\right)$, together with the estimate of the features is of great value. For example, this knowledge can be used to classify the target, as explained in the following section.

D. Target Classification

In many sensor systems, the target features are measured in order to determine the target typology. For example, in demining applications, the features of objects detected are needed to determine whether they are mines or clutter. For moving targets, features are usually referred to as state variables, and their dynamic estimate may be used not only for classification, but also for *identification*. Through identification, a dynamic model of the target can be built based on the behavior of observable state variables over time. The BN approach developed here can be extended to dynamic targets, which is the topic of work in progress.

In BN target classification, it is assumed that the typology of the target is related to its features. However, the precise relationships are unknown due to the complexity of the sensors, the targets, and their environment. For example, a frequent pattern is that an object that is large, buried, and cylindrical and is detected under favorable environmental conditions (e.g., wet, clay, uniform, nonmagnetic soil) typically is a mine. However, the existing patterns are not always known *a priori*. Thus, even after the features can be estimated from the measurements, the targets may not be easily classified. Since the target typology influences its actual features, a BN classifier can be built by adding a parent node representing the target's typology to the BN sensor model(s). If target classification involves multiple variables, a set L of parent nodes can be added to $\mathcal{X}$.

For simplicity, let a single variable x_l denote the target typology, with r_l mutually-exclusive classifications. Then, x_l is a parent to all variables in T and to no other variable in $\mathcal{X}$, since the sensor characteristics and operating conditions are independent of x_l, and viceversa. It follows that the BN classifier must have the structure shown in Fig. 3. Its CPTs, $p(T|x_l)$ are learned from $\mathcal{D}$ using the technique described in Section III-A2. If multiple sensors are involved, the BN classifier is modified to account for fusion by connecting the nodes in T to the BN models of all the sensors. With this approach, the patterns between target typology, features, sensor measurements and environmental conditions are established and represented by a BN. Therefore, they can be used to infer the features and the typology of unknown targets, as well as the confidence levels associated with these estimates in the form of probabilities.

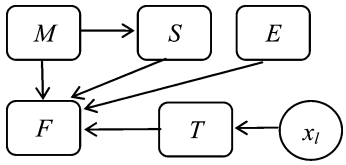

Fig. 3. BN classifier.

IV. Application to Landmine Detection and Classification

The BN modeling approach is used to estimate the features of objects buried underground and to classify them as mines or clutter. Since World War II, numerous conflicts resulted in the planting of millions of landmines that cause about 10 000 casualties each year, in nearly 60 countries [20]. Thanks to the increasing availability of advanced sensor technologies and signal processing software the effectiveness of future demining systems can be greatly enhanced by the development of intelligent sensor fusion and management systems. An *intelligent system* is characterized by the ability to learn over time how to optimize goals in complex, nonlinear environments whose dynamics must ultimately be learned during operation (i.e., on line) [21]. The objective of demining systems is to clear minefields safely, rapidly, and at low cost. These objectives translate into the need for maximizing the probability of detection of buried mines, and for minimizing classification errors, or *false alarms*, where cans or other debris are mistaken for mines.

In a typical demining system GPR, EMI, and IR sensors may be employed to measure the depth d, size z, shape s, and metal content c of mines buried in heterogeneous soils, under various environmental conditions. Because they rely on different operating principles and may function in different modes, these sensors may be more or less effective depending on the environmental conditions. Also, by providing complementary information about the mines, the demining system performance can be significantly improved through sensor fusion.

A. Simulated Demining System

In order to test the methodology on a realistic and controllable test bed a simulation of the landmine sensing system, which includes sensors, targets, soils, and meteorological conditions, is developed. A grid is superimposed on the minefield dividing it into *bins* that contain at most one target (i.e., a mine or clutter). GPR, EMI, and IR measurements of the target features are reproduced and deteriorated according to the operating conditions and surrounding environment. Anti-tank mines (ATM), anti-personnel mines (APM), unexploded ordnance (UXO), and clutter objects (CLUT) have been reproduced based on the Ordata Database [22].

Soil composition (e.g., clay or sand), soil characteristics (e.g., magnetic properties, moisture, uniformity), vegetation, and time-varying meteorological conditions are modeled according to [20], [23], and [24] and can be placed at random or at user-specified positions in the minefield. When a sensor detects an object, the simulation uses actual target features and environmental/meteorological conditions to generate sensor data with random noise and errors that are commensurate to the given situation [see [25] for more details].

TABLE I
LIST OF NODES IN BAYESIAN NETWORK MODELS OF GPR, EMI, AND IR SENSORS

Set:	Node:	Instantiations and corresponding range:
M	GPR mode: m_{GPR}	depth search, resolution search, anti ground-bounce-effect search
	EMI mode: m_{EMI}	high-sensitivity search, low-sensitivity search, shape search
	IR mode: m_{IR}	surface-mine search, shallow-buried-mine search
E	Soil moisture (%): s_r	dry [0-10], wet (10-40], saturated (> 40)
	Soil composition: s_c	very-sandy, sandy, high-clay, clay, silt
	Soil uniformity: s_u	yes, no
	Magnetic soil: s_g	yes, no
	Vegetation: v	no-vegetation, sparse, dense
	Weather: w	clear, overcast, raining
	Illumination	low (7-10 a.m. and 6-9 p.m.), medium (10-1 p.m.), high (1-6 p.m.)
T	Depth (cm): d	surface [0], shallow-buried (0-12], buried (12-60], deep-buried (> 60)
	Size (cm): z	small (2-13], medium (13-24], large (24-40], extra-large (> 40)
	Shape: s	cylinder, box, sphere, long-slender, irregular
	Metal content (gr): c	no-metal [0-3], low-metal (3-200], high-metal (> 200)

