Available online at www.sciencedirect.com
ScienceDirect
Information Fusion 7 (2006) 342–345
INFORMATION FUSION
www.elsevier.com/locate/inffus

Guest editorial

FUSION 2004: The seventh international conference on information fusion

This special issue of the journal Information Fusion consists of a few selected, substantially extended, and thoroughly revised papers from among the many that were originally presented at FUSION 2004, the Seventh International Conference on Information Fusion held 28 June–1 July 2004 in Stockholm, Sweden (http://www.fusion2004.org). This series of conferences was started in 1998, almost coincident with the founding of their sponsoring professional society, the International Society of Information Fusion (http://www.isif.org). The FUSION 2004 conference received 279 original submissions. Each submission was reviewed by at least three reviewers and 171 papers were accepted for the conference.

From a set of the highest ranked papers that came out of the conference review process, the guest editors selected those papers they found most interesting and innovative and invited their authors to submit an extended version for regular review with the journal. Each article was reviewed by three reviewers according to the journal's standards. Based on this rigorous journal peer-review process, all the papers were further revised to satisfactorily address the reviewer comments. One of the papers is co-authored by the guest editors and was handled by the Editor-in-Chief. Eventually, thirteen papers were accepted for this special issue.

What does emerge from these accounts is that information fusion, after a few years as an emerging research area and community, today is able to present a wide range of not only interesting but also quite mature applications and a no less broad palette of relevant methodologies. Classical methods, such as Kalman filtering, are finding new useful variants and applications (such as evidenced in the papers by Bilenne and Koch) and new methodological approaches claim their grounds (as shown, e.g., by the paper by Thorsen and Oxley on the use of category theory in information fusion performance assessment). Bridges are built across previously non-negotiable methodological gorges, such as shown by the paper by Smets and Ristic, where the problem of joint tracking and classification is successfully treated using a new methodological combination of belief function theory and Kalman filtering. In addition, solid cross-disciplinary links between information fusion and robotics (see, e.g., the paper by Makarenko and Durrant-Whyte) are being established, as between information fusion and database technology (Chang, Jungert and Li). Doherty, Łukaszewicz and Szałas reveal another somewhat unexpected link in their paper, which uses rough sets to manage uncertain information in a complex real-world application, achieving intelligent behaviour from an unmanned vehicle required to maneuver in a complex environment.

The special issue begins with the paper by Thorsen and Oxley. An organisation that is building a fusion system to detect or classify objects will want to get the best possible result for the money spent. Receiver operating characteristics (ROCs) can be developed for such systems. This paper proposes a functional defined on ROC curves as a method of quantifying the performance of a classification system. This functional then allows for the development of a cogent definition of what is fusion (i.e., results of fusion rules in general) and what the paper terms *fusors*, a subcategory of fusion rules which rely on qualitative differences between the fused and non-fused information. It is shown how, by choosing a particular such functional, an analyst who is investigating competing classification systems would be able to evaluate their performance with respect to a given set of quality requirements. The generic description of fusion and fusion systems on which this performance quantification concept is based uses category theory, a branch of mathematics useful for demonstrating mathematical relationships and properties of mathematical constructs.

The second contribution, by Haenni and Hartmann, concerns the problem of partially reliable information sources. They define a general model of partially reliable sources within the framework of Dempster–Shafer theory with a number of possible instantiations. Specific well-known models, such as two Bayesian models and two models where it is possible but not mandatory to include prior information, turn out as important special cases. In the general model every information source i delivers some

1566-2535/$ - see front matter
doi:10.1016/j.inffus.2006.03.004

information about the set $\{HYP, REP_i\}$, where the possible values of HYP are Hyp "the hypothesis is true" and $\neg Hyp$, and the possible values of REP_i are Rep_i "a positive report" and $\neg Rep_i$. As the only common variable among the information sources is HYP it is possible to marginalize each information source to $\{HYP\}$ before combining all of them. After combination it is possible to draw conclusions about the hypothesis under consideration. Starting from the general model the authors develop a concise and comprehensive model taxonomy using probabilistic argumentation. They describe and classify 21 different models of partially reliable information sources. In one model they propose to model the influence of a prior probability by a continuous parameter, corresponding to the situation where the holder of prior knowledge is not totally sure about his or her opinion. Four models are further explored in case studies. However, the authors conclude that the particular choice of model depends crucially on the specific application at hand and the nature of the available information.

The next contribution is by Niu, Varshney and Cheng. They propose a distributed detection and decision fusion scheme for a wireless sensor network (WSN) consisting of a large number of sensors. At the fusion center, the total number of detections reported by local sensors is employed for hypothesis testing. Based on the assumption that the received signal power decays as the distance from the target increases, system level detection performance measures, specifically probability of detection and false alarm, are derived analytically through approximation by using the central limit theorem. It is shown that for all the different system parameters explored, this fusion rule is equivalent to the optimal fusion rule, which requires much more prior information. To achieve a better system level detection performance, the decision threshold at the local sensor level should be designed optimally. Algorithms and guidelines for selecting this optimal decision threshold at the local sensor level are provided. The conference version of this paper won the Best Paper Award at FUSION 2004.

