



Journal of Advances in Information Fusion

A semi-annual archival publication of the International Society of Information Fusion

Regular Papers	Page
Experiences and Challenges in Automated Support for Intelligence in Asymmetric Operations <i>Joachim Biermann, University of Skövde, Sweden</i> <i>Pontus Hörling, Swedish Defense Research Agency, Sweden</i> <i>Lauro Snidaro, Italian Ministry of University and Research, Italy</i>	101
The Probability Generating Functional for Finite Point Processes, and Its Application to the Comparison of PHD and Intensity Filters <i>Roy Streit, University of Massachusetts, USA</i>	119
Establishment of Human Performance Baseline for Image Fusion Algorithms in the LWIR and SWIR Spectra <i>Christopher Howell, U.S. Army's Night Vision and Electronics Sensors Directorate, USA</i> <i>Steve Moyer, U.S. Army's Night Vision and Electronics Sensors Directorate, USA</i>	133
Testing Under Communication Constraints <i>Sora Choi, University of Connecticut, USA</i> <i>Balakumar Balasingam, University of Connecticut, USA</i> <i>Peter Willett, University of Connecticut, USA</i>	143

*From the
Editor-In-Chief*

*Journal
Publication of
Conference
Papers*



INTERNATIONAL SOCIETY OF INFORMATION FUSION

The International Society of Information Fusion (ISIF) is the premier professional society and global information resource for multidisciplinary approaches for theoretical and applied INFORMATION FUSION technologies. Technical areas of interest include target tracking, detection theory, applications for information fusion methods, image fusion, fusion systems architectures and management issues, classification, learning, data mining, Bayesian and reasoning methods.

JOURNAL OF ADVANCES IN INFORMATION FUSION: JUNE 2013

Editor-In-Chief	W. Dale Blair	Georgia Tech Research Institute, Atlanta, Georgia, USA; 404-407-7934; dale.blair@gtri.gatech.edu
Associate	Uwe D. Hanebeck	Karlsruhe Institute of Technology (KIT), Germany; +49-721-608-3909; uwe.hanebeck@ieee.org
Administrative Editor	Robert Lynch	Naval Undersea Warfare Center, Newport, Rhode Island, USA; 401-832-8663; robert.s.lynch@navy.mil
Associate	Ruixin Niu	Virginia Commonwealth University, Richmond, Virginia, USA; 804-828-0030; rniu@vcu.edu

EDITORS FOR TECHNICAL AREAS

Tracking	Stefano Coraluppi	Compunetix, Inc., Monroeville, PA, USA; 412-858-1746; stefano.coraluppi@compunetix.com
Associate	Peter Willett	University of Connecticut, Storrs, Connecticut, USA; 860-486-2195; willett@enr.uconn.edu
Associate	Huimin Chen	University of New Orleans, New Orleans, Louisiana, USA; 504-280-1280; hchen2@uno.edu
Detection	Pramod Varshney	Syracuse University, Syracuse, New York, USA; 315-443-1060; varshney@syr.edu
Fusion Applications	Ben Slocumb	Numerica Corporation; Loveland, Colorado, USA; 970-461-2000; bjslocumb@numerica.us
Image Fusion	Lex Toet	TNO, Soesterberg, 3769de, Netherlands; +31346356237; lex.toet@tno.nl
Fusion Architectures and Management Issues	Chee Chong	BAE Systems, Los Altos, California, USA; 650-210-8822; chee.chong@baesystems.com
Classification, Learning, Data Mining	Müjdat Çetin	Sabancı University, Turkey; +90-216-483-9594; mçetin@sabancıuniv.edu
Associate	Pierre Valin	Defence R&D Canada Valcartier, Quebec, G3J 1X5, Canada; 418-844-4000 ext 4428; pierre.valin@drdc-rddc.gc.ca
Bayesian and Other Reasoning Methods	Shozo Mori	BAE Systems, Los Altos, California, USA; 650-210-8823; shozo.mori@baesystems.com
Associate	Jean Dezert	ONERA, Chatillon, 92320, France; +33146734990; jean.dezert@onera.fr

Manuscripts are submitted at <http://jaif.msubmit.net>. If in doubt about the proper editorial area of a contribution, submit it under the unknown area.

INTERNATIONAL SOCIETY OF INFORMATION FUSION

Wolfgang Koch, *President*

Darin Dunham, *President-elect*

Stefano Coraluppi, *Secretary*

Chee Chong, *Treasurer*

Yaakov Bar-Shalom, *Vice President Publications*

Robert Lynch, *Vice President Communications*

Uwe Hanebeck, *Vice President Conferences*


Pierre Valin (Acting Elisa Shahbazian), *VP Membership*

Sten F. Andler, *Vice President Working Groups*

Journal of Advances in Information Fusion (ISSN 1557-6418) is published semi-annually by the International Society of Information Fusion. The responsibility for the contents rests upon the authors and not upon ISIF, the Society, or its members. ISIF is a California Nonprofit Public Benefit Corporation at P.O. Box 4631, Mountain View, California 94040. **Copyright and Reprint Permissions:** Abstracting is permitted with credit to the source. For all other copying, reprint, or republication permissions, contact the Administrative Editor. Copyright© 2013 ISIF, Inc.

From the Editor-in-Chief:

December 2013



Journal Publication of Conference Papers

Recently, the publication of material in conference papers in peer-reviewed, archival journals has been a topic of great debate. Historically, say before 2000, the distribution and accessibility to conference papers were limited to those individuals who attended the conference or worked for institutions that purchased the proceedings. Furthermore, the peer-review process for conferences was often weak or nonexistent. Collection and distribution of the papers, collection of the reviews, and management of the whole process with the web-based tools of the time made implementing a peer-review process for a conference difficult. Thus, shortcuts such as the review of extended abstracts were often used in place of a review of the completed paper. Hence, given the significant differences in accessibility and peer-review, journal publications served a much different purpose than conference papers in those days and the question of duplicate publication was less of a question than it is in 2013.

Since 2013, the situation has changed significantly. Conference papers are now readily available electronically via IEEE Xplore or other web-based archiving systems. Thus, conference papers and journal papers have similar accessibility today. Also, the availability of web-based tools for managing the peer-review process and the desire for better conference papers have pushed most conference organizers to obtain peer-review of full manuscripts. Hence, what are the differences between the papers published in conference proceedings and those published in a peer-reviewed journal? How should JAIF policies reflect the situation of today?

Papers published in most conference proceedings are quite different than papers found in JAIF and most IEEE journals. First, a conference paper has passed a single review with an accept/reject decision and little or no verification of the requested changes. On the other hand, a (meaningfully peer-reviewed) journal paper has passed multiple review cycles with the reviewers and editors verifying that the referees' comments have been addressed adequately before it is published. The peer-review process of a journal ensures that the paper is of high quality and usually leads to a better paper and that reflects positively on the authors. Also, the length of

a conference papers is often limited so that the peer-review can be accomplished in a timely fashion. On the other hand, a journal paper can be longer and it will often include a more thorough review of the related literature and presentation of the contribution.

The JAIF editorial board has reviewed the policies of other journals and is establishing the following policy concerning the publication of material previously published in conference papers.

While direct submission of a conference paper by its author to JAIF is not acceptable, submission of an appropriately enhanced version of the manuscript is acceptable. Declaration that a manuscript has similar content to previously published conference papers is expected at time of submission and this will not affect the appropriateness

of the manuscript for journal publication. The submission to JAIF is expected to include at least 30% new material or be an integration of multiple conference papers into a comprehensive treatment of the problem under study. The overall quality of the submission to JAIF should be better with respect to quality of the explanations, literature review, derivations, examples, and illustrations. All associated conference papers should be cited in the submission.

This policy is very similar to that followed informally by the JAIF Editorial Board. The Operations Manual for JAIF has been updated to formally reflect this policy.

William Dale Blair
Editor in Chief

Experiences and Challenges in Automated Support for Intelligence in Asymmetric Operations

JOACHIM BIERMANN
PONTUS HÖRLING
LAURO SNIDARO

This paper presents some experiences and findings of the NATO RTO Task Group on Information Fusion in Asymmetric Operations. It briefly describes the functional processing steps in military intelligence presenting the underlying aspects of information processing and fusion and revealing main challenges for automatic support of the required functionalities. The extraction and structuring of relevant information from unstructured text documents is shown to be one of the fundamental steps where human operators need assistance. As an example of the state of the art the interactive tools PARANOID and CoALA are presented. They provide the basic information and knowledge structure for all subsequent information processing like Link Analysis and Social Network Analysis. The use and benefit of CoALA will be illustrated by results from a military experiment. Finally, with respect to further research, open questions and new approaches for the support of intelligence production are discussed concerning automatic information structuring and discovery as well as pattern and behaviour based threat assessment. In relation to pattern based threat assessment a third tool, called Impactorium that is developed for threat assessment in military as well as in civilian environments, will be briefly described.

Manuscript received August 26, 2010; released for publication February 8, 2012.

Refereeing of this contribution was handled by Chee-Yee Chong.

Authors' addresses: Joachim Biermann, Dept. for Sensor Data and Information Fusion, Fraunhofer FKIE, 53343 Wachtberg, Germany, E-mail: (joachim.biermann@fkie.fraunhofer.de); Pontus Hörling, Dept. for Information and Aeronautical Systems, Swedish Defence Research Agency (FOI), SE-164 90 Stockholm, Sweden, E-mail: (hoerling@foi.se); Lauro Snidaro, Dept. of Mathematics and Computer Science, University of Udine, Udine, Italy, E-mail: (lauro.snidaro@uniud.it).

1557-6418/13/\$17.00 © 2013 JAIF

1 INTRODUCTION

The conduct of intelligence is an essential task not only in military command and control but also for homeland security and disaster management. A most accurate awareness of the actual situation, including an assessment of the potential development and threats, is essential prior to all decisions and own activities. An intelligence cell needs the capability to collect, process, and disseminate a wide variety of data and information produced by the full spectrum of technical sensors, human intelligence, i.e. intelligence gathered from humans via observations, interviews etc. (HUMINT), and socio-political sources. There are a number of major challenges for the conduct of intelligence: first, there is a danger that the processing capability will be overtaken by the sheer volume of information that is available in very large quantities and various formats. Second and especially true for asymmetric threats, by its nature the collected information and knowledge mainly are unstructured, typically provided as text documents. Therefore, as an inevitable precondition for being processed automatically, relevant information aspects have to be extracted and structured efficiently so that this type of input can be readily and efficiently exploited for all of its intelligence value without loss of rigor [6] [29]. The urgent requirement for reasoning methods and procedure which give automated support to the further analysis and integration of structured semantic information defines a further challenge. Shortcomings in the ability to make deductions about missing and conflicting information and the current inability to support automatic context based correlation and reasoning about vast amounts of information are drawbacks to providing a coherent overview of the unfolding events.

This paper, extending the concepts presented in [8], describes some results and findings of a series of NATO Research and Technology Organization (RTO) Task Groups on Information Fusion of which the authors are members. By revisiting the intelligence process with particular attention paid to collation and analysis, the requirements for automated support are exposed and examples of existing solutions are presented.

1.1 Structure of the paper

Section 2 will explain the main steps of information processing in intelligence and explains some of the major challenges with respect to support and automation. A short description of heuristic human intelligence processing is given and two main aspects for support are presented. Section 3 then introduces two tool suites for automatic Collation and Link Analysis which focus on the ideas for support presented in Section 2. Finally, some findings from a military trial testing on one of the tools are presented discussing its benefit to the military users. Section 4 and 5 discuss further aspects of intelligence processing which are still unsolved with

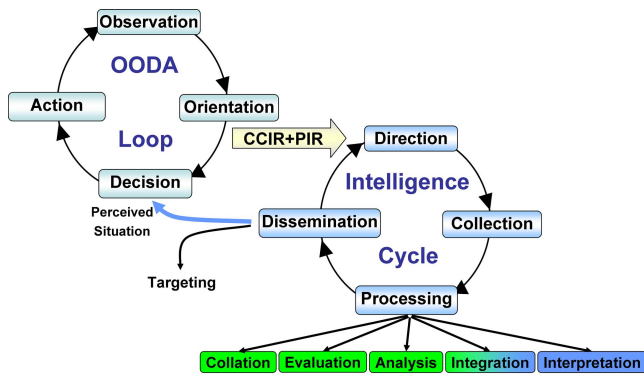


Fig. 1. Intelligence Cycle interfacing with OODA Loop

respect to automation, shortly referring to a third tool. In Section 6 we conclude our findings.

2 THE INTELLIGENCE CONTEXT

Intelligence processing is an important part of Command and Control (C2) because a completely accurate situational awareness of the situation is essential prior to all decisions and successful activities. In order to fulfil the requirements of all the various users in the military area and to provide an appropriate picture of the Area of Operations or Interest in the most timely and reliable fashion, intelligence cells have to process and evaluate all incoming information. A wide variety of information produced by the full spectrum of sensors and human sources has to be collected, filtered, processed and disseminated. The final goal for intelligence is often to provide, roughly speaking, a decision support for assessing our room for manoeuvre in the current situation. So, besides giving timely situational awareness, these working processes are also expected to give well founded assessments on opportunities for own forces, people and the infrastructure we need to protect as well as threats and risks against them. This will be discussed more in detail in Section 5.