TABLE II
COMPOSITION OF DATABASE USED TO TRAIN THE BAYESIAN NETWORK SENSOR MODELS AND CLASSIFIER

Variable:	Number of training cases containing each instantiation:
Target category:	APM 1320, ATM 2456, UXO 1224, CLUT 3000
Soil moisture:	dry 2339, wet 4555, saturated 1606
Soil composition:	very-sandy 1743, sandy 1724, high-clay 1688, clay 1647, silt 1698
Magnetic soil:	yes 1971, no 6529
Vegetation:	no-vegetation 2796, sparse 2756, dense 2948
Depth:	surface 3136, shallow-buried 2981, buried 1083, deep-buried 800
Size:	small 3102, medium 2510, large 2316, extra-large 72
Shape:	cylinder 3676, box 688, sphere 1320, long-slender 1224, irregular 1320
Metal content:	no-metal 2184, low-metal 4129, high-metal 1687

B. BN Modeling of GPR, EMI, and IR Sensors

The approach described in Section III-A is used to obtain BN models of three sensor types: GPR, EMI, and IR. The set of nodes $\mathcal{X}$ and the initial structure $\mathcal{A}_0$ of each model are determined based on expert knowledge of the sensors and their operating principles. The measured features, corresponding to the target characteristics $T = \{d, z, s, c\}$, are denoted by $F = \{d_m, z_m, s_m, c_m\}$. In this case, the chosen variables in T and F have the same instantiations, as listed in Table I. However, in general, there need not be a direct correspondence between nodes in T and F. The following three sections describe how the remaining BN nodes (Table I) and initial structure are determined for each sensor model.

Following the selection of the nodes and instantiations, the BN structure is refined and the CPTs are learned from a database $\mathcal{D}$ of sensor measurements. In this application, each case in $\mathcal{D}$ contains the measurements that are collected from one of the sensors (GPR, EMI, or IR) from a target with known features (T) and typology (x_l), under known environmental conditions (E), and sensor parameters (M); thus, $\mathcal{D} = \{\mathcal{D}_{\mathrm{GPR}}, \mathcal{D}_{\mathrm{EMI}}, \mathcal{D}_{\mathrm{IR}}\}$. $\mathcal{D}$ is referred to as *training set*, and, here, it is generated through the demining simulation by uniform sampling of the Ordata Database [22], which contains over 5000 explosive items. In this research, $\mathcal{D}$ contains 5000 mines and its composition, described in Table II, is consistent with that found in this existing database. Also, 3000 clutter objects are reproduced to emulate metallic debris, cans, and plastic objects of regular shapes that resemble anti-personnel mines. All of these targets (Table II) are buried in a simulated minefield whose environmental conditions are sampled uniformly from 660 possible configurations. Once sampled, each environmental configuration is placed randomly in zones that are two- to eight-bin wide, resulting in the distribution shown in Table II. Each simulated demining sensor is used to measure the features of these targets, under the assigned conditions and using everyone of its modes, finally producing the data $\mathcal{D}_{\mathrm{GPR}}$, $\mathcal{D}_{\mathrm{EMI}}$, and $\mathcal{D}_{\mathrm{IR}}$. The influence of the training-set size on the BN model performance is discussed in Section V-A. The minimum number of cases required to model the GPR, EMI, and IR sensors is found to be approximately 2000. When the training data is limited, proper BN learning requires a modified K2 algorithm that exploits *a priori* conditional independencies between the nodes to reduce the search space of feasible structures (e.g., [27]).

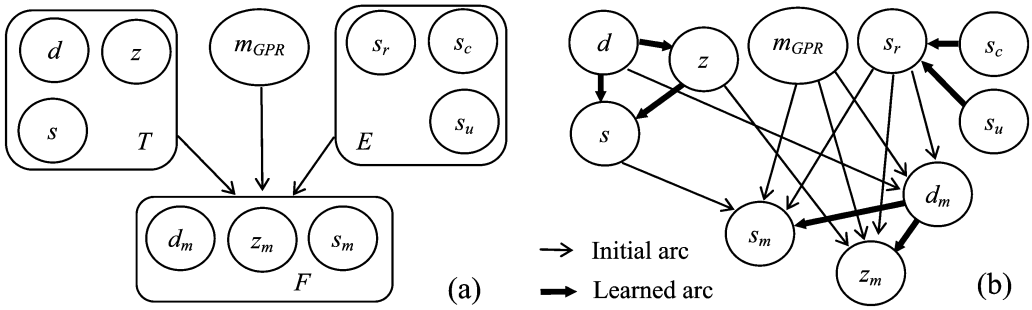

Fig. 4. (a) Initial BN structure based on GPR variable sets (b) and final GPR BN structure learned from data.

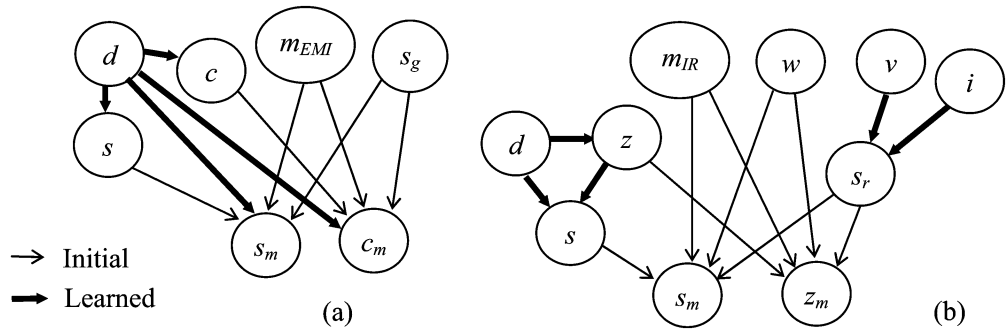

Fig. 5. Initial and final BN structure of EMI (a) and IR (b) sensor models.

1) GPR: GPR sensors emit radio waves that penetrate the ground and process their reflections at the boundaries of materials characterized by different refraction indexes. An image of underground vertical slices and of any objects buried within is obtained by sensing discontinuities in electrical properties. The measured object size, shape, and depth can be obtained through signal processing techniques [28], such as edge extraction. The frequency of the radio wave and its bandwidth are the sensor parameters that determine the search mode m_{GPR}, and must be tuned to achieve the best trade-off between penetration depth, which increases at lower frequencies, and image resolution, which improves at higher frequencies. Very high frequencies may be required to overcome the so-called ground-bounce effect (GBE), which may cause shallow-buried mines to be missed because of ground discontinuities [20]. The influence of soil properties, such as, uniformity, moisture, and composition, on the sensor performance is learned from the GPR data, $\mathcal{D}_{\mathrm{GPR}}$.