The fourth paper by Oxenham, Challa and Morelande concern an automatic target identification problem in a distributed network-centric environment where sensors deliver disparate types of uncertainty. They propose two novel Bayesian and generalized Bayesian distributed target identification algorithms for fusing target identity estimates generated by local heterogeneous data fusion systems. These may deliver their estimates either as finite probability distributions or as belief functions. The decentralized Bayesian and decentralized TBM generalized Bayesian approaches proposed are evaluated together with standard centralized Bayesian and centralized Dempster's rule approaches against 20 different two-sensor scenarios. The latter two approaches serve as benchmarks when given probabilistic data from both sensors and evidential data from both sensors, respectively. While the results of the four algorithms show remarkably similar rates of convergence for most scenarios the generalized Bayesian approach may be preferred based on its low computational cost and its graceful degradation when nodes of the network are lost.

The presentation by Makarenko and Durrant-Whyte concerns large numbers of autonomous sensing platforms connected into a network that promises better spatial coverage, higher responsiveness, survivability and robustness compared to a single platform solution. The Active Sensor Network (ASN) is a software framework for scalable, autonomous, and cooperative information gathering, combining decentralized information fusion and decision-making algorithms in a common architecture and concrete implementation. The focus of this paper is on the algorithmic side of the ASN framework, describing its probabilistic information fusion (Bayesian Decentralized Data Fusion, BDDF) and decision making (Bayesian Decentralized Decision Making, BDDM) algorithms. The data fusion layer of the ASN leads to a synchronised view of the state of the environment. Based on this belief, sensing platforms can employ BDDM to make appropriate individual action choices in order to maximise the team utility function.

The next four papers concern different novel approaches to target tracking.

The sixth paper by Maskell, Everitt, Wright and Briers investigates how to use graphical models for handling multi-target out-of-sequence data association. They review different approaches within a Bayesian framework and propose an architecture that orthogonalises the data association problem and the out-of-sequence problem. This makes it possible to use any combination of solutions to these two problems. A stochastic dynamic system is introduced that makes it possible to describe previous approaches to solving the problems associated with out-of-sequence measurements. In Fig. 1 it is drawn as a graphical model where the nodes represent variables and the arrows direct dependencies between variables.

The authors demonstrate that this model can accommodate the wide range of existing algorithms for processing out-of-sequence measurements. Furthermore, this framework can easily be extended to describe data association in the same context to solve out-of-sequence data association problems. The association problem is modeled by supplementing the graphical model with another set of nodes, Fig. 2.

In a multi-target tracking scenario the extended framework is shown to exhibit good tracking performance.

This is followed by the paper by Särkkä, Vehtari and Lampinen. They propose a Rao-Blackwellized particle filtering algorithm for tracking an unknown and time varying number of targets. The method is based on formulating probabilistic stochastic process models for target states, data associations, and track birth and death. These stochastic processes are tracked using particle filtering where the efficiency of the Monte Carlo sampling is improved by Rao-Blackwellization. The algorithm generates estimates of data association probabilities that can be used for approximating the probability of a hypothesis that

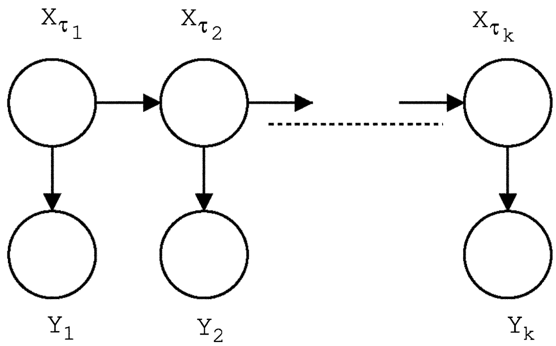

Fig. 1. Stochastic dynamic system drawn as a graphical model.

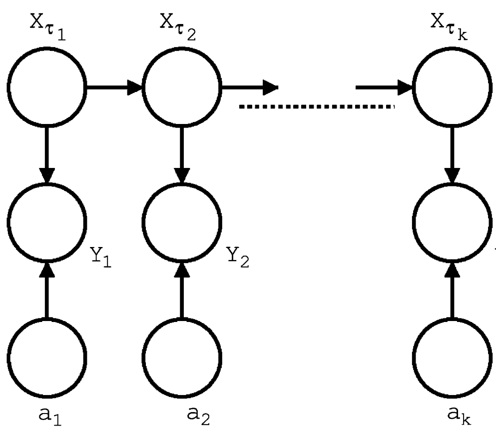

Fig. 2. Association variables a_k govern the measurement-to-target association at the kth iteration.

the target has disappeared from the surveillance area. In order to estimate the number of targets they model the birth and death processes in such a way that track formation and termination are based on the rules determined by the estimation algorithm for the probability model. The method is tested on two different tracking scenarios with an unknown number of targets. The estimation of the number of targets is shown to have good performance with only a slight delay after the disappearance of the signals before they disappear from estimation. Estimated target trajectories follow true target trajectories quite well.