2.1 Process flow and functional steps of the Intelligence Cycle

The processing of information for the production of intelligence is performed in a structured and systematic series of operations which is called the Intelligence Cycle. It includes four stages, *Direction—Collection—Processing—Dissemination*, which are defined by the NATO Glossary of Terms and Definitions (AAP-6) [33] as follows:

Direction: “Determination of intelligence requirements, planning the collection effort, issuance of orders and requests to collection agencies and maintenance of a continuous check on the productivity of such agencies”

Collection: “The exploitation of sources by collection agencies and the delivery of the information obtained to the appropriate processing unit for use in the production of intelligence”

Processing: “The production of intelligence through collation, evaluation, analysis, integration and interpretation of information and/or other intelligence.”

Dissemination: “The timely conveyance of intelligence, in an appropriate form and by any suitable means, to those who need it.”

These four discrete stages are conducted culminating in the distribution of the finished intelligence product. The representation of the military intelligence function in Figure 1 is strongly connected with the so called OODA (Observe, Orient, Decide, Act) Loop as intelligence is an integral part within the military command and control cycle. Boyd introduced the notion “O-O-D-A” and he stated “The process of observation-orientation-decision-action represents what takes place during the command and control process—which means that the O-O-D-A loop can be thought of being the C&C loop.” [3]. The intelligence effort is “Directed” by the Commander’s Critical Information Requirements (CCIR) from which his Priority Intelligence Requirements (PIR) are derived. Eventually providing the military commander with a most timely and comprehensive situational picture the intelligence cycle supports both the *Orientation* and the *Decision* phase.

The Processing phase is the most essential part with respect to information fusion issues. It is defined as “The production of intelligence through *Collation, Evaluation, Analysis, Integration* and *Interpretation* of information and/or other intelligence.” [33]. It is a structural series of actions where the information, which has been collected in response to the directions of the commander (CCIR, PIR), is converted into intelligence products. A more detailed discussion on the principles of heuristic intelligence processing can be found in [6]. It is here, that the intelligence staff needs automation to be more effective in its work. In cooperation with an international group of military experts and based on realistic asymmetric scenarios, the established heuristic procedures of intelligence processing have been analysed to understand the approach of the human experts and their cognitive processes in order to adapt their reasoning principles and methods for automated fusion concepts and procedures.

Figure 2 [31] illustrates the relation of the different processes supporting the Intelligence Cycle (shown in Figure 1) now organised as functional flow and having its focus on the Processing phase. Specifically, Collation is presented in more detail and the steps Analysis, Integration and Interpretation are grouped and renamed to “Link Analysis” which is, in some nations, how Analysts name their job. As mentioned before, the CCIR and other information requirements of the commander and his staff initiate the intelligence processing (see ① in Figure 2). Incoming information first has to be digitised, if necessary, logged and stored into a data management system. This part is covered by ② in Figure 2. The main function of such document management relates to the

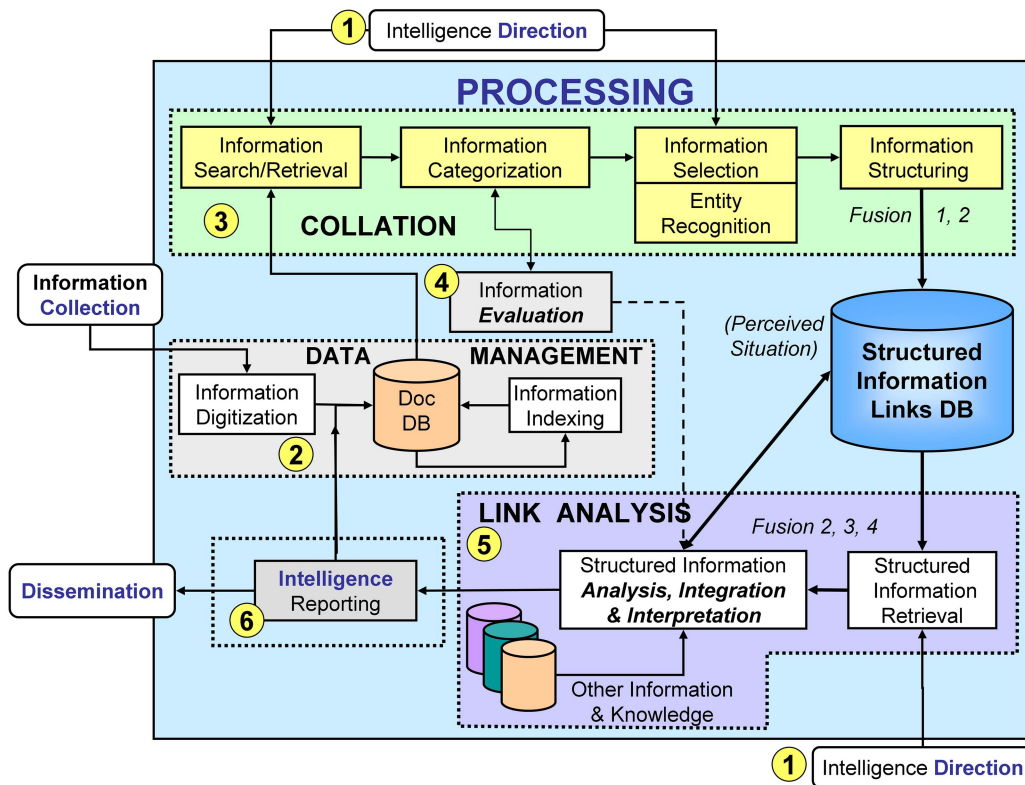


Fig. 2. Intelligence processing from unstructured to structured information

ability of registering and storing structured and unstructured documents in a document database, and of discovering knowledge. The function of knowledge discovery refers to the different ways of searching and retrieving information from large information sources with interactive capabilities of guiding the user through the process. It exploits structures such as semantic networks, ontology, and meta-data to establish links between domain models and information sources, and helps users to find relevant information. These functions directly support the Collation process described below (see also Section 4).

During the *Collation* process (indicated by ③ in Figure 2) information is decomposed into individual items which are grouped according to categories relevant to the context of the mission and cross-referenced with previously processed information items.

From an operational context, it is known that especially in asymmetric operations much of the incoming information is to be found within text documents and is often not in a format suitable for machine manipulation. Therefore, any automated support of the collation step essentially requires the extraction of relevant information from incoming unstructured pieces of semantic information as well as the structured representation of these newly processed information items. One of the main purposes and benefits of the tools described in chapter 3 actually is to support this information structuring (see section 3.1.3).

The *Evaluation* of the reliability of sources and the credibility of collected information is done by intelligence analysts as soon as the relevant information has been extracted and can be annotated directly as a tag to the piece of information or the document (④ in Figure 2). In the context of semantic information from HUMINT and Open Source INTElligence (OS-INT) sources, newspapers, journals, the web, blogs, twitter etc., evaluation is most often a very experience-based task with highly subjective results. For these reasons *Evaluation* was not regarded to be done or supported automatically.

Analysis: "...information is subjected to review in order to identify significant facts for subsequent interpretation" [33]. It consists of a number of interacting sub-processes resulting in the analyst answering questions like: "Who/What is it?," "What does it mean?," "Why is it happening?" etc. in order to recognize indicators and warnings.

Integration: "...analysed information or intelligence is selected and combined into a pattern in the course of the production of further intelligence" [33]. It is the process of building pictures of the current and of the predictive situations from all the gathered and analysed information.

Symbol ⑤ indicates where Analysis and Integration of information are conducted. In practise they are very often performed as one combined step and they are not conducted as separated parts of the overall process flow. It is here that intelligence is produced and the fusion of

information takes place. The notion “Fusion 1, 2” in ③ and “Fusion 2, 3, 4” in ⑤ used in Figure 2 refers to the level of data fusion as it is defined by the data fusion model of the US Joint Directors of Laboratories (JDL) [28].

A further important requirement for an intelligence processing system is to be able to support link discovery and analysis. This approach (compare so-called “Story-telling” [44]) depends on the capacity of the system to automatically or semi-automatically allow the identification of a specific object of information and all of its related categories such as the location, the time, the cause, the originator, the subject, etc. Once those links are enabled, identified and validated, analysts will obtain a better and more focused image of the situation. Disparate pieces of information that had little or no value when considered independently could have a whole new meaning when combined and linked to form a pattern. Link creation is carried out during the “Information Structuring” process found in ③ and link discovery and analysis during the “Structured Information Analysis and Interpretation” process found in ⑤ with the information stored into the Structured Information Links Data Base shown in Figure 1. Link analysis is a capability that can support both the collation and analysis processes. We will discuss link discovery and analysis further in Section 3 of this paper.

To summarise: by categorising, classifying, indexing and cross-referencing all information appropriately the intelligence organization avoids losing important information and context. Disciplined and methodical collation enables further analysis to be efficiently performed using link analysis among other techniques. Information systems support this approach depending on their capacity to automatically or semi-automatically allow identification of specific information and all of its related categories such as the location, the time, the cause, the originator, the subject, etc. Once those links are identified and validated, analysts are given better bases for understanding the different key factors influencing the overall situation. Disparate pieces of information that have little or no value when considered independently could have a whole new meaning when combined and linked together thus allowing the emergence of potential key patterns.

2.2 Challenges and main requirements for automated information processing

All the different processes shown in Figure 2 are relevant to the conduct of intelligence but the three processing steps shown in Table 1 [7] were determined as those ones which, on the one hand, are central to the conduct of intelligence, and, on the other hand, were supposed to be capable of being automated.

To be able to build a system for (semi-) automated intelligence processing and decision support incorporating these functionalities, at least the following requirements and challenges have to be met:

TABLE 1
Required functionality for automated information processing in intelligence

Step	Required Functionalities		
<i>Collation</i>	Semantic text extraction	Categorisation	Information structuring
<i>Analysis</i>	Classification Identification	Correlation	Link analysis
<i>Integration</i>	Pattern matching	Aggregation	Fusion

1. Semantic access to all input information

Within the *Collation* step operators have to deal with a continuum of different types of information and all available input information and data should be used for the production of a reliable and most comprehensive operational picture as a base for situation awareness. Therefore it is necessary to be able to get full semantic access to the content of all unstructured text documents. The Battle Management Language (BML) is an unambiguous language which, among others, is used to provide for situational awareness and a shared, common operational picture. It is a promising linguistic approach to structure free text information for decision support and automatic information fusion [37], [36] (see also section 2.4 for BML use in knowledge representation). For a wider use of BML in military Command & Control and simulation see e.g. [35] [10]. In Section 3 two software tools are presented which support the interactive extraction of text from documents for categorising and structuring information relevant for respective information requirements (see esp. Subsection 3.1.3). These tools support the heuristic operating procedures of the human operators.

2. Understanding human reasoning in intelligence processing

Within *Analysis* and *Integration* significant information has to be found within the information set which is provided as the result of the *Collation* step. This significant information has to be put together to a situational picture according to the information requests given by the respective commander. Human analysts develop an appropriate view of the theatre which means they have a mental model of all relevant aspects of own and hostile operations as well as of the activities of the environment. Link Analysis is a technique well known by intelligence analysts and other security organizations that allows for the detection and visualization of interrelated topics to help resolving the “effects-to-cause” puzzles (see Figure 3) which arises when trying to put together all pieces of fractional information to form a coherent and reliable picture of the real situation. Ongoing research in this area is discussed in Sections 4 and 5.