The relevant GPR nodes $(\mathcal{X}_{\mathrm{GPR}})$ and initial structure illustrated in Fig. 4 are chosen based on prior expert knowledge [29] using the architecture in Fig. 2. All nodes and instantiations are selected according to the procedure in Section III-A1 and are defined in Table I. Then, BN learning algorithms are employed to determine the additional arcs [shown in bold in Fig. 4(b)] and the CPTs that maximize the scoring metric in (3) for the given sensor data $(\mathcal{D}_{\mathrm{GPR}})$. From the refined structure it can be seen that additional relationships between the target features, measurements, and environmental variables are revealed by the sensor data. Also, the results presented in Section V show that this structure estimates target features more accurately then the one based solely on expert knowledge.

2) EMI: EMI sensors detect metal objects by measuring the secondary magnetic field produced by eddy currents that are induced by a time-varying magnetic field [20]. The primary field is produced by an electrical current flowing in a transmit coil of wire, and the secondary field is sensed by measuring the voltage induced in a receiving coil. Detection of metal content is declared when the measured voltage exceeds a chosen threshold that is determined by whether the sensor is in a high or low sensitivity mode. The main disadvantage of EMI sensors is their inability to detect plastic mines and mines buried in magnetic soils. However, if their measurements are augmented by other sensors that do not rely on metal content, such as GPR, the overall system performance can greatly benefit from their employment.

The EMI sensor mode is represented by the variable m_{EMI}. The other relevant nodes $(\mathcal{X}_{\mathrm{EMI}})$ are determined from expert knowledge of EMI detectors, such as, the Geonics EM-61 and the AN/PSS-12 [30], as defined in Table I. The initial arcs [Fig. 5(a)] also are determined from these studies, and the final BN model is learned from EMI sensor data, $\mathcal{D}_{\mathrm{EMI}}$, as shown in Fig. 5(a). This model is then used for inference, fusion, and classification, as demonstrated in Section V.

3) IR: IR sensors detect anomalies in IR radiation that is either emitted by mines, soil, or vegetation. Based on the location of the sensor, the radiation data can be processed to build an image of an horizontal area and to estimate the depth of the object therein. Images can be obtained for depths up to 12 cm. Therefore, IR sensors only can be used to obtain the size and shape of surface or shallow-buried objects. Because they rely on temperature variations, their performance also is highly influenced by environmental conditions, such as, time, weather, vegetation, and soil properties.

The mode of IR sensors m_{IR} influences the measured target features (F) and is uniquely determined by its height above the ground. Hence, the height needs not be included as a sensor state

TABLE III
PERFORMANCE OF STRUCTURAL LEARNING ALGORITHMS FOR GPR BN MODEL

Algorithm:	Initial structure	Structural EM	Expert-order K2	*E*-order K2	*F*-order K2
$\bar{J}_d$:	0.1588	0.1706	0.1728	0.1721	0.1684
$\bar{J}_z$:	0.1529	0.1338	0.1346	0.1360	0.1360
$\bar{J}_s$:	0.3669	0.2507	0.2434	0.2515	0.2537

variable in the model. Based on the IR working principles and detailed studies of Agema Thermovision 900 sensors [16], [24], [29], the relevant variables (Table I) and initial dependencies are determined, as shown in Fig. 5(b). As in the case of GPR and EMI sensors, the initial structure is refined and the CPTs are learned from the simulated IR sensor data, $\mathcal{D}_{\mathrm{IR}}$, by means of BN learning algorithms. The final BN model is shown in Fig. 5(b), where the arcs learned from the data are highlighted in bold. For this sensor, the data reveals that the vegetation and illumination only influence the measurements through the soil-moisture variable.

Section V-A describes a comparative study of BN learning algorithms applied to GPR sensor modeling. Based on this study, a BN model is obtained for each sensor type (Figs. 4–5) using the K2 structural learning algorithm and the EM parameter learning algorithm. These three models are then used to estimate target features based on single-sensor measurements (Section V-B), and to perform GPR/IR and GPR/IR/EMI sensor fusion (Section V-C) and target classification (Section V-D).

V. RESULTS

The demining system simulation (Section IV-A) is used to produce both the training set $\mathcal{D}$, described in Section IV-B, and a *validation set*, $\mathcal{V}$, of cases with sensor data. The validation set contains cases that are not used for training, i.e., $\mathcal{D} \cap \mathcal{V} = \emptyset$, and is used to test the performance of the BN model under new conditions, such as those that might be encountered during actual operation. Similarly, when a database is available from experiments it can be partitioned into the two sets $\mathcal{D}$ and $\mathcal{V}$ that are used to train and test the BN models, respectively. Suppose the BN model of a sensor is used to infer a variable $x_f \in \mathcal{X}$ given evidence from a case in the validation set, e.g., $\mathcal{C}_v \in \mathcal{V}$, where, $\mathcal{C}_v = \{x_i^v \mid x_i = x_i^v, \forall\, x_i \in \mathcal{X}\}$. Then, the true value of the inferred variable, x_f^v, is known from $\mathcal{C}_v$. Inference in the BN model provides not only the most likely value of x_f, but also a probability distribution over all of its possible state values, i.e., $p\left(x_f^m \mid e\right)$, for $m = 1, \ldots, r_f$.

An error metric that takes into account the BN confidence level is obtained by weighting the distance from the true value of the inferred variable by the respective probability

$$J_f = \mathbf{p}_f \cdot \mathbf{g}_f \tag{7}$$

$\mathbf{p}_f$ is a $1 \times r_f$ vector of probabilities obtained by the inference algorithm: $\mathbf{p}_f \equiv \left[p\left(x_f^1 \mid e\right) \cdots p\left(x_f^{r_f} \mid e\right)\right]$. $\mathbf{g}_f$ is a $1 \times r_f$ vector containing the distance g_f between the true value of x_f (i.e., x_f^v) and all of its possible instantiations: $\mathbf{g}_f \equiv \left[g_f\left(x_f^1, x_f^v\right) \cdots g_f\left(x_f^{r_f}, x_f^v\right)\right]$. The distance g_f is a discrete metric that is feature specific and satisfies the following conditions: $J_f = 0.5$, when $p(x_f \mid e)$ is uniform, and $J_f = 1$ is the maximum value of (7) [25]. When the same feature or variable x_f is inferred for t targets, an average error metric $\overline{J}_f = J_f/t$ is used to evaluate performance.