Smets and Ristic develop an approach to joint tracking and classification (JTC) based on kinematic data using the Kalman filter and belief functions as understood in the transferable belief model (TBM). While the TBM solution to the tracking phase of JTC is essentially similar to the one achieved within the probabilistic framework the classification phase differs significantly from the classical framework. Here, they distinguish between the observed target behaviour and the underlying target classes, which are usually not in one-to-one correspondence. As with the probabilistic case the classification is based on the likelihood function. In TBM the likelihood function is equated to the conditional plausibility of the observation given the hypothesis. A scenario where three different target behaviours correspond to subsets of three different target classes is analyzed using probability functions and belief functions, respectively. In this scenario the first behaviour allows all target classes, the second behaviour allows two target classes and the third behaviour allows one of these two target classes. A Kalman filter is used for behaviour 1, and two different interactive multiple-model (IMM) filters are used for behaviours 2 and 3. While the problem is modeled probabilistically in three alternative ways they all show unsatisfactory classification results in some way. When the same problem is modeled using the TBM classifier a behaviour that allows several object classes will have equal pignistic class probabilities, while classes that are not allowed have a zero pignistic probability.

Exploiting 'negative' sensor evidence for target tracking is the subject of the paper by Koch. In addition to the sensor measurements, this approach also includes refined models describing sensor performance. By including sensor performance together with the sensor measurements in the processing chain it becomes possible to exploit 'negative' sensor evidence, i.e., to draw the conclusions that can be made from expected but missing sensor observations. Such 'negative' evidence will often appear as an artificial sensor observation. Including this information in the processing chain can improve the target position and velocity estimates. In order to do so within the Bayesian formalism, a formulation of a problem-specific likelihood function is needed. From a missing but expected sensor observation we may conclude that the target seems to move in such a way that it is buried in the clutter and infer information about its kinematic state. This may for instance be especially useful to help with early detection of stopping targets in ground target tracking, improved tracking performance of possible unresolved group targets, etc. The use of 'negative' evidence is studied in three different applications: group tracking, tracking with electronically scanned array (ESA) radar and ground moving target indicator (GMTI) tracking.

The tenth contribution by Bilenne deals with the problem of dynamic state estimation of continuous-time systems from discrete-time measurements in the context of high-integrity applications. The estimation scheme presented here is equivalent to the Kalman filter, with the difference that the data is not processed immediately, but collected in sets in preparation for a slightly delayed batched processing. This strategy is particularly suitable for fault detection, because the estimator naturally takes into account the cross-correlations of close-in-time measurements and the decisions can be based on more data. Dynamic tools for detecting faults and sensor failures are introduced. A new method for limiting the complexity of computing the posterior distributions in an integrity-oriented context is presented. Finally, the estimator is tested on a typical rail navigation problem.

This is followed by the presentation by Doherty, Łukaszewicz and Szałas which concerns real world applications where robots and software agents often have to be equipped with higher level cognitive functions that enable them to reason, act, and perceive in changing, incompletely known and unpredictable environments. In such a dynamically changing environment even a single agent may have varying abilities to perceive its environment. The situation becomes even more complex when different agents have different perceptual abilities and need to communicate and reason about their perceptions with each other. The authors propose a framework that provides agents with the ability to fuse both low and high level *approximate* knowledge in the context of dynamically changing environments while taking account of agents' heterogeneous and contextually limited perceptual capabilities. It is assumed that each agent has one or more approximate databases where approximate relations are represented using lower and upper approximations on sets. Approximate relations are represented by generalizations of rough sets.

The next to last contribution is by Chang, Jungert and Li. In this work, an approach to information fusion using a progressive query language and an interactive reasoner is introduced. The system basically consists of a query processor with fusion capability and a reasoner with learning capability. The query processor first executes a query to produce some initial results. If these are uninformative, the reasoner guided by the user may create a more elaborate query by means of some rule. The query is evaluated to produce a more informative answer. Novel in this approach is that application-dependent fusion rules may be specified by the user and subsequently learned by the reasoner. Examples are drawn from multi-sensor information fusion applications.

The issue ends with a paper by the guest editors and their colleagues. It describes work towards integrating fusion methodologies for achieving a dynamic common operational picture in a ground warfare scenario on the tactical level. This includes a new method of force aggregation based on Dempster–Shafer clustering and template matching, a new particle filtering method for multi-sensor multi-target tracking of an unknown number of objects, and a new resource allocation method. These are integrated to create, dynamically update, and maintain components of a tactical situation picture.

As guest editors of this special issue, we wish to express our sincere appreciation to Dr. Belur V. Dasarathy, the Editor-in-Chief, and to all authors and all reviewers for making this special issue possible.

Guest Editors
Per Svensson
Johan Schubert
Department of Data and Information Fusion,
Division of Command and Control Systems,
Swedish Defence Research Agency,
SE-164 90 Stockholm, Sweden
E-mail addresses: per.svensson@foi.se (P. Svensson),
johan.schubert@foi.se (J. Schubert)

Available online 11 May 2006