The analysts have to solve many different puzzles at the same time. The underlying problem is the same as in risk assessment and threat detection in civil secu-

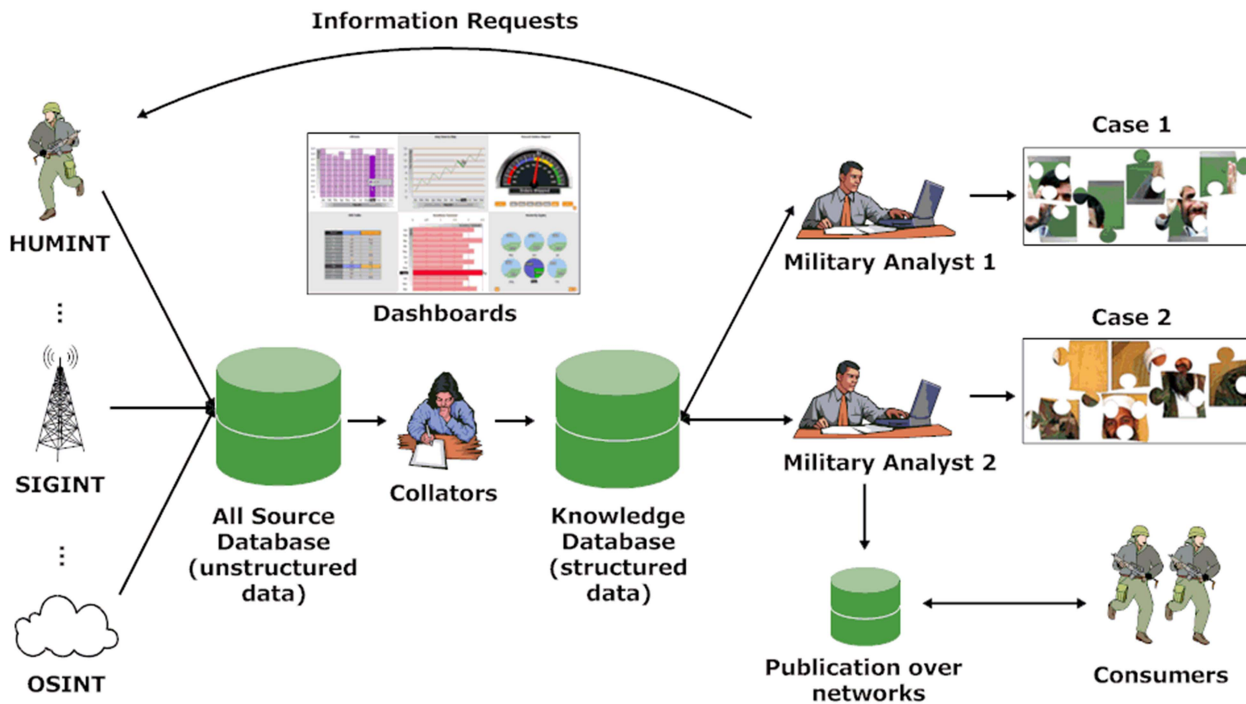


Fig. 3. Solving the many incomplete puzzles in intelligence [48]

rity. Police officers tasked to determine the structure of criminal groups, the territories or business areas they control, and the tactics, logistics and communication they use will follow the same reasoning principles as in military intelligence. Some basic aspects of cognition in heuristic intelligence processing are given in Section 2.3.

3. Knowledge representation and methods for automatic reasoning

To process semantic information from text documents the functionalities listed in Table 1 require an ontological model of the domain of application. It has to be established in a formal and structured way as a sufficiently appropriate and consistent approximation of the real situation and activities. Furthermore, reasoning methods operating on this model have to be developed that are able to cope with incomplete and imperfect information. One successful approach can be seen in using the principles of default or pattern based reasoning deducing from structured semantic information. Some arguments about this topic are given in Subsection 2.4.

There are more challenges in building support systems for intelligence processing, e.g. how to cope with assessing or estimating uncertainty in text information or dealing with deceptive or wrong information. The tools presented in this paper have not been designed to support these highly experience based and subjective information processing steps. Approaches to deal with the before mentioned questions can be found e.g. in [38] [24] [52].

2.3 Heuristic intelligence processing and default reasoning

Intelligence cells will never get all the relevant information they request, but they will be bombarded by partial, false, unreliable, irrelevant, and redundant pieces of information which they will have to filter according to the information and intelligence requirements given to them by the commander. The human brain permanently selects, relates and inserts relevant information guided by its internal, mental abstract model of the world to further develop the perception of the reality which can be understood as a concrete instantiation of the abstract world model. Human analysts use their experience gained from “similar” problems in order to fuse all data and information into a reasonable picture of the situation thus deciphering the meaning of all pieces of input data. This describes roughly the main steps and aspects of cognitive reasoning based on patterns, schemata, learning and experience as it can be found e.g. in [1]. The principles of “default reasoning” are not depending on the specific problem area but are a general human problem solving paradigm.

In heuristic intelligence processing the following constraints have to be stated with respect to the available data, information and knowledge:

- Usually only general and incomplete information is available about the structure, the activities, the situation and the intent of adversaries and other involved factions

- Information gained by reconnaissance is imperfect and incomplete

This means that because of lacking knowledge and information, neither the mental model of the situation (“world”), which analysts have, is perfect, nor are they provided with a perfect actual view on this world by the incoming information stream. In order to deduce a most reliable picture of the situation, in spite of these weak preconditions, analysts practise a method of heuristic reasoning which relies on the assumptions that

- Actors operating in a professional way according to their doctrines, rules, principles, modus operandi, and standards
- Effective operations are planned and organised for the benefit of the mission
- There are conditional dependencies among the phases and steps of an operation
- The real situation picture can be deduced from characteristic patterns of activity or of state (called templates or schemata)
- Templates can be recognised from significant information

It is common military experience and expert knowledge that the production of intelligence can be done by integrating current information based on the assumption of default behaviour. Behaviour modelling, such as doctrinal templating [4], is a descriptive, qualitative method of knowledge representation. The elements of the situation are not only described by their attributes but also by their relations and dynamic behaviour as well as their operational potential and assumed tactical intentions.

For supporting intelligence systems two different kinds of models are relevant [6]:

- a) Behaviour models describing the tactics of potential adversary factions and all necessary pre-conditions for their hostile activities;
- b) Models describing “normality,” the common and unsuspecting behaviour of defined subgroups of the population or other elements and groups relevant to the situation

In case a), analysis is the task of detecting special indicators of activities or status which define by their combination a potentially evolving threat. This approach is used e.g. in low and high intensity conflicts of military or paramilitary type. Threats, like an ambush or an Improvised Explosive Device (IED) attack, are complex sequences and interrelations of different activities. Each of them having their own structure (pattern) and being combined they form the high-level pattern of the final threat.

As a consequence, a pattern to be used in automatic reasoning has to consider and incorporate all significant and characteristic factors in order to build up a template for analysis and integration. Little and Rogova

[26] discuss the formal ontological structure of threats as holistic phenomena possessing three interrelated parts: intentions, capabilities and opportunities (further elaborated on here in section 5, where we will discuss the Indicator concept). They show how these facets of a threat are related to one another, as well as to states of vulnerability. These aspects have to be covered by surveillance and reconnaissance to provide the intelligence staff with corresponding information.

In the above mentioned case b) concerning “normality,” the task is to detect deviations from patterns of “normal” behaviour. Snidaro, et al. [45] give an informal definition of an anomalous event as “a deviation from common patterns of activity.” This method is used e.g. in combating terrorism to be able to define indicators of suspicious activities [7]. It is an increasingly important topic for decision support, since it can give hints to the intelligence staff towards where more analysis or information is needed.

Intelligence operators often have to deal with information sources that can provide a sequence of unreliable observations or reports due to unfavourable sensing conditions or limited/erroneous projection of real-world observables. In the case of human sources, information could even be deliberately incomplete, erroneous or deceptive. Poor quality information and unreliable sources can have disruptive effects on the fusion process. For this reason an automatic Situation Assessment (SA) system for intelligence purposes should take into account both the reliability of the sources [41], and the quality of their data [40] to regulate the fusion process accordingly. While (automatically) weighting or pruning information is far from being a trivial process, these topics are being actively researched by the fusion community.

2.4 Knowledge representation for automatic information processing

In a decision support system all actual and background information, including the models of behaviour and normality, has to be processed automatically. The representation of information and knowledge can be based on an ontology of the domain. Ontology, as a semantic description of all objects and classes (or categories) and their relations, incorporates taxonomies, attributes of the objects and respective values and constraints, rules and schemata, representing the behaviour defaults [23] [43]. Schemata can be well represented as so called feature-value matrices (FVM). These are sets of features (or attributes) and value pairs. For schemata, on the top level, the features denote the thematic roles of the represented object or class and the values are feature-value matrices themselves that pool the information about the object that fulfils the respective role. From the mathematical view, a feature-value matrix is a finite set of pairs. Each pair consists of a feature and a value. A feature is always a symbol, a value, however,

can be a symbol or a feature-value matrix itself. In addition, the following uniqueness condition holds: every feature has a unique value. In other words, in a matrix there cannot be two pairs which share the feature but not the value. However, different features may have the same value.

These matrices have many beneficial properties. First, objects can be represented which are not specified completely. This is important for fusing partial information. Second, the matrices obey XML schemata which allow further automatic processing. Third, the matrix representation allows ‘unification,’ a standard computational linguistic algorithm for merging information which we regard beneficial in information fusion [43]. FVM are used beneficially as a technical representation form within the linguistic BML approach to knowledge representation.

3 TOOLS SUPPORTING INFORMATION PROCESSING

Coping with the more or less diffuse adversaries in asymmetric warfare as well as unfolding the structure of criminal networks has resulted in the introduction of several dedicated tools for intelligence analysis in the last decade. They give the ability for analysing social and semantic networks, spatio-temporal pattern analysis, all with different abilities of visualization [9]. One example of the former is Palantir [21] and Detica NetReveal. Other tools focus even more on Visual Analysis such as NetLens [18], IN-SPIRE [39] and Jigsaw [19]. However, such tools still have their limitations [44].

Two examples of existing support tools for actual intelligence processing are presented in greater detail here. The special features which relate to the before mentioned process flow and required functionality are highlighted. In Subsection 3.2 some results of a military trial on intelligence processing using one of the tools are given and the requirements of the military analysts with respect to more elaborated automatic support for analysis and integration are presented.

3.1 Interactive tools for Collation and Link Analysis

CoALA and PARANOID [48] are products of a close and intensive collaboration effort between Defence Research and Development Canada (DRDC), Quebec, Canada, and the Dutch research organization TNO Defence, Safety and Security, Den Haag, The Netherlands. They have been developed in parallel to the activities of the NATO RTO research Task Groups on Information Fusion active since 2000 and are related to the results of these groups. CoALA is based on PARANOID and it was supposed to be in operation in 2009. For different reasons the fielding had been postponed to 2010. These tool suites provide the intelligence personnel with a functionality that supports the collation of free-text documents. It does so by supporting interactive extraction of relevant information from free

text source documents and storing that information to a structured database to be further analysed and related to other items of information, thus creating intelligence. In brief the general characteristics of the tools are:

- Rapid collation of unstructured text information into pertinent intelligence products
- Identification of hidden patterns and connections within information to focus analysis on counterterrorism, organised crime, threat assessment and incidents
- Collaborative collection and analysis enabled

A more detailed description of the information processing approach underlying both tools and the implemented functionalities can be found in [30].

3.1.1 PARANOID

PARANOID (Program for Analysis Retrieval And Navigation On Intelligence Data), developed by TNO. In this tool techniques for searching, storing and analysing information are being implemented and tested. This tool suite supports the process of specifying the total functionality for an operational processing system for intelligence, such that it reflects the workflow of intelligence staff. PARANOID processes information in support of Peace Support Operations (PSO), but is equally applicable to other areas such as counter-terrorism operations, the fight against fraud, and the acquisition of business intelligence.

The functions of PARANOID reflect the workflow in the intelligence process, starting with the definition of information need through to the storage of the intelligence products. Three main functional areas have been defined:

Profiles: In this function the user is able to define certain factors, such as time and space definitions, certain types of events, and particular individuals that have to be taken into account while processing the incoming information.

Documents: This function carries out a range of different operations on all incoming information. One example is the storage and transformation of structured and unstructured data from documents into a structured database, carried out by applying different information extraction techniques.

Analysis: There is a need for different types of analyses to be able to support the different sub-processes of Processing: link analysis, pattern recognition, trend analysis and threat/risk analysis. There is also a need to be able to visualise the data and results. This should be possible not only by using a geographical information system, but also through a number of innovative ways of navigating through a network of different types of related data and information.

3.1.2 CoALA

The Collation and Link Analysis (CoALA) tool is an evolutionary specialized collation tool suite for intelli-

IntObject association to relevant topics / categories in intelligence reports

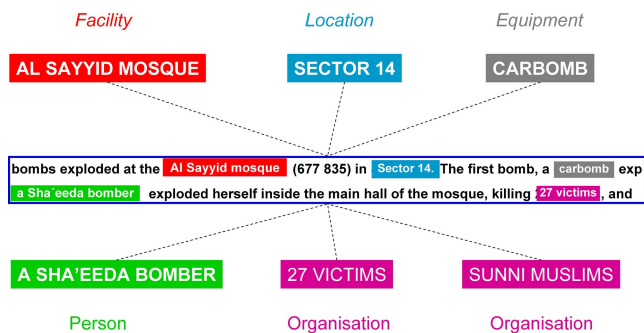


Fig. 4. Interactive information extraction and structuring

gence analysts based on PARANOID and developed by DRDC. It provides expert applications to exploit unstructured information and populate a structured intelligence database that allows detailed analysis and production of intelligence. CoALA version (1.0) provides a suite of functions that support the collation and analysis process. The key functions are as follows:

Document management: Basic document management functions such as importing, registering, storing and disposing of documents.