The validation minefield $\mathcal{V}$ used to verify the BN approach developed in the previous sections is divided into 10×12 bins and contains 65 mines and 30 clutter-objects with position and features assigned randomly (as explained in Section IV-A). The information associated with each bin constitutes a case $\mathcal{C}_v$ in $\mathcal{V}$. The field's surface is 30% magnetic, 46% saturated, and 50% covered in dense vegetation, thus presenting unfavorable conditions for EMI, GPR, and IR sensors, respectively. All of the sensor measurements are collected while the field is in clear weather and medium illumination. Together with the sensor mode, each bin's environmental conditions and measurements constitute the evidence used for inference in the BN sensor models.

A. Learning the BN Sensor Models

The performance of two structural learning algorithms is investigated using the error metric defined above. After selecting the GPR BN nodes and initial structure (Section IV), the structural K2 and EM algorithms are used to refine the structure based on $\mathcal{D}_{\mathrm{GPR}}$. Both the K2 and the EM structures are tested by performing inference on resolution-search GPR measurements collected from the validation field $\mathcal{V}$, leading to the error metrics in Table III for the mine features d, z, and s. The K2 algorithm allows the user to specify the initial order of the BN nodes, from parents to children. Table III shows the performance of BN structures obtained by choosing E and F as the first parents (E order and F order, respectively), and by using an expert-based order that is obtained from Fig. 4. As can be expected, using expert knowledge to initialize the learning procedure not only improves the BN performance, but also produces a BN structure that better reflects the sensor operating principles and has lower complexity (i.e., a smaller number of arcs). Although Table III shows a small improvement with respect to EM learning, the K2 algorithm is considered to be better suited for this application because it incorporates expert knowledge and displays lower computational times. Similar results are obtained by learning EMI and IR sensor model structures. Hence, the K2 algorithm is used to determine the final structures of the three BN sensor models (shown in Figs. 4–5).

The EM parameter learning algorithm is used to learn the BN parameters, or CPTs, from the database of sensor data, $\mathcal{D}$, described in Section IV-B. The size of this training set is chosen to be commensurate with the size of the on-line Ordata Database [22]. Hence, $\mathcal{D}$ contains GPR, EMI, and IR sensor measurements and environmental information for 5000 mines and

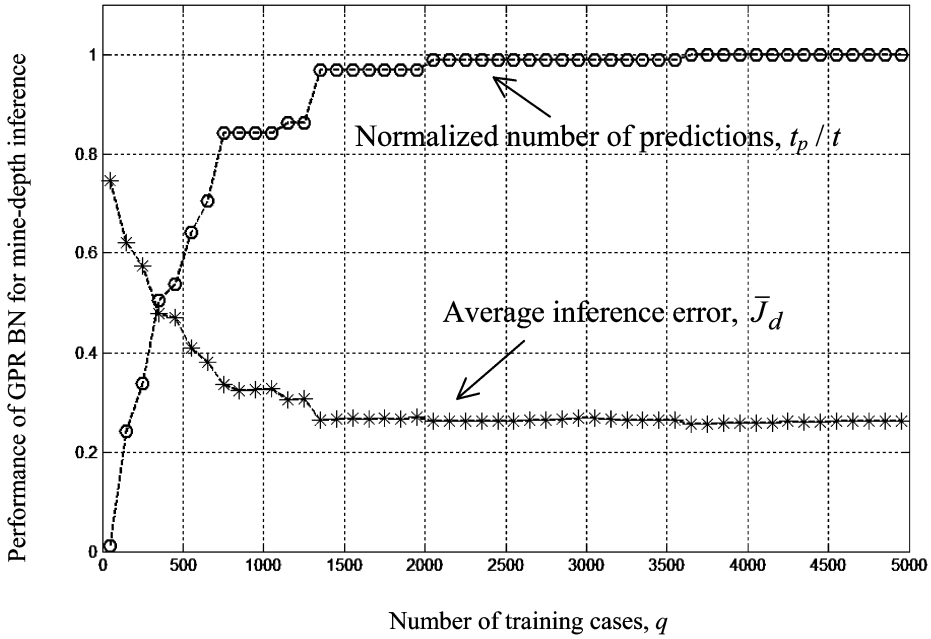

Fig. 6. Effect of training-set size over the inference performance of the GPR BN model learned.

3000 clutter objects. If a training set is inadequate in size and in variation of significant variables, the learned BN model will display poor inference performance over the validation set. That is, the inference algorithm may be unable to provide a prediction for some of the cases in $\mathcal{V}$. Instead, a set of zero marginal probabilities is produced for one or more of the inferred features indicating that the algorithm failed. Also, the predictive accuracy, which is evaluated through the error metric $\bar{J}_f$, may not be satisfactory for the sensor application, indicating that more training cases may be needed to learn a better structure and CPTs.