Information Management: CoALA includes a structured knowledge database that provides a means to record common pieces of information and intelligence in an organized fashion that support the retrieval of that information and intelligence. With respect to the amount of cases to be managed during a military mission there is a function to get detailed information about the record management. This management is closely tied to the management of Priority Information Requirements (PIR) and other Information Requirements (IR) related to the management of intelligence collection.

Data Collation: Capabilities that allow pieces of information to be related to each other, grouped in related categories, and stored into the knowledge database.

In detail CoALA provides an information categorization tool. This function is used to identify the various objects found in the text and associated them with a class of objects predefined by the user. Based on the intelligence analysts' experience, the following classes had been chosen: person; equipment, facility, organization, event and geo-object (location, map reference, coordinates). To establish a structured information set there is a function to create, manage and visualize relationships between objects in a central knowledge database.

Data Analysis: To conduct link, pattern, geospatial and temporal analysis of information and intelligence. The results are stored into the knowledge database. There is a function to analyze and create working assumptions supported by visualization of the content of the master

knowledge database. This visualization is done using various tools, including:

- Charts of information objects and their interrelations (Link Chart);
- Timeline charts for events and their interrelationships;
- Matrix of links and relationships existing between various types of objects. The most common example is the matrix of what is known in an organization;
- Basic geospatial visualization of geo-referenced objects (GIS);

Intelligence production management: Simple means to capture and manage the IR/PIR list and to link the intelligence production back to it. The tool allows for any intelligence products (assessments, analytical charts, briefings and reports) to be stored in the knowledge database with references to all of its supporting material. Furthermore CoALA (v1.0) provided

- a limited printing tool.
- a function for exporting Link Charts to the commercial i2 Analysts' Notebook product.
- collaborative work in real time via a common database (MS SQL 2005).

3.1.3 Information extraction and structuring

One of the core concepts for good analysis in both tools is the collation concept: the extraction of relevant information from unstructured information into structured knowledge. The extraction of information is predominantly done by interactively tagging relevant parts of sentences from documents ("Statements") and linking them to the so called "Intelligence Objects" or "IntObjects."

Int Objects are elements of categories of domain items as Persons, Organisations, Location, Equipment and Facilities. Figure 4 shows an example of a statement (in the rectangle) that is linked to other IntObjects. The Statement contains different IntObjects that are linked in a standard way ("related to"). Figure 4 gives an example how the relationships between IntObjects, like between the Person "A Sha'eeda Bomber" and the location "Sector 14" is established by extracting and tagging the single information items.

The newly established set of structured information is to be integrated into the knowledge base (KB) which represents the so far perceived situation. The KB is searched for already existing IntObjects which are the same or may be the same as one of the elements of the newly structured information set. Figure 5 shows that two IntObjects "A Sha'eeda Bomber" and "Carbomb" are already known within the KB. They are offered to the operator to verify and confirm that the already known IntObjects in the knowledge network are identical to those ones which are part of the newly structured information set. If this is true the new IntObject structure is merged into the KB unifying the identical IntObjects which results in more comprehensive and/or

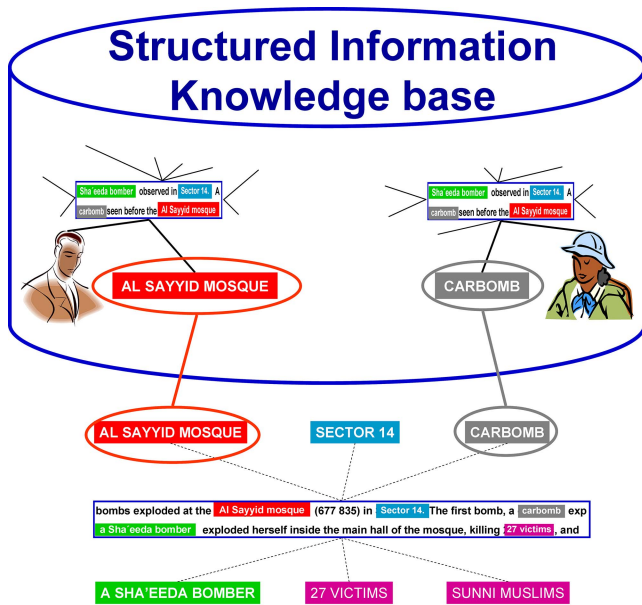


Fig. 5. Integrating the new structure into the knowledge base

more precise and/better confirmed knowledge about the respective IntObjects. The graphical representation of the result is shown in Figure 5. The special benefit of this information integration lies in the now established connection between the two persons, shown by the red line in Figure 6. These two actors now are related to one another by this merged-in information structure.

3.2 Expert trial on intelligence processing

The investigation of the RTO Task Group had been carried out with the support of an international group of military advisors. They focused on the structure and process flow of the conduct of intelligence, the human cognitive methods and practical procedures on how to process the collected information and available knowledge. This analysis was based on several scenarios, starting with a more conventional low intensive operation dealing with the Kosovo conflict and finally using an Iraq-type asymmetric operation. The insight gained into the main character of the conduct of intelligence did not change over the varying conflicts and the necessary steps which have to be done in the course of the production of intelligence seemed to remain the same. This is at least true for the more abstract point of view of a paper work analysis. But there was no certified and detailed information on how the processing of intelligence is carried out under real conditions by analyst experts of the intelligence branch. In particular there was only little information about the detailed breakdown and organisation of the work, the sharing of information and partitioning of responsibility.

Up to the mid of the last decade intelligence cells in operation have been using mostly standard office tools to manage information and data without any specific functionality and support for exploiting its intelligence

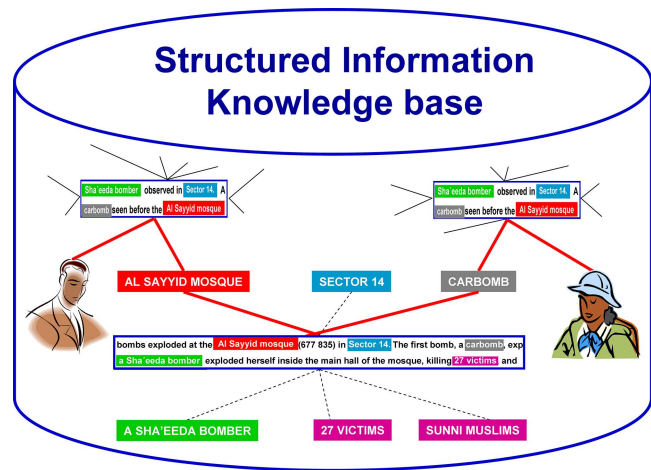


Fig. 6. The merged information reveals new relations in the knowledge network

value. It had been recognised that collators do not perform their tasks efficiently under these circumstances. They tend to transfer the burden to intelligence analysts who have to complete the collation process. In order to carry out a knowledge elicitation a Subject Matter Experts (SME) trial with domain experts coming from Afghanistan was arranged by the Canadian Forces. Six intelligence experts, using the CoALA tool suite, constituted an All Source Intelligence Cell. Their task was to work on a set of CCIRs and PIRs based on the context of a so far unknown asymmetric scenario. The intelligence trial was performed:

- to analyse whether the military understanding of the conventional processing steps in intelligence had been carried over to asymmetric operations, to observe the real workflow and processing steps of the experts,
- to observe the processing of unstructured text information carried out by experts experienced in asymmetric operations using the support of an interactive collation tool,
- to analyse the way of human deduction and main reasoning principles of the experts,
- to observe in how far and to which extent the supporting functionality provided by CoALA are accepted by military experts,
- to validate the usability, capability and potential of the interactive supporting tool CoALA, the acceptance by the intelligence experts from the Canadian Forces and to get recommendations for further enhancement and development of collation, analysis and integration functionalities.

The behaviour of the experts which was observed first was caused by the fact that during the last years, the All-Source Information Cells (ASICs) have been overwhelmed by unstructured text information. Often no IT support has been available or IT support was too unqualified or unusable for the task of structuring this input according to established and proven intelligence requirements. It therefore was natural for the *collators* in

our trial to take the task just to read the messages, identify important information mentioned within the messages, inform the analysts about the interesting observations and organize the messages in a way that they could be accessed easily when required. The linking of intelligence objects was only carried out by the *analysts*. They were re-reading the messages directed to them by the collators and, which from the point of view of the individual analyst, were of interest to the very Priority Intelligence Requirement (PIR) or Information Requirement (IR) he/she was working on. Therefore, at the beginning of the trial, the collators were told by the leading officer just to tag the statements and other intelligence objects, but not to link the statement to objects. This easily could and should have been done using the Collation tool CoALA to keep the connection between newly created intelligence objects and the constituting and justifying statement or message. Later on, the analysts started complaining about the fact that they only could retrieve “standalone” intelligence objects, as there had been no links established to be analysed. The analysts were almost doing the Collation process again. Therefore, after some time, the collators were asked to establish the links between processed statements and other intelligence objects. Establishing this “new” work schedule, the ASIC personnel only returned to the well defined and commonly performed procedure and work share as they used to be before the overwhelming flood of information degraded the role of the collators to just tagging the information. It was observed that PIRs were the leading factor in directing the information processing for intelligence. Processing of any information which could not be related to the list of CCIRs and PIRs was not observed. Nevertheless it would be of interest how analysts cope with developments outside the scope of interest.

Despite several difficulties the SMEs encountered during the trial, they were able to reach their “operational” goals. They answered the CCIRs and they were able to give detailed recommendations and assessments on the CoALA functionality, although they might not have fully experienced and tested the full potential of the CoALA concept of information extraction and structuring.

The semi-automated tool functions which support the extraction and structuring of information are going far beyond the low level requirements of the NATO AAP-6 [33] Collation step definition, which only claims for “grouping of information.” The support given by information structuring tools like those presented above will enable users to establish and persistently keep relations between pieces of information and to give the rationale for these associations. By this, a complex information structure is being developed in a cooperative way to be used commonly by all users in the ASIC. This persistent knowledge gives insight in the dynamically developing situation and serves as bases to all further

intelligence processing, as it has to be done during the analysis and integration step.

The assessment of the military intelligence experts concerning the usability of the CoALA tool was as follows: Despite the fact that it was still under development, the members of the intelligence branch considered it usable enough. The tool was perceived as being easy to learn and intuitive because of the new paradigm introduced to directly support the process of transforming unstructured information to structured information. It followed user interface architecture best practices but did not yet implement all the Microsoft Windows GUI practices. One function that collators really appreciated was that the text document processing tool retains objects once initially categorized. This speeded up collation work. The tool also supported collaborative work via a central server and knowledge database. This let the user know immediately if a certain object was already retained and recorded in the system, thus preventing pointless duplication.

4 AUTOMATIC KNOWLEDGE DISCOVERY

This section discusses other available technologies and current research directions for automatic knowledge extraction from data. The early stages of intelligence processing are largely deductive detection processes, performed by intelligence operators who look for relevant information in the intelligence information repository assisted by software applications that support information search and retrieval. Indexing and cross-referencing are processes that can be performed automatically, even by off-the-shelf software, as the documents are filed in the database. These simple steps already add value to the database as they provide means for retrieval of digitised documents and navigation capabilities within the information repository. COTS (commercial off-the-shelf) software is therefore already available to support the “Information Indexing” phase of the “Data Management” processing block of Figure 2. The specialized tools described in Section 3.1 offer dedicated functionalities to ease this early stage of the intelligence operator’s work.

Knowledge extraction from unstructured data is of course a topic of paramount importance not only for intelligence processing. For example, the topic is very relevant to contemporary search engines that aim at indexing all sorts of documents and media. To be mentioned is the Unstructured Information Management Architecture (UIMA) standard for the development, discovery, composition, and deployment of multi-modal analytics for unstructured information search technologies [5]. These efforts are in line with the processing steps involved in the “Collation” phase.

However, it would be extremely valuable for intelligence purposes to have a document management system which is able to perform batch knowledge discovery. That is, to automatically mine the data with the purpose

of aggregating, linking and relating information without a specific directive from the operator. This early form of knowledge discovery is called *structure discovery* and could provide precious “new” information to the operator as it could hint on hidden or unknown patterns of relations in the data. This abductive discovery process aims at finding the best explanation of relationships that describe the data. This batch processing would automate the “Information Structuring” phase of the Collation step.

Once structured information has been extracted from data, data mining techniques can be used to discover knowledge from data. An example for this kind of discovery is detecting patterns, associations, and correlations that occur frequently. Frequent patterns can include item sets, subsequences, and substructures. Subsequences can refer to sequential patterns (e.g. a temporal sequence of events) while the presence of substructures (e.g. graphs, trees, lattices) in the data can suggest interesting relational patterns among entities. In particular, graph mining techniques are mostly based either on the Apriori or pattern-growth algorithms [17]. The above mentioned data mining techniques work with structured data which are typically organized in databases and discover relations typically on a statistical basis according to the number of occurrences of certain patterns. Recent research efforts are focused on methods and techniques for handling inherently uncertain and compositionally structured data. In the intelligence context, all sources of information are likely to provide data affected by some degree of uncertainty. It could be measurable uncertainty as in the case of sensory data measuring some physical quantity or much less measurable as in the case of a human observer.

Statistical Relational Learning (SRL) is an emerging research area that, building on probability and statistics, aims to represent reason and learn in domains with complex relational and rich probabilistic structure [15]. SRL encompasses a number of formalisms that have been proposed over the years, from early attempts within the inductive logic programming community [25], to Machine Learning techniques such as Bayesian Networks, Markov Networks and Conditional Random Fields, to mixed (and more recent) approaches such as Markov Logic Networks [15]. The latter is an example of the emerging trend in SRL combining first order logic and Markov Networks into a new formalism.

Pre-processing of intelligence information could be an interesting application of SRL techniques given their ability to model (possibly uncertain) dependencies between related instances. Distinguishing “significant” information from “noise” in the continuous flow of input data to any reasoning system is a key step to be carried out in order to seed the interpretation process. This initial partition can be performed by matching a priori defined models. This implies that a certain number of possible “explanations” of observed data should already be available to the system as (human) expert

provided knowledge. An unsupervised probabilistic approach, on the contrary, can succeed in identifying a structural model (which can explain regularities in the data) using just a very basic form of prior knowledge if none at all. Recent studies in cognitive sciences show how achieving significant degree of success in “comprehension” needs discerning the underlying regularities in the world. This process seems to require some (inductive-abductive) constraints in order to cope with sparsity and noise in data and information [20]. According to these cognitive theories, the best the human mind can do in inferring from available data is to make the “best possible guess” guided by prior probabilities about which world structures are most likely.

A Bayesian approach seems to mimic human reasoning over structures, relations and links, and it is possible to provide a detailed computational account of how a number of basic structural forms can be inferred from various types of data (feature sets, similarity matrices, and relations). This can be applied to different areas of interest, covering higher-level problems like inferring causal structures, learning about hidden properties or objects, and interpreting the meaning of words [20]. As already mentioned, the process of deductive detection of patterns or “significant” information implies already having a model according to which data can be judged as such, that is having some strong a priori assumption over the situation under investigation. This is what is needed by logic-based approaches or graph matching algorithms used in data mining.

The algorithm proposed by Kemp and Tenenbaum’s exploits Bayesian inference to identify a hierarchical model that best accounts for the observed data and generates candidate models from graph grammars. The model with maximum posterior probability given the data [20] is taken as the most likely explanation of observed patterns. This framework allows alternative forms to compete with one another to explain any given set of data rather than requiring an a priori assumption about the form appropriate for a specific dataset. For example, the technique allows inferring structure from relational data as in the case of frequency of communications between a group of persons leading to the discovery of social cliques or hierarchical tree structures (eventually discovering lead roles within an organization). Discovered structures can dynamically be adjusted as new information is collected and filed in the database. A similar approach could be applied as a batch pre-processing to intelligence data greatly augmenting the value of the information contained in the repository as it can direct the attention of the collation operator and provide precious clues for later higher-level processing by intelligence analysts.

The support of later stages of information processing could benefit from the use of graphical models to express the probabilistic consequences of causal relationships. The scientific research community is currently

discussing whether these models could serve as the basis for learning causal relationships from data. The prospect would be to have a Bayesian learner working backwards from observed patterns of correlation (or statistical dependency) to make probabilistic inferences about the underlying causal structures likely to have generated those observed data. This process would be very similar to what is intended as creative abduction [34]. It is possible to use the basic principles of Bayesian inference over data which is represented by samples from an unknown causal graphical model and the hypotheses to be evaluated are different candidate graphical models. A brief survey of some cognitive approaches which we believe might be considered to support the automation of information fusion tasks for intelligence analysis is given in [12] [13] where also links to historical philosophical foundations are given.

5 THREAT AND RISK ASSESSMENT

There are lots of definitions of the concepts of *threat* and *risk* on, for instance, the web. They have in common that they both represent something that might happen in the future that will influence us in a negative way. One common definition, that for statistical decision theory, is found on the Wikipedia: *risk* = (Probability of an event occurring) X (Impact of the event if it occurs). For several (infinitely) possible and observable events it is, simply put, a sum (integral), over the impacts (loss function) X the probabilities (probability density function) for the observable events. Something that influences us in a positive way could on the other hand give an *opportunity* for own action. In practice, there are normally a limited set of events that can reasonably be expected and where the impact can be estimated. This is the case we discuss below.

In the military case, there is often a more or less well defined “adversary” which imposes a *threat*. Here, the threat can be formalized as a combination of the adversary’s *capability* to attack us, their *intensions* to attack, and if they can find an *opportunity* to attack us [26]. As earlier described in the paper, in a tool like CoALA a network of IntObjects and their relations are continuously built up to reflect the semantic content of a set of intelligence reports in many different “qualitative” dimensions besides the more “quantitative” time and space. Now, how could this network and the patterns emerging in it be used for forensic (history), situation (now) and threat (future) assessment? We will do this by introducing the concept of *Indicators*.

5.1 Definition and usage of an Indicator

An indicator in its most general form can be defined as something that signals (indicates) the presence of something else. Here its definition is limited to a direct observation of a maybe seemingly less relevant event or a state that can indicate something more serious (primary)

Event *has happened*, *is going on*, or *is about to happen*

State *has been realized*, *is becoming realized* or *will be realized*

Hence, the indicator is a secondary effect of the past, present or future primary event or state, simply called the *primary* below, which has not been observed directly so far. The indicator concept can be used both for detecting present or forecasting future primaries that might be threats or opportunities, as well as used in abductive reasoning [16] [49]. In the last case, indicators are regarded as consequences or effects of primaries that have already happened, as in forensic investigations. Experienced persons can often assess, or hypothesize about, what has been or is going on, or if the risk of something happening is increased, by taking notion of such indicators.

A single observation like “There are no people in the square when it is normally crowded.” can be an indicator as well as the fused result of several different observations leading to some conclusion or hypothesis like “There seem to have been repeated correlated money transfers from X to a known IED expert Y via Z’s account in bank W.” The primaries in these two cases can be a forthcoming shooting on the square, being known to the local population, and a forthcoming bomb attack, respectively.

How can we, using indicators, formalize this building of hypotheses about primaries? Imagine that a decision maker has a “*monitor list*” of primaries in the form of events or state changes that, in our asymmetric scenario, are regarded as more important than others to prevent, or exploit. They might impose plausible and extra serious threats, or positive opportunities. This list is assumed to be compiled by SMEs that in some way are familiar with the situation at large. The list, perhaps sorted according to priority or probability, might contain many primaries, more or less related to each other. Intelligence reports on the presence, or explicit absence, of indicators must now be exploited in order to somehow assess the probabilities of the different primaries in this list.

5.2 Coupling Indicators to Primaries

The couplings between indicators and a primary can for example, but not necessarily, be achieved via a Bayesian Network (BN) built by an experienced person who knows which indicators tend to influence a primary, and which indicators are more important than others and should have higher weights. A combination of indicators with different weights, and maybe also observed absence of expected indicators (Negative Information) feed into a BN, and if the output is higher than some threshold, an alert corresponding to the primary modelled by that BN as a root node is issued. A BN could be built, and be extended or modified during a mission when situation-specific knowledge grows, or several BN fragments managed separately by domain

experts could be put together to a BN tailored to match the specific mission or case [49]. There are other ways than BNs to couple the influence of indicators to primaries, but to obtain trust in the system, it must be easy to understand why a certain primary might suddenly be alerted in the system by, for instance, “clicking” on it in a Graphical User Interface. If inferences cannot be followed easily it would render the tool useless; no reasonable decision maker could take decisive decisions based on threat alerts generated by a tool whose way of functioning is regarded anything similar to a black box. Then the inference path used must be displayed in an easy-to-understand way. Furthermore, there will often be an interest to study the history of a monitored feature (primary or indicator) and see how its probability has changed in time during the inflow of intelligence reports. It is often difficult to judge about an absolute probability number, should it be, say, 0.4 or 0.6, so the rate of change can be more interesting.

Indicator weights can besides a preset importance level also be related to the frequency of similar observations as well as a preset value in the leaves of the BNs on how much a certain category of observations affects an indicator. As well, an observed indicator of a future primary must have a decay time constant associated with it depending on what it is assumed to indicate, or if the indicator itself is more of a state than an event. An indicator (explosives found) typical of a discrete upcoming event (bombing) of course decays more quickly than one (bad harvests) typical for a more permanent state (famine) and must soon enough decrease its influence on BN’s representing discrete events. What decay times to set for different types of indicators of course varies, and has to be judged by SMEs.

5.3 How to display the risk—Impactorium

So, it would be of great benefit to have a mechanism that continuously shows if the estimated risk has increased that some primary is realized. At FOI, a tool called Impactorium [46] [14] [2] [11] has been developed.¹ Impactorium has a display idea based on the so called Impact Matrix (IM). In the enterprise world, the IM has been used for risk visualisation for a long time. An IM is a 2D plot area with a “coordinate system” for the primaries where the (horizontal) X dimension represents the severity (impact or consequence) *if* it happens, and the (vertical) Y dimension the a priori probability for it to happen. We do not elaborate on the X dimension more in this paper, and the impacts have to be judged by SMEs and are normally not a subject for change in time. The primaries in the assumed monitor list mentioned above, on the other hand, are assessed concerning their probabilities using incoming intelligence reports as

¹This tool was not described in Section 3.1 above since it was not a part of the SME trial with PARANOID/CoALA and is not really a dedicated collation and link analysis tool.

sources to BNs. They should now be moved along the Y axis of the IM according to their updated probabilities. In Impactorium we have tested a slightly different visualization technique; monitored features are positioned along the Y dimension according to their a priori probabilities. When the monitored features change probability by the influence of the indicators in the BN due to incoming intelligence, their symbols in the IM are initially not shifted in position along the Y axis. Rather they automatically change colour continuously between green (improbable) and red (strong alert). When corroborating information might later be received, they can be shifted accordingly in the IM. If not received, their colour fades back to their earlier states with the earlier mentioned decay time constant. Primaries, indicators and intermediate nodes in the BN can of course also be displayed in a “monitor list,” sorted according to present probability.

The tool allows for Impactorium clients to access intelligence reports in a common database, as well as to design or use pre-designed BNs to connect indicators to primaries. Instead of BNs, simple mathematical operators like *mean*, *min* or *max* can be applied in the network nodes. Different analysts can focus on subsets of primaries by keeping them and their associated indicators in personal analysis object baskets, much like the IntObject basket in CoALA. As Impactorium is a semi-automatic tool, the operator is now able to browse the inference chains in the BNs to see details on why the alert emerged. In situations of time pressure the visualization could be done on the fly where the BN issuing the alert is visualized automatically and the most important nodes for the alert in it is highlighted.

User studies have been performed with Impactorium as well [32] [47] and the tool has since then been further developed concerning web service API and user interfaces. Impactorium is still something of a research test bench, and not yet an as well developed product as CoALA or PARANOID), but a plan for how it could be implemented in the Swedish Armed Forces has been produced.

5.4 Relating Impactorium to CoALA

As mentioned, BNs are one way to link observations or intelligence reports via indicators to potential primaries, which is the way it is done today in Impactorium. Another way would be to continuously monitor the structure of a semantic network while it is built up as in CoALA. Instead of letting one or several, maybe fused, intelligence reports trig one or several indicators (as is the case today in Impactorium), one could try to identify patterns in the CoALA network that are known beforehand to indicate threatening situations. This could be done for instance by graph matching techniques [42] [27]. Impactorium currently has a somewhat more causal event-chain analysis approach, but extending it with pattern-recognition techniques would be very interesting; this is elaborated a bit in the next

section. A maybe semi-automatic Mixed-Initiative reasoning, pattern-recognition functionality for identifying such network structures should then be the equivalent to the BNs causing certain types of observations to trig indicators in Impactorium today. How could this be achieved? Experienced people have models to which they compare a new situation they are confronted with, and to link cause and effect. Earlier experienced cases, maybe in different mixings, serve as models used to assess the type and characteristics of the new situation. This can to a large extent be compared to case-based reasoning. An implementation of this mental model-building and matching process into some algorithm, following the ideas in Section 4 of this paper, would make it possible to obtain a coupling between the outputs of a fusion level 2 tool like CoALA to the input of a level 3 risk analysis tool like Impactorium.