These conditions are illustrated by learning the CPTs of the GPR sensor model [Fig. 4(b)] using training sets of different size that are sampled uniformly from $\mathcal{D}_{\mathrm{GPR}}$. The resulting error metric $\bar{J}_d$ and number of predictions t_p obtained by testing these BNs over $\mathcal{V}$ are plotted with respect to the number of training cases, q, in Fig. 6. As the the size of the training set increases, the number of predictions approaches the total number of targets $(t = 95)$, such that if an adequate value of q is used the BN always is capable of providing an estimate for the target features and $t_p = t$. The improved accuracy of these estimates is reflected in the reduction of the error $\bar{J}_d$ with increasing q. When the BN is unable to make a prediction, the bin is assigned an error $J_d = 0.75$, for convention, because this case is considered to be worse than one with a uniform probability distribution $(J_d = 0.5)$, but it is better than a case where the BN provided an incorrect estimate with a high confidence level $(J_d = 1)$. Initially, when q is very small and almost no predictions are made, $\bar{J}_d \approx 0.75$. As $t_p \rightarrow t$, and the size of the training set becomes adequate, the average error $\bar{J}_d$ represents the accuracy of the predictions provided by the BN. Eventually, this error approaches an asymptote, which here is $\bar{J}_d \approx 0.2$, such that beyond a certain value of q the BN accuracy can no longer be improved by increasing the size of the training set. Therefore, if this level of performance is unsatisfactory, the BN sensor model must be improved by changing one or more of the following: the set of nodes $\mathcal{X}$, the instantiations of one or more nodes, e.g., $\{x_i : x_i^1, \ldots, x_i^{r_i}\}$, the structure $\mathcal{A}$, or the composition of the training set (see [25] for additional experiments). From Fig. 6 it can be concluded that, for the GPR BN model, $q \geq 2000$ provides acceptable performance, while $q \geq 3600$ is required for maximum performance. Thus, if the sensor data available is limited to fewer cases, a different learning algorithm may be required to increase the number and accuracy of the predictions without changing the value of q, as shown in [27].

B. Inference From Single-Sensor Measurements

The BN model of a sensor can be used to infer the features of an unknown target from the measurements of a single sensor, given its mode and environmental conditions. This procedure is illustrated by simulating a GPR in resolution-search mode that measures the features of the objects buried in $\mathcal{V}$. Using the GPR BN model in Fig. 4(b), the size, shape, and depth of these objects are inferred and the corresponding confidence levels are computed. In Fig. 7(b), the inferred depth of each target and its probability are plotted with respect to their location in the minefield. In this figure, the bins are numbered in the x and y directions, and the depth value is represented by the chromatic scale in the legend, with white indicating an empty bin. These inference results can be compared to the actual mine depth [Fig. 7(a)], which is available from each validation case $\mathcal{C}_v \in \mathcal{V}$. The corresponding error metric J_d, computed from (7), is plotted in Fig. 7(c). The error metric defined in (7) varies between 0 and 1, where $J_f = 1$ represents an incorrect estimate that is provided by the BN with a high confidence level. $J_f = 0$ represents a perfect estimate, and $J_f = 0.5$ corresponds to a feature inferred with a uniform probability distribution, for which any of its possible values are equally likely. Since the prior probabilities of the mine features are uniform, if $J_f < 0.5$ the estimated feature has value, and if $J_f < 0.2$ the estimated feature is close to the real one. Sensor failure indicates that an object was not detected due

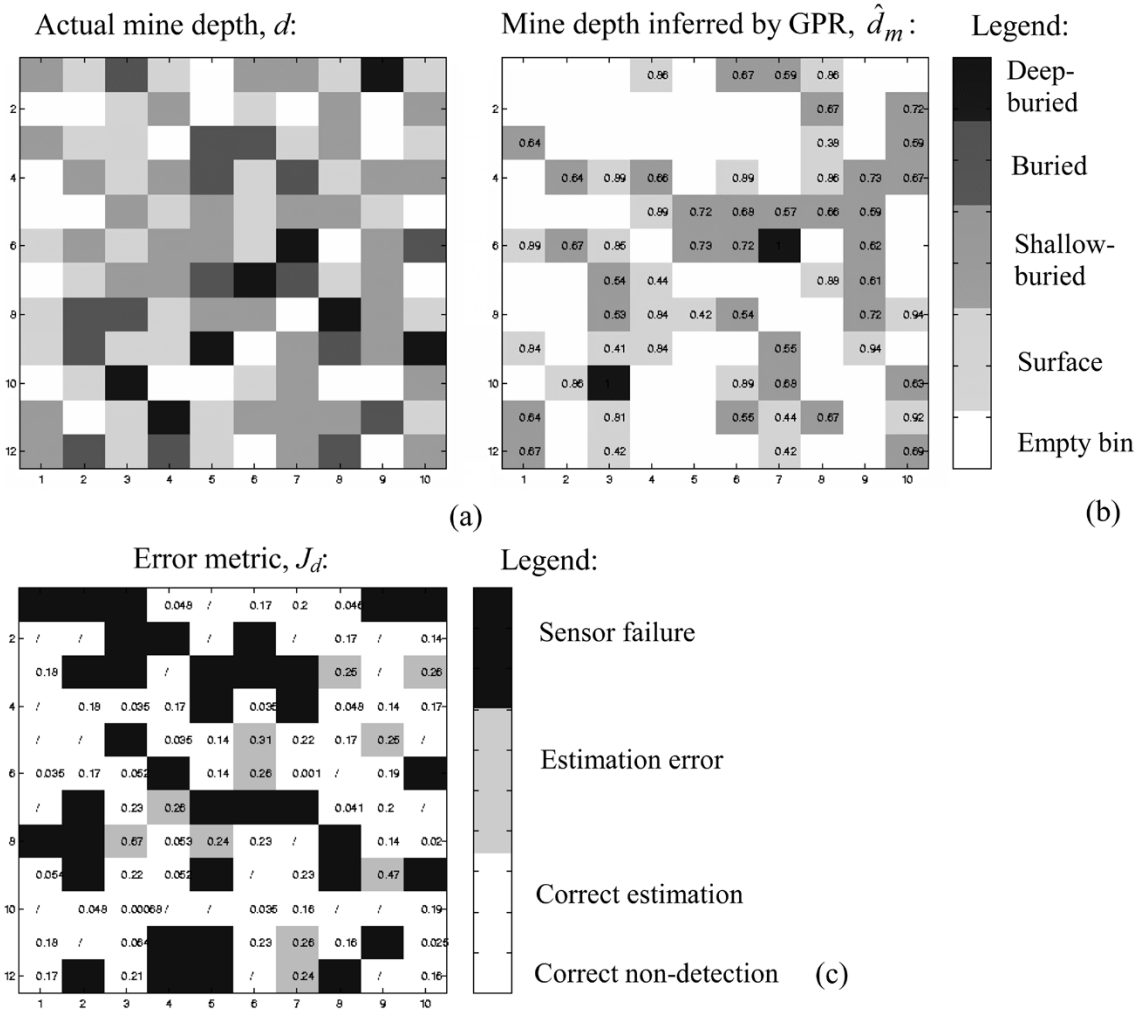

Fig. 7. (b) Mine depth inferred by GPR BN model based on the measurements performed in resolution-search mode on the mine field in (a). (c) Respective error distribution.

TABLE IV
AVERAGE ERROR METRIC FOR MINE FEATURES INFERRED BY THE GPR BN MODEL

GPR Mode:	$\bar{J}_d$:	$\bar{J}_z$:	$\bar{J}_s$:
Depth search	0.18792	0.16840	0.31245
Resolution search	0.15945	0.09220	0.29483
Anti GBE search	0.16793	0.09134	0.27247

TABLE V
AVERAGE ERROR METRIC FOR MINE FEATURES INFERRED BY THE EMI BN MODEL

EMI Mode:	$\bar{J}_c$:	$\bar{J}_s$:
High-sensitivity search	0.39444	n.a.
Low-sensitivity search	0.38617	n.a.
Shape search	0.38393	0.34580

TABLE VI
AVERAGE ERROR METRIC FOR MINE FEATURES INFERRED BY THE IR BN MODEL

IR Mode:	$\bar{J}_z$:	$\bar{J}_s$:
Surface-mine search	0.19345	0.40312
Shallow-buried-mine search	0.08307	0.30590

to the simulated GBE, a physical limitation of GPR sensors that can only be overcome through sensor fusion.

Through a similar process, the performance of the GPR, EMI, and IR BN models without fusion is tested over $\mathcal{V}$ using different sensor modes. The average error metrics are listed in Tables IV–VI according to the mode and to the features measured by each sensor. These results are consistent with the expert knowledge of demining sensor systems. For instance, shape is more difficult to estimate than depth or size, causing $\bar{J}_s$ to be greater than $\bar{J}_d$ and $\bar{J}_z$, except when the shape-search mode is used. Also, the fact that EMI measurements generally are less reliable than those of GPR and IR is reflected in the higher error metrics that are obtained from the EMI BN model, as shown in Table V. However, EMI sensors may still be useful to complement other measurements, as demonstrated by the fusion results described in the next section. Additional numerical experiments [25] show that by accounting for environmental and operating conditions in processing the measurements of a single sensor

TABLE VII
SIZE AND SHAPE ERROR METRICS BEFORE AND AFTER FUSION OF GPR, EMI, AND IR DATA

	Sensor Fusion (sensors used)	**GPR Only (fusion improvement)**	**EMI Only (fusion improvement)**	**IR Only (fusion improvement)**
$\bar{J}_z$:	0.02769 (GPR/IR)	0.07898 (64.9%))	n.a.	0.08307 (66.7%)
$\bar{J}_s$:	0.21072 (GPR/EMI/IR)	0.36891 (42.8%)	0.29627 (28.9%)	0.31469 (33.0%)

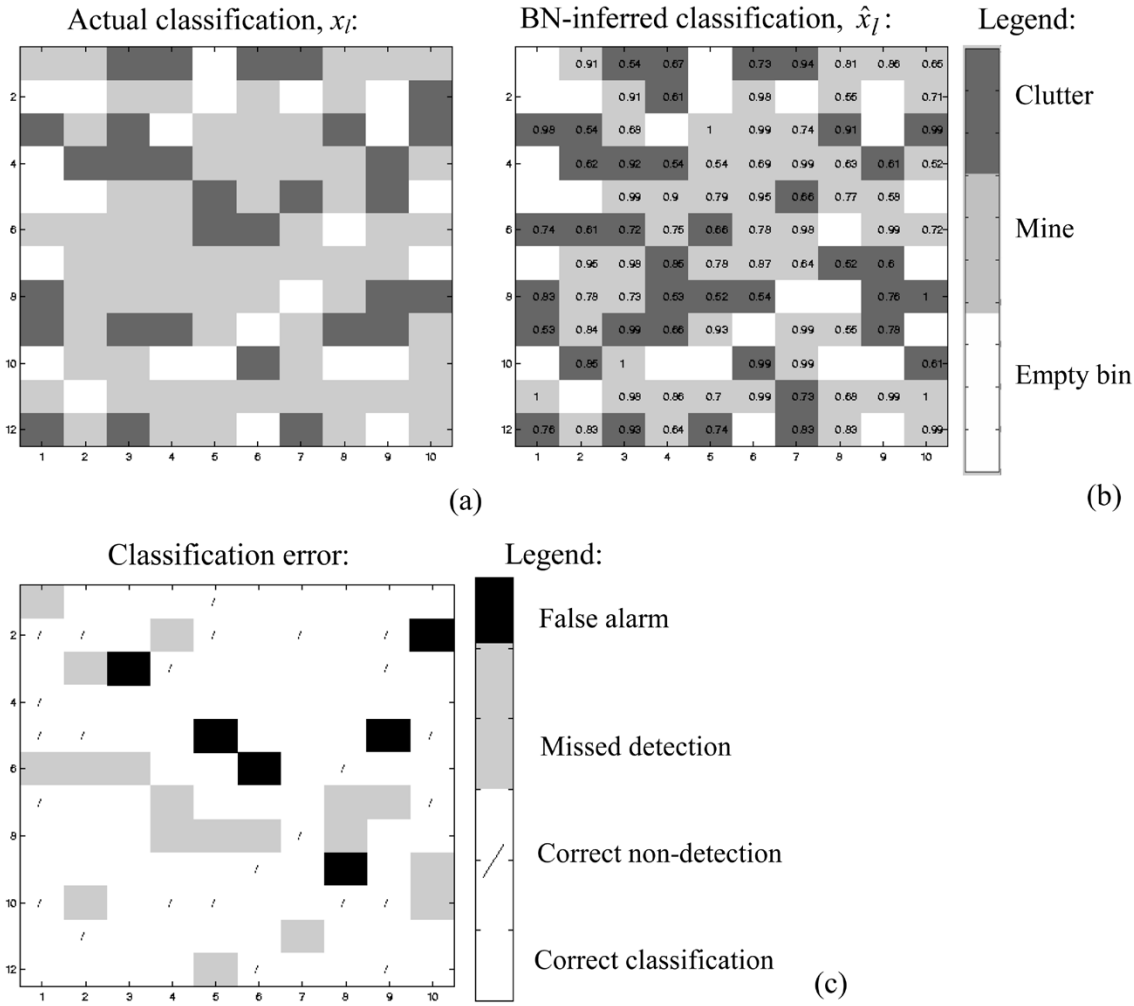

Fig. 8. (b) Object type inferred by BN classifier based on GPR, EMI, and IR measurements obtained from the mine field in (a). Respective error distribution (c).