5.5 Other applications of Impactorium

Impactorium have so far been developed as an impact assessment tool for two cases: assessing potential *ongoing*, but still not directly observed, primaries, and potential *future* ones. An example of the latter is [47]. An example on the former is the work going on in the EU FP7 projects “Support” and “Contain” for assessing and controlling threats in sea port areas and against containers in the container logistics chain, respectively. These works are still unpublished, but use a method very similar to the idea described here: Port surveillance sensor- or container status sensor reports are tagged semantically; in this case using RDF [51] triples, describing the observed events on a semantic level appropriate for the observations done by the specific sensors. An ontology, for example described as an RDF schema, defines the event types that could be “instantiated” by the sensors. Triples expressed in XML from different sensors in the port stream into a RDF stream complex event processor [50] [22] which, from these triples and further entailments done using implicit knowledge in the ontology, builds semantic networks that describe the situations to which the sensors have contributed with observations. Predefined queries on these networks, defined using SPARQL (an RDF network query language) acts as the patterns to be searched for in the growing network. This idea is very similar to a monitoring function for networks in CoALA described in the previous subsection. When a pattern matches, perhaps within a spatial and temporal window, an alert is issued, and analysts using Impactorium can via web services subscribe on such events depending on what type of alert is relevant for their respective role (representing customs, a freight company, port security etc.).

We would also like to include “forensic” reasoning: the primary might already have happened in the *past*, and we want to reason abductively from the indicators which have followed as consequences of the hypothesis of a primary similar to [49]. This is to say that Impactorium should allow the operator to enter anywhere

in the temporal chain of intelligence reports and associated alerts of indicators and follow the inferences done by the tool.

As mentioned, primaries can be events or states. States can change discretely or more continuously. Most military actions are executed as activities together aiming at some higher goals, or *effects*. This is a central concept in the Effect-Based Approach to Operations (EBAO) paradigm, but has of course always been important in military thinking. An effect is in principle a change of state as a result of non-planned events or planned actions or activities. The way Impactorium works can help personnel that are responsible for monitoring these changes of states to do this monitoring, and when reporting on it, be more clear and concise about what causal chains are assumed to be the reason for the change. In this case the assessed outcome of executed actions is fed into the indicators of a BN or similar network that defines the influence of the success or failure of actions on the primary, here being the effect one wants to obtain. This is a mode of usage more operative than tactical and it suits threat assessments that are done by intelligence staffs rather than by the staff responsible for direct execution of military activities.

6 CONCLUSIONS

Tools like CoALA or PARANOID are accepted and appreciated by the military community. They give support for the processing and exploitation of unstructured semantic information as well as for some additional functionality analysing the established structured information set. However, up to now these interactive tools mainly just assist the human operators in their semantic exploitation of the information and their reasoning about the meaning and the consequences of the determined situation. To support situation awareness and threat and impact assessment more research on the discovery and update of behaviour pattern and system structures as well as on the principles of pattern and behaviour based reasoning, especially for imperfect data and information has to be performed. How to then alert and focus users on emerging threats and risks found accordingly, like in the Impactorium tool, is another important issue.

REFERENCES

- [1] J. R. Anderson
Cognitive Psychology and its Implications.
Palgrave Macmillan, 6th Edition, 2005.
- [2] M. Andersson et al.
The FOI C4ISR Demonstrator 2008.
In *Proceedings of the 12th International Conference on Information Fusion*, Seattle, WA, USA, July 2009, pp. 1604–1612.
- [3] W. S. Angerman
Coming Full Circle With Boyd’s Ooda Loop Ideas: An Analysis Of Innovation Diffusion And Evolution,
Thesis, AFIT/GIR/ENV/04M-01, Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio, USA, March 2004, <http://www.dict.mil/cgi-bin/GetTRDoc?AD=ADA425228>.

- [4] R. T. Antony
Principles of Data Fusion Automation,
Artech House, Inc., Norwood, MA, 1995.
- [5] Apache UIMA project
Homepage: <http://uima.apache.org/index.html>.
- [6] J. Biermann
Understanding Military Information Processing—An Approach to Supporting the Production of Intelligence in Defence and Security
in: Shahbazian, Elisa; Rogova, Galina; De Weert, Michael J. (Eds.) “Harbour Protection Through Data Fusion Technologies,” *Proceedings of the NATO Advanced Research Workshop on Data Fusion Technologies for Harbour Protection* Tallinn, Estonia 27 June–1 July 2005, Springer, 2009.
- [7] J. Biermann et al.
From Unstructured to Structured Information in Military Intelligence: Some Steps to Improve Information Fusion. NATO RTO SCI-158 Panel Symposium on “Systems, Concepts and Integration (SCI) Methods and Technologies for Defence Against Terrorism,” 25 to 27 October 2004, London, UK.
- [8] J. Biermann, P. Hörling, and L. Snidaro
Automated Support for Intelligence in Asymmetric Operations: Requirements and Experimental Results.
In *Proceedings of the Twelfth International Conference on Information Fusion*, Seattle, U.S., July 6–9, 2009, pp. 1592–1599.
- [9] J. Bohannon
Counterterrorism’s New Tool: ‘Metanetwork’ Analysis.
Science, Vol. 325, pp. 409–411, 24 July 2009.
- [10] A. Brook
UK Experiences of Using Coalition Battle Management Language.
In *15th International Symposium on Distributed Simulation and Real Time Applications (DS-RT)*, Salford, UK, 2011.
- [11] J. Brynielsson, A. Horndahl, F. Johansson, L. Kaati, C. Mårtensson, and P. Svenson
Analysis of weak signals for detecting lone wolf terrorists.
In *Proceedings of the European Intelligence and Security Informatics Conference (EISIC 2012)*, 2012, pp. 197–204.
- [12] G. Ferrin, L. Snidaro, and G. L. Foresti
Structuring Relations for Fusion in Intelligence.
In *Proceedings of the 12th International Conference on Information Fusion*, Seattle, U.S., July 6–9, 2009, pp. 1621–1626.
- [13] G. Ferrin, L. Snidaro, and G. L. Foresti
Revisiting the role of abductive inference in fusion domain.
In *Proceedings of the 13th International Conference on Information Fusion*, Edinburgh, U.K., July 26–29, 2010.
- [14] R. Forsgren, L. Kaati, C. Mårtensson, P. Svenson, and E. Tjörnhammar
An overview of the Impactorium tools 2008.
SWIFT 2008—Skövde Workshop on Information Fusion Topics, Skövde, Sweden (2008), http://www.his.se/PageFiles/6801/Forsgren_SWIFT2008.pdf.
- [15] L. Getoor and B. Taskar (Eds.)
Statistical Relational Learning, Adaptive Computation and Machine Learning.
MIT Press, Cambridge, Massachusetts, 2007.
- [16] T. Goan, N. Kartha, and R. Kaneshiro
Effective Behavioral Modeling and Prediction Even When Few Exemplars are Available.
Proc. SPIE Vol. 6235, Signal Processing, Sensor Fusion and target recognition XV, 623511 (2006).
- [17] J. Han and M. Kamber
Data Mining: Concepts and Techniques.
Morgan Kaufmann, 2006.
- [18] H. Kang, C. Plaisant, B. Lee, and B. B. Bederson
NetLens: Iterative Exploration of Content-actor Network Data.
Information Visualization, vol. 6, n. 1, 2007, pp. 18–31.
- [19] Y. Kang, C. Gorg, and J. Stasko
How can visual analytics assist investigative analysis? Design implications from an evaluation.
IEEE Transactions on Visualization and Computer Graphics, Vol. 17, no. 5, 2011, pp. 570–583.
- [20] C. Kemp and J. B. Tenenbaum
The discovery of structural form.
Proceedings of the National Academy of Sciences, Vol. 105, n. 31, pp. 10687–10692, 2008.
- [21] H. Khurana, J. Basney, M. Bakht, M. Freemon, V. Welch, and R. Butler
Palantir: A Framework for Collaborative Incident Response and Investigation.
In *Proceedings of the 8th Symposium on Identity and Trust on the Internet*, 2009, pp. 38–51.
- [22] S. C. Komazec
Sparkwave: Continuous Schema-Enhanced Pattern Matching over RDF Data Streams.
In *Proceedings of the 6th ACM International Conference on Distributed Event-Based Systems (2012)*, pp. 58–68.
- [23] K. Kruger, J. Biermann, M. Frey, and U. Schade
Automatic Information Fusion from Diverse Sources.
In *Proceedings of the Military Communications and Systems Conference (MCC) 2007*, Bonn, 2007.
- [24] K. B. Laskey, D. A. Schum, P. C. G. Costa, and Terry Janssen
Ontology of Evidence.
In *Ontology For the Intelligence Community*, 2008/12/3, p. 20–62, <http://ceur-ws.org/Vol-440/Proceedings.pdf#page=20;2013-05-05>.
- [25] N. Lavrac and S. Dzeroski
Inductive Logic Programming: Techniques and Applications.
Ellis Horwood, New York, 1994.
- [26] E. G. Little and G. L. Rogova
An Ontological Analysis of Threat and Vulnerability.
In *Proceedings of the 9th International Conference on Information Fusion*, Florence, Italy, July 2006.
- [27] E. Little, K. Sambhoos, and J. Llinas
Enhancing Graph matching Techniques with Ontologies.
In *Proceedings of the 11th International Conference on Information Fusion*, July 2008, Cologne, Germany, pp. 1890–1897.
- [28] J. Llinas, C. Bowman, G. Rogova G., A. Steinberg, E. Waltz, and F. White
Revisiting the JDL Data Fusion Model II.
In: *Proceedings of the 7th International Conference on Information Fusion*, June 2004, Stockholm, Sweden, pp. 1218–1230.
- [29] M. D. Moskal, M. Sudit, and K. Sambhoos
The role of information fusion in providing analytical rigor for intelligence analysis.
In *Proceedings of the 14th International Conference on Information Fusion*, 5–8 July 2011 pp. 1279–1284.
- [30] NATO AC/323(IST-065)TP/423 RTO-TR-IST-065 Technical Report on “Information Fusion in Asymmetric Operations,” November 2012.
- [31] NATO RTGonIFD IST-038/RTG-016
Information Fusion: Supporting Military Intelligence by Structuring Information
Proceedings of the NATO RTO IST Final Demonstration on Information Fusion For Command Support (IST-055/CDT), 10 November 2005, The Hague, The Netherlands.
- [32] M. Nilsson, J. van Laere, T. Ziemke, P. Berggren, and B. Kylesten
A user study of the Impact matrix, a fusion based decision support for enhanced situation awareness.
In *Proceedings of the 11th Int. Conference on Information Fusion*, Cologne, 2008, pp. 440–447.

- [33] North Atlantic Treaty Organization NATO Standardization Agency (NSA)
AAP-06 Edition 2013 "NATO Glossary of Terms and Definitions" (English and French).
<http://nsa.nato.int/nsa/nsdd/APdetails.html?APNo=1689&LA=EN> (15.5.2013).
- [34] C. S. Peirce
Collected Papers of Charles Sanders Peirce.
 Harvard University Press, 1931.
- [35] J. M. Pullen
 Battle Management Language Enables Rapid Integration of Command & Control with Simulation.
 In *14th International Symposium on Distributed Simulation and Real Time Applications (DS-RT)*, Fairfax, VA, USA, 2010.
- [36] K. Rein and U. Schade
 Battle Management Language as a "Lingua Franca" for situation awareness.
Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA), New Orleans, LA, USA, 2012.
- [37] K. Rein, U. Schade, and M. R. Hieb
 Battle Management Language (BML) As an Enabler.
 International Conference on Communications (ICC '09), Dresden, 2009.
- [38] K. Rein, U. Schade, and S. Kawaletz
 Uncertainty Estimation in the Fusion of Text-Based Information for Situation Awareness.
 In E. Hüllermeier, R. Kruse, F. Hoffmann (Eds.) *Information Processing and Management of Uncertainty, In Knowledge-Based Systems. Applications, Vol. 81 Communications in Computer and Information Science*, Springer Berlin Heidelberg, 2010.
- [39] PNNL. Pacific Northwest National Laboratory
 INSPIRE Visual Document Analysis.
 Last accessed: May 20, 2013, <http://in-spire.pnl.gov>.
- [40] G. Rogova and E. Bossá
 Information quality in information fusion.
 In *Proceedings of the 13th International Conference on Information Fusion*, Edinburgh, U.K., July 2010.
- [41] G. Rogova and V. Nimier
 Reliability in information fusion: literature survey.
 In *Proceedings of the 7th International Conference on Information Fusion*, ISIF, June 2004, Stockholm, Sweden, pp. 1158–1165.
- [42] K. Sambhoos, J. Llinas, and E. Little
 Graphical Methods for Real-Time Fusion and Estimation with Soft Message Data.
 In *Proceedings of the 11th International Conference on Information Fusion*, Cologne, 2008, pp. 1621–1628.
- [43] U. Schade, J. Biermann, M. Frey, and K. Kruger
 From Battle Management Language (BML) to automatic information fusion.
 In: V. Popovich, M. Schrenk, and K. Korolenko (Eds.) *Information Fusion and Geographical Information Systems (Lecture Notes in Geoinformation and Cartography)*, pp. 84–95, Berlin: Springer.
- [44] M. Shahriar Hossain, P. Butler, A. P. Boedihardjo, and N. Ramakrishnan
 Storytelling in entity networks to support intelligence analysts.
 In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2012, pp. 1375–1383.
- [45] L. Snidaro, I. Visentini, K. Bryan, and G. L. Foresti
 Markov Logic Networks for context integration and situation assessment in maritime domain.
 In *Proceedings of the 15th International Conference on Information Fusion*, Singapore, July 09–12, 2012, pp. 1534–1539.
- [46] P. Svenson, T. Berg, P. Hörling, M. Malm, and C. Mårtensson
 Using the impact matrix for predictive situational awareness.
 In *Proceedings of the 10th Int. Conference on Information Fusion*, Québec, Canada, July 2007.
- [47] P. Svenson, R. Forsgren, B. Kylesten, P. Berggren, W. Rong Fah, M. S. Choo, and J. K. Yew Hann
 Swedish-Singapore studies of Bayesian modeling techniques for tactical intelligence analysis.
 In *Proceedings of the 13th International Conference on Information Fusion*, Edinburgh, UK, July 2010.
- [48] G. Thibault, M. Gareau, and F. Le May
 Intelligence Collation in Asymmetric Conflict: A Canadian Armed Forces Perspective.
Proc. 10th International Conference on Information Fusion, Québec, Canada, July 2007.
- [49] S. P. van Gosliga and I. van de Voorde
 Hypothesis Management Framework: a flexible design pattern for belief networks in decision support systems.
 UAI 2008, Helsinki, Finland, 2008.
- [50] Wikipedia
 Complex Event Processing
http://en.wikipedia.org/wiki/Complex_event_processing.
- [51] Wikipedia
 Resource Description Framework:
http://en.wikipedia.org/wiki/Resource_Description_Framework.
- [52] E. J. Wright and K. B. Laskey
 Credibility Models for Multi-Source Fusion.
 In *Proceedings of the 9th International Conference on Information Fusion*, Florence, Italy, July 2006, pp. 274–280.

Joachim Biermann received his degree in numerical mathematics (Dipl. Math.) from the Department of Mathematics of the University of Cologne in 1982. From 1982 to 1987 he was assistant professor with the same department working in the area of numerical mathematics, especially on methods for parallel and vector computing, finite-element multigrid methods for fluid dynamic problems, and computer science.

Since 1988 he is with the Research Institute for Communication, Information Processing, and Ergonomics (FKIE) an institute of the Fraunhofer Gesellschaft and has been mainly working on knowledge-based exploitation and information fusion in the context of intelligence processing and decision support. Since 1992 he has been actively participating in several NATO activities on information fusion in intelligence, particularly in all NATO RTO IST Task Groups on Information Fusion (1999–2009). In 2003 he became deputy department head of the department of Sensor data and Information Fusion at the FKIE and leader of the research group ‘Threat Recognition for Defence and Security.’

His actual research interest is related to the analysis and modelling of heuristic information processing procedures, especially for unconventional or asymmetric threat, the integration of context and background information into multi sensor data fusion and the development and application of according fusion methods in defence and security applications.

He is a member of the Armed Forces Communications and Electronics Association (AFCEA), member of the Advisory Board of the ‘‘Information Fusion Research Programme’’ of the University of Skövde, Sweden, and a member of the Board of Directors of the International Society of Information Fusion (ISIF).



Pontus Hörling received the M.Sc. degree in physics and chemistry from Uppsala University, Uppsala, Sweden in 1988, and the Ph.D. degree in physics from the Royal Institute of Technology, Stockholm, Sweden, in 1996. Since then he has been with Swedish Defense Research Agency (FOI), the last 8 years as Senior Scientist. Here, he has mainly been involved with computer science, software development and simulation. His research interests are information processing for command and control systems, sensor- and information fusion, effects-based planning and simulation of such processes. Besides being involved in the NATO RTO task group that delivered the results reported on in this article, he has been the Swedish PoC in the NATO STO (former RTO) IST Panel since spring 2007.





Lauro Snidaro received his M.Sc. and Ph.D. in Computer Science from University of Udine in 2002 and 2006 respectively. Since 2008 he is an Assistant Professor at the Dept. of Mathematics and Computer Science, Univ. of Udine. His main interests include Data/Information Fusion, Computer Vision, Pattern Recognition, Machine Learning, Multimedia, Video understanding and annotation. He actively publishes in international journals and conferences and has co-authored more than 70 papers. He was the appointed Italian member of the NATO Task Group on Information Fusion Demonstrator (TGonIFD) IST-038/RTG-016 for 2003–05, appointed Italian member of the NATO Information Technology Panel Task Group IST-065/RTG-028 on “Information Fusion in Asymmetric Operations” (RTGonIFAO) for 2007–2009, and is the appointed Italian member of the NATO Exploratory Group on “Information Filtering and Multi-Source Information Fusion” for 2011–2013. He cooperates with international industries and research centres and is involved in several international projects on multi-sensor data fusion. He has been invited speaker at international symposiums and meetings in Armenia, Bulgaria, Spain, United States. In 2011 he was co-chair of the 16th International Conference on Image Analysis and Processing (ICIAP). He has co-organized several special sessions at the International Conference on Information Fusion and the IEEE International Conference on Advanced Video and Signal-Based Surveillance. He is special issue organizer and editorial board member of the Information Fusion journal. He serves as a computer science project reviewer for the Italian Ministry of University and Research. He regularly serves as a reviewer for 19 international journals and 18 conferences.

The Probability Generating Functional for Finite Point Processes, and Its Application to the Comparison of PHD and Intensity Filters

ROY STREIT

Probability generating functionals (PGFLs) for finite point processes are used to derive the probability hypothesis density (PHD) filter and intensity filter (iFilter) for multitarget tracking. Presenting them in a common PGFL framework makes manifest their similarities and differences. A significant difference is their measurement model—the PHD filter uses an exogenous clutter model and the iFilter uses an endogenous scattering model.

Manuscript received January 1, 2012; revised July 1, 2012 and February 14, 2013; released for publication April 14, 2013.

Refereeing of this contribution was handled by Peter Willett.

Authors' address: Metron, Inc., 1818 Library St., Reston, VA 20190

1557-6418/13/\$17.00 © 2013 JAIF

1. INTRODUCTION

Many radar and sonar sensor systems generate several point measurements at every scan. Some measurements are due to targets and others are due to clutter, or scatterers, in the sensor field of view. The multi-target tracking problem is to estimate the number of targets and their states given the measurements. The multi-hypothesis tracking (MHT) method for solving this problem is based on two widely accepted assumptions: 1) targets are points; and 2) sensors generate at most one measurement per target per scan. The second is called the “at most one measurement per target” rule. It is the cause of the intrinsically high computational complexity of optimal MHT algorithms and, consequently, the reason so many diverse kinds of alternative suboptimal algorithms are widely studied.

This paper concerns the class of multitarget tracking filters based on finite point process models for multiple target states and sensor measurement sets. Two specific kinds of filters are discussed—the PHD (probability hypothesis density) filter and the iFilter (intensity filter). Many of the differences between these filters are due to the different models of the measurement set. Contrasting these two filters in this way has the added benefit of revealing the fundamental importance of the classical methods of finite point processes for tracking applications.

Section 2 provides background on finite point processes and reviews their application to multitarget tracking filters. The next two sections are largely didactic. Section 3 defines the probability generating functional (PGFL) of a single point process. Basic results related to the PGFL are derived there. PGFLs play a central role—they characterize the probability structures underpinning the filters. Section 4 defines the bivariate PGFL of two finite point processes. The general Bayes posterior point process is defined, and its PGFL is derived from the bivariate PGFL.

Section 5 derives the PHD filter and iFilter as examples of the general Bayes posterior point process. The PHD filter uses a traditional clutter model, while the iFilter uses a scattering model. These modeling differences manifest themselves in the PGFLs of the filters, thus exposing the similarities and differences between them. Conclusions and concluding remarks are given in Section 6.

2. BACKGROUND

PGFLs for finite point processes were introduced in 1962 by Moyal [11]. In this seminal paper, Moyal noted the connection between PGFLs and probability generating functions (PGFs) of discrete random variables. He defined functional derivatives of the PGFL and used them (see (15) below) to prove that the PGFL characterizes the point process. He defined the factorial moments using PGFLs. Moyal applied his functional

calculus to stochastic population processes, establishing the connection to the classical theory of branching processes (see [1] for more background). Moyal investigated cluster processes and multiplicative processes, which are processes whose PGFL factors as in (42) below. He also studied time-dependent Markovian multiplicative population processes.

Branching processes and point process theory were studied extensively by Harris in 1963 [6]. According to the authoritative text by Daley and Vere-Jones [5, p. 1], point process theory “reached a definitive form in the now classic treatments by Moyal (1962) and Harris (1963).”

Mahler applied PGFLs to multitarget tracking problems in a series of papers; see [9] and [10] and the references therein. In this corpus he uses the FISST (*f*inite set *s*tatistics) calculus to derive the PHD filter. He introduced random finite set (RFS) models for multitarget state, as well as the idea of recursively approximating the Bayes posterior process by a Poisson point process (PPP). The term PHD was coined by Stein and Winter [14], who viewed the process of evidence accrual as additive, as opposed to multiplicative. The reformulation of the PHD using random finite sets is due to Mahler [8]. The PGFL of the Bayes posterior finite point process takes an attractive form (see (28) below). The same form was derived for the PHD tracking filter in [10, Sec. 14.8.2]; however, that result is specific to the tracking application.

An exact expression is given for the probability generating function (PGF) of the distribution of the number of points in the Bayes posterior process *before* the PPP approximation of the multitarget state. The result is a straightforward consequence of the connection between the PGF and the PGFL of the posterior process, but nonetheless it may be new. These discrete distributions provide insight into the nature of the exogenous and endogenous measurement models, as well as the PPP approximation to multitarget state.

The distinction drawn between exogenous and endogenous measurement models is perhaps new, but the use of the augmented state space, denoted below by S^+ , in tracking applications dates to at least 1986 (see [7]). (More general augmented state spaces are used by Chen, et al. [3] for dynamic clutter modeling.) The iFilter was derived by Streit and Stone [15] using a direct enumeration of measurements to targets that avoids PGFLs. Their Bayesian method is based on well-known properties of PPPs [16]. The PGFL derivation of the iFilter presented in this paper is new. The iFilter was first referred to by that name in 2010.

The relationship between medical imaging algorithms and the PHD and iFilter was first discussed in [17]. The similarity between them and the famous

Shepp-Vardi algorithm (1982) for positron emission tomography (PET) is remarkable. The relationship arises because PET uses PPP models for the image—the spatial distribution of a radioisotope, i.e., the intensity function of radioisotope decays. The connection to the classic Richardson-Lucy (1972/1974) algorithm for image restoration problems is also pointed out in [16].

3. PROBABILITY GENERATING FUNCTIONALS

The event space $\mathcal{E}(S)$ of the finite point process Ξ is the set of all ordered pairs of the form $\xi = (n, \{s_1, \dots, s_n\})$, $s_i \in S$. For $n = 0$, the event is $(0, \emptyset)$. For $n \geq 1$ the event corresponds to $n!$ equally likely, ordered events of the form $(n, s_{\sigma(1)}, \dots, s_{\sigma(n)})$, $\sigma \in \text{Sym}(n)$, where $\text{Sym}(n)$ denotes the set of all permutations of the first n positive integers. The space S can be very general, but is typically a specified subset of \mathbb{R}^d , $d \geq 1$. In physics, n is called the canonical number, the collection $\mathcal{E}_n(S)$ of all subsets of S with n points is the n th canonical ensemble, and the space $\mathcal{E}(S)$ is the grand canonical ensemble.

A functional is, in general, merely a name for an operator whose input is a function and output is a (real or complex) number. For example, definite integrals are functionals. PGFLs for general finite point processes were defined by Moyal [11, Sec. 4] as a generalization of PGFs for multivariate discrete random variables. He showed that PGFLs characterize the point process via its functional derivatives. The results presented in this section are due to Moyal. The presentation here is didactic in style and intended to be widely accessible.

3.1. Definition of the PGFL

Let Ξ be a random variable with outcomes $\xi \in \mathcal{E}(S)$. Define $\Xi = (N, X)$, where N is the canonical number and X is the set of points in the random canonical ensemble $\mathcal{E}_N(S)$. The PGFL of Ξ is defined for real-valued functions h on the state space S as

$$G^\Xi[h] = \sum_{n=0}^{\infty} p_N^\Xi(n) \int_{S^n} \left(\prod_{i=1}^n h(s_i) \right) p_{X|N}^\Xi(s_1, \dots, s_n | n) ds_1 \cdots ds_n \quad (1)$$

where $p_N^\Xi(n)$ is the distribution (probability mass function or discrete pdf) of N , and $p_{X|N}^\Xi(s_1, \dots, s_n | n)$ is the pdf of the points (s_1, \dots, s_n) conditioned on $N = n$. For $n = 0$, $p_{X|N}^\Xi(\cdot | n) = 1$ and $\prod_{i=1}^n h(s_i)$ is defined to be one. Simply put, the PGFL is the expectation of the random product $\prod_{i=1}^N h(s_i)$. The PGFL is evaluated only for functions h such that the integrals and the sum in are absolutely convergent. It is sufficient [11] to require that $|h(s)| \leq 1$ for $s \in S$. No physical units are associated with the values of $h(s)$, so the integrals in (1) are unitless and the sum is dimensionally consistent.

A finite point process Ξ is a PPP if the canonical number N is Poisson distributed with mean $\mu =$

$\int_S f^\Xi(s) ds < \infty$, where $f^\Xi(s) \geq 0$ is the intensity function, and points are independently and identically distributed in S with pdf $f^\Xi(s)/\mu$. Thus, $p_N^\Xi(n) = e^{-\mu} \mu^n / n!$ and $p_{X|N}^\Xi(s_1, \dots, s_n | n) = \mu^{-n} \prod_{i=1}^n f^\Xi(s_i)$. Direct calculation shows that

$$G^\Xi[h] = \exp \left[- \int_S f^\Xi(s) ds + \int_S h(s) f^\Xi(s) ds \right]. \quad (2)$$

The PGFL (2) is log-linear, that is, $\log(G^\Xi[h]/G^\Xi[0])$ is linear in h . For further discussion of PPPs and their applications, see [16].

3.2. Functional Derivatives of the PGFL

The finite set statistics (FISST) calculus concerns functional differentiation of PGFLs, where functional differentiation has exactly the same meaning as in the Calculus of Variations. The functional derivative of $G^\Xi[h]$ with respect to the variation w is defined by

$$\begin{aligned} \frac{\partial G^\Xi}{\partial w}[h] &= \lim_{\varepsilon \rightarrow 0^+} \frac{d}{d\varepsilon} G^\Xi[h + \varepsilon w] \\ &= \lim_{\varepsilon \rightarrow 0^+} \frac{G^\Xi[h + \varepsilon w] - G^\Xi[h]}{\varepsilon}. \end{aligned} \quad (3)$$

Here, w is a bounded real-valued function on S . (It will be specified shortly.) From (1),

$$\begin{aligned} G^\Xi[h + \varepsilon w] &= \sum_{n=0}^{\infty} p_N^\Xi(n) \int_{S^n} \prod_{i=1}^n [h(s_i) + \varepsilon w(s_i)] \\ &\quad \times p_{X|N}^\Xi(s_1, \dots, s_n | n) ds_1 \cdots ds_n. \end{aligned} \quad (4)$$

Moyal [11, Sec. 4] proves that (4) is an analytic function of ε in some open region of the complex plane containing the origin. Using (3) gives, since analyticity in ε justifies interchanging the sum and the derivative,

$$\begin{aligned} \frac{\partial G^\Xi}{\partial w}[h] &= \sum_{n=1}^{\infty} p_N^\Xi(n) \sum_{k=1}^n \int_{S^n} w(s_k) \prod_{i=1, i \neq k}^n h(s_i) \\ &\quad \times p_{X|N}^\Xi(s_1, \dots, s_n | n) ds_1 \cdots ds_n. \end{aligned} \quad (5)$$

The outermost sum starts at $n = 1$ because the derivative with respect to ε of the $n = 0$ term is zero. The innermost sum over $i \neq k$ arises from the product rule for ordinary differentiation. The Dirac delta function $\delta_x(s) \equiv \delta(s - x)$ is called an ‘‘impulse (point mass) at $s = x \in S$.’’ Specifying the variation to be $w(s) = \delta_x(s)$ gives the functional derivative

$$\begin{aligned} \frac{\partial G^\Xi}{\partial x}[h] &\equiv \frac{\partial G^\Xi}{\partial \delta_x}[h] = \left. \frac{\partial G^\Xi}{\partial w}[h] \right|_{w(\cdot) = \delta_x(\cdot)} \\ &= \sum_{n=1}^{\infty} p_N^\Xi(n) \sum_{k=1}^n \int_{S^n} \delta_x(s_k) \prod_{i=1, i \neq k}^n h(s_i) \\ &\quad \times p_{X|N}^\Xi(s_1, \dots, s_n | n) ds_1 \cdots ds_n. \end{aligned} \quad (6)$$

(Alternatively, specifying the variation to be a function in a test sequence for the delta function and taking the limit gives the same result.) Using the sampling property of the Dirac delta function, the argument symmetries of $p_{X|N}^\Xi(\cdot)$, and relabeling arguments appropriately gives

$$\begin{aligned} \frac{\partial G^\Xi}{\partial x}[h] &= \sum_{n=1}^{\infty} p_N^\Xi(n) n \int_{S^{n-1}} \prod_{i=2}^n h(s_i) \\ &\quad \times p_{X|N}^\Xi(x, s_2, \dots, s_n | n) ds_2 \cdots ds_n \end{aligned} \quad (7)$$

where the product in (7) is taken equal to one for $n = 1$. The integrals are over S^{n-1} , not S^n . Note that the derivative is a functional.

It is important to keep in mind that taking the variation w to be equal to the Dirac delta function δ_x makes the derivative $\partial G^\Xi[h]/\partial x$ depend on the point x even though $G^\Xi[h]$ itself does not. For this reason, the functional derivative (7) is referred to in this paper as the derivative with respect to an impulse at x , not simply as the derivative with respect to x .

Derivatives of the PGFL with respect to any finite number of distinct impulses extract, or decode, the pdf of Ξ from its PGFL. To find the functional derivative with respect to $x_2 \neq x_1$, start with (7) by replacing x with x_1 and h with $h + \varepsilon w$. This gives

$$\begin{aligned} \frac{\partial G^\Xi}{\partial x_1}[h + \varepsilon w] &= \sum_{n=1}^{\infty} p_N^\Xi(n) n \int_{S^{n-1}} \prod_{i=2}^n [h(s_i) + \varepsilon w(s_i)] \\ &\quad \times p_{X|N}^\Xi(x_1, s_2, \dots, s_n | n) ds_2 \cdots ds_n. \end{aligned} \quad (8)$$

Differentiating with respect to ε and setting $\varepsilon = 0$ gives

$$\begin{aligned} \frac{\partial}{\partial w} \left(\frac{\partial G^\Xi}{\partial x_1}[h] \right) &= \sum_{n=2}^{\infty} p_N^\Xi(n) n \sum_{k=2}^n \int_{S^{n-1}} w(s_k) \prod_{i=2, i \neq k}^n h(s_i) \\ &\quad \times p_{X|N}^\Xi(x_1, s_2, \dots, s_n | n) ds_2 \cdots ds_n. \end{aligned} \quad (9)$$

Substituting the variation $w(s) = \delta_{x_2}(s)$, where $x_2 \neq x_1$, and using symmetry properties of $p_{X|N}^\Xi(\cdot)$ gives the functional derivative,

$$\begin{aligned} \frac{\partial^2 G^\Xi}{\partial x_2 \partial x_1}[h] &= \sum_{n=2}^{\infty} p_N^\Xi(n) n(n-1) \int_{S^{n-2}} \prod_{i=3}^n h(s_i) \\ &\quad \times p_{X|N}^\Xi(x_1, x_2, s_3, \dots, s_n | n) ds_3 \cdots ds_n \end{aligned} \quad (10)$$

where, for $n = 2$, the product is equal to one. The integrals are now over S^{n-2} .

Functional derivatives of the PGFL with respect to the variations w_1, \dots, w_n are defined recursively as

above, or equivalently as

$$\frac{\partial^n G^\Xi}{\partial w_1 \cdots \partial w_n} [h] = \frac{\partial^n G^\Xi}{\partial \varepsilon_1 \cdots \partial \varepsilon_n} \left[h + \sum_{j=1}^n \varepsilon_j w_j \right]_{\varepsilon_1 = \cdots = \varepsilon_n = 0} \quad (11)$$

The functional derivative with respect to impulses at distinct points x_1, \dots, x_n is

$$\begin{aligned} \frac{\partial^n G^\Xi}{\partial x_1 \cdots \partial x_n} [h] &\equiv \frac{\partial^n G^\Xi}{\partial w_1 \cdots \partial w_n} [h] \Big|_{w_1 = \delta_{x_1}, \dots, w_n = \delta_{x_n}} \\ &= \sum_{k=n}^{\infty} p_N^\Xi(k) k(k-1) \cdots (k-n+1) \\ &\quad \times \int_{s^{k-n}} \left(\prod_{i=n+1}^k h(s_i) \right) \\ &\quad \times p_{X|N}^\Xi(x_1, \dots, x_n, s_{n+1}, \dots, s_k | k) ds_{n+1} \cdots ds_k \end{aligned} \quad (12)$$

where for $k = n$ the product is equal to one. The order of differentiation is immaterial. For convenience, the derivative for $n = 0$ is defined to be $G^\Xi[h]$. The derivative (12) is first order with respect to the distinct points x_1, \dots, x_n .

3.3. Event Likelihood

Evaluating (7) and (10) for $h(\cdot) \equiv 0$ gives, respectively,

$$\frac{\partial G^\Xi}{\partial x} [0] = p_N^\Xi(1) p_{X|N}^\Xi(x | N = 1) \quad (13)$$

and

$$\frac{\partial^2 G^\Xi}{\partial x_1 \partial x_2} [0] = \frac{\partial^2 G^\Xi}{\partial x_2 \partial x_1} [0] = 2! p_N^\Xi(2) p_{X|N}^\Xi(x_1, x_2 | N = 2). \quad (14)$$

In words, the derivative evaluated at $h \equiv 0$ is the pdf of the event $\xi = (1, \{x\})$, and the derivative with respect to impulses at x_1 and x_2 is the pdf of the event $\xi = (2, \{x_1, x_2\})$ or, equivalently, $2!$ times the pdf of the ordered event $(2, x_1, x_2)$. From (12), for $n \geq 1$ distinct impulses,

$$\begin{aligned} \frac{\partial^n G^\Xi}{\partial x_1 \cdots \partial x_n} [0] &= n! p_N^\Xi(n) p_{X|N}^\Xi(x_1, \dots, x_n | n) \\ &= n! p^\Xi(n, x_1, \dots, x_n) \\ &= p^\Xi(n, \{x_1, \dots, x_n\}) \end{aligned} \quad (15)$$

where $p^\Xi(n, \{x_1, \dots, x_n\})$ is the pdf of Ξ for unordered events and $p^\Xi(n, x_1, \dots, x_n)$ is pdf for the corresponding ordered event.

The derivatives (15) show that a finite point process is characterized by its PGFL. This fact is important

because it means that a finite point process can be defined by deriving its PGFL.

3.4. Factorial Moments

The first moment of Ξ is the special case of (7) with $h(s) \equiv 1$:

$$\begin{aligned} m_{[1]}^\Xi(x) &= \frac{\partial G^\Xi}{\partial x} [1] \\ &= \sum_{n=1}^{\infty} n p_N^\Xi(n) \int_{S^{n-1}} p_{X|N}^\Xi(x, s_2, \dots, s_n | n) ds_2 \cdots ds_n. \end{aligned} \quad (16)$$

For PPPs it is straightforward to verify from the PGFL (2) that the intensity function $f^\Xi(x)$ is identical to the first moment, i.e., $f^\Xi(x) = m_{[1]}^\Xi(x)$. For this reason the first moment of a finite point process is often called the intensity function.

Substituting $h(s) \equiv 1$ into (12) gives the n th factorial moment,

$$\begin{aligned} m_{[n]}^\Xi(x_1, \dots, x_n) &\equiv \frac{\partial^n G^\Xi}{\partial x_1 \cdots \partial x_n} [1] \\ &= \sum_{k=n}^{\infty} p_N^\Xi(k) k(k-1) \cdots (k-n+1) \\ &\quad \times \int_{s^{k-n}} p_{X|N}^\Xi(x_1, \dots, x_n, s_{n+1}, \dots, s_k | k) ds_{n+1} \cdots ds_k \end{aligned} \quad (17)$$

where for $k = n$ the conditional pdf is $p_{X|N}^\Xi(x_1, \dots, x_n | n)$. Factorial moments can be interpreted as multi-point intensity functions (when points are distinct with probability one). To see this, note that (17) can be written intuitively as [4, eq. (5.4.12)]

$$\begin{aligned} m_{[n]}^\Xi(x_1, \dots, x_n) dx_1 \cdots dx_n &= \Pr \left[\begin{array}{l} \text{exactly one point of the process is} \\ \text{located in each infinitesimal subset} \\ [x_i, x_i + dx_i], i = 1, \dots, n \end{array} \right]. \end{aligned} \quad (18)$$

For $n = 1$ and $n = 2$, for distinct points $x, y \in S$,

$$\begin{aligned} m_{[1]}^\Xi(x) dx &= \Pr[\text{exactly one point in } [x, x + dx)] \\ m_{[2]}^\Xi(x, y) dx dy &= \Pr[\text{exactly one point in } [x, x + dx] \\ &\quad \text{and one point in } [y, y + dy)