without fusion, the number of targets whose features were estimated correctly increased on average by 10%. A much greater improvement is obtained by using the BN models for sensor fusion, as demonstrated in Section V-C.

C. Heterogeneous Sensor Fusion

BN feature-level fusion (Section III-C) is used to obtain a probability distribution for the target features based on the evidence from all three demining sensors and from known environmental conditions, i.e., $p(T \mid e_{\mathrm{GPR}}, e_{\mathrm{EMI}}, e_{\mathrm{IR}})$. By fusing GPR and EMI target-size measurements over the validation field $\mathcal{V}$ the average error metric is improved by 64.9% with respect to using only GPR measurements, and by 66.7% with respect to using only EMI measurements (Table VII). Fusion of GPR, EMI, and IR shape measurements also improves the average error metric with respect to using any of these three sensors alone without fusion. For example, when the GPR measurements are fused with those of the EMI and IR sensors, the error in estimating the shape is decreased by 42.8% (Table VII).

The results also show that BN fusion drastically reduces sensor failure by exploiting the complementarity of the sensors. For instance, failure to estimate the targets shape in $\mathcal{V}$ occurs only for five of the 95 objects after BN fusion is performed. Instead, when GPR, EMI, and IR measurements are processed individually failure to estimate the target shape occurs for 19, 40, and 64 objects, respectively. The reason is that EMI sensors

are only able to detect metal objects, and IR sensors are only able to detect surface or shallow-buried objects. However, GPR sensors also can detect plastic, buried, and deep-buried objects (Fig. 7) and with the availability of EMI and IR measurements the GPR estimate of size and shape becomes more accurate, as shown in Table VII.

D. Target Classification

The BN classifier developed in Section III-D and illustrated in Fig. 3 is used to infer the typology of each object in $\mathcal{V}$, based on the features estimated from the fused GPR, EMI, and IR measurements. The objects' actual classification is shown in Fig. 8(a), where a grey scale is used to represent clutter, mines, and empty bins. According to the technique presented in Section III-D, the fused probability distribution of the target features, $p(T \mid e_{\mathrm{GPR}}, e_{\mathrm{EMI}}, e_{\mathrm{IR}})$, can be used as soft evidence in the BN classifier to infer the target typology, by computing $p(x_l \mid e_{\mathrm{GPR}}, e_{\mathrm{EMI}}, e_{\mathrm{IR}})$. The most probable typology and its corresponding confidence level are obtained from this distribution and are plotted in Fig. 8(b). In this case, among the 95 objects buried in the validation field $(\mathcal{V})$ six clutter bins were erroneously classified as mines (false alarms), and 18 mines were missed (missed detections): three due to sensor failure and 15 due to incorrect classification, as illustrated in Fig. 8(c). These results are representative of extensive numerical simulations carried out in [25]. For validation fields with challenging environmental conditions (such as $\mathcal{V}$), the average accuracy of the BN classifier is approximately 75%. As can be expected, when milder conditions are simulated, both the classification and the feature-inference errors are significantly decreased on average.

VI. Conclusion

A novel, unified BN approach to sensor modeling, fusion, and target classification is developed and demonstrated on a demining application. The approach uses *a priori* expert knowledge of the sensor's operating principles and available databases of actual sensor data to build a probabilistic model of the measurement process. The demining sensor system investigated is comprised of GPR, EMI, and IR sensors that measure the shape, size, depth, and metal content of objects buried in a field with heterogeneous soil and environmental conditions. The results show that the BN models are capable of inferring target features by systematically taking into account single or fused sensor measurements and known environmental conditions. The BN modeling approach is combined with the D-S fusion technique in order to exploit the complementarity of the sensors. The sensor models also are used to develop a BN classifier that estimates the target typology based on the inferred features and its environment. By systematically combining heterogeneous sensor measurements, the accuracy of the feature estimates is improved by up to 64% and the number of objects detected and classified by the demining system is increased by up to 62%, with respect to single-sensor measurements. Also, the tools developed in this paper allow the operator or sensor manager to make informed decisions pertaining the quality of target estimation and classification, by computing corresponding confidence levels that are based on the sensor's physical limitations and operational environment.

References

[1] E. Waltz and J. Llinas, *Multisensor Data Fusion*. Norwood, MA: Artech House, 1990.

[2] R. R. Brooks and S. S. Iyengar, *Multi-Sensor Fusion: Fundamentals and Applications With Software*. Englewood Cliffs, NJ: Prentice-Hall, 1998.

[3] F. Cremer, K. Schutte, J. G. M. Schavemaker, and E. den Breejen, "Toward on operational sensor-fusion system for anti-personnel landmine detection," presented at the SPIE Detection and Remediation Technologies for Mines and Minelike Targets V Conf., vol. 4038, A. C. Dubey, J. F. Harvey, J. T. Broach, and R. E. Dugan, Eds., 2000.

[4] P. B. W. Schwering, "Model based feature fusion approach," presented at the SPIE Detection and Remediation Technologies for Mines and Minelike Targets IV Conf., vol. 4394, A. C. Dubey, J. F. Harvey, J. T. Broach, and V. George, Eds., 2001.

[5] A. P. Dempster, N. M. Laird, and D. B. Rubin, "A generalization of bayesian inference," *J. Roy. Stat. Soc. B*, vol. 39, pp. 1–39, 1968.

[6] G. Shafer, *A Mathematical Theory of Evidence*. Princeton, NJ: Princeton Univ. Press, 1976.

[7] F. V. Jensen, *Bayesian Networks and Decision Graphs*. New York: Springer-Verlag, 2001.

[8] G. F. Cooper and E. Herskovits, "A bayesian method for the induction of probabilistic networks from data," *Mach. Learn.*, vol. 9, pp. 309–347, 1992.

[9] R. E. Bellman, *Dynamic Programming*. Princeton, NJ: Princeton Univ. Press, 1957.

[10] R. Shachter and C. Kenley, "Gaussian influence diagrams," *Management Sci.*, vol. 35, pp. 527–50, 1989.

[11] S. L. Lauritzen, "Propagation of probabilities, means and variances in mixed graphical association models," *J. Amer. Stat. Assoc.*, vol. 87, pp. 1098–108, 1992.

[12] W. L. Buntine, "Theory refinement on Bayesian networks," in *Proc. 7th Conf. Uncertainty in AI*, 1991, pp. 52–60.

[13] D. Heckerman, D. Geiger, and D. M. Chickering, "Learning Bayesian networks: The combination of knowledge and statistical data," *Mach. Learn.*, vol. 20, pp. 197–243, 1995.

[14] S. Chang and M. F. Ruane, "Extraction of ground penetrating radar for mine detection," *Proc. SPIE*, vol. 5089, no. 2, pp. 1201–1209, 2003.

[15] T. G. Savelyev, L. Van Kempen, and H. Sahli, "GPR anti-personnel mine detection: Improved deconvolution and time-frequency feature extraction," *Proc. SPIE*, vol. 5046, pp. 232–241, 2003.

[16] W. Messelink *et al.*, "Feature-based detection of landmines in infrared images," *Proc. SPIE*, vol. 4742, no. I, pp. 108–119, 2002.

[17] [Online]. Available: http://www.ai.mit.edu/murphyk/Software/BNT/bnt.html

[18] (1998) *Learning in Graphical Models* [Online]

[19] A. P. Dempster, "Upper and lower probabilities induced by a multivalued mapping," *Ann. Math. Stat.*, vol. 38, pp. 325–339, 1967.

[20] J. MacDonald *et al.*, *Alternatives for Landmine Detection*. Stokie, IL: Rand McNally, 2003.

[21] D. A. White and D. A. Sofge, *Handbook of Intelligent Control: Neural, Fuzzy, and Adaptive Approaches*. New York: Van Nostrand Reinhold, 1992.

[22] [Online]. Available: http://maic.jmu.edu/ordata/mission.asp

[23] L. R. Pasion, S. D. Billings, and D. W. Oldenburg, "Evaluating the effects of magnetic soils on TEM measurements for UXO detection," presented at the 72nd Annu. Meeting Soc. Exploration Geophysicists, Salt Lake City, UT, 2002.

[24] R. L. Van Dam *et al.*, "Soil effects on thermal signatures of buried nonmetallic landmines, detection and remediation technologies for mines and minelike targets VIII," *Proc. SPIE*, vol. 5089, pp. 1210–1218, 2003.

[25] A. Vaghi, "Sensor Management by a Graphical Model Approach," Laurea thesis, Dept. Elect. Eng., Politecnico di Milano, Milano, Italy, 2004.

[26] N. Friedman, "Learning belief networks in the presence of missing values and hidden variables," in *Proc. 14th Int. Conf. Machine Learning*, San Francisco, CA, 1997, pp. 125–133.

[27] K. C. Baumbgartner, S. Ferrari, and G. C. Salfati, "Bayesian network modeling of criminal behavior for criminal profiling," in *Proc. Conf. Decision Control*, 2005, pp. 2702–2709.

[28] J. K. Paik *et al.*, "Image processing-based mine detection techniques using multiple sensors: A review," *Subsurf. Sens. Technol. Appl., Int. J.*, vol. 3, no. 3, pp. 203–252, 2002.
[29] "Final Technical Report Forward Looking Infrared Detector and Radar Antenna Array Optimization—Vehical Mounted Mine Detection Systems," Geo-centers, 1997.
[30] Genomics EM 61 Specifics. Geo-centers. [Online]. Available: http://www.geonics.com
[31] J. K. Paik *et al.*, "Image processing-based mine detection techniques using multiple sensors: A review," *Subsurf. Sens. Technol. Appl., Int. J.*, vol. 3, pp. 203–252, 2002.
[32] T. Miller *et al.*, "Effects of soil moisture on landmine detection using ground penetrating radar," in *Proc. SPIE Detection and Remediation Technologies for Mines and Minelike Targets VII Conf.*, vol. 4742, 2002, pp. 281–290.

Silvia Ferrari (M'01) received the B.S. degree from Embry-Riddle Aeronautical University, Bunnell, FL, and the M.A. and Ph.D. degrees from Princeton University, Princeton, NJ.

She is Assistant Professor of mechanical engineering and materials science at Duke University, Durham, NC, where she directs the Laboratory for Intelligent Systems and Controls. Her principal research interests are robust adaptive control of aircraft, learning in network models, approximate dynamic programming, and distributed sensor management.

Dr. Ferrari is a member of the ASME and AIAA. She is the recipient of the NC Space Grant Consortium Research Seed Award (2003), the ONR Young Investigator Award (2004), and of the NSF CAREER Award (2005).

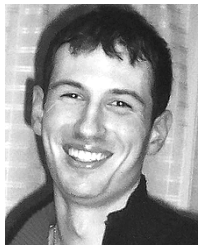

Alberto Vaghi received the M.S. degree in electrical engineering and information science from the Politecnico di Milano, Milano, Italy, and the Diplôme d'Ingénieur from the Supélec, Paris, France, in 2004.

From 2003 to 2004, he was a Visiting Researcher in the Laboratory for Intelligent Systems and Controls, Duke University, Durham, NC.

Mr. Vaghi participated in the Top Industrial Managers for Europe (TIME) program in 2003. He also received the ISU scholarships (1999 to 2004), and the Special ISU Award (2002) for students enrolled in the TIME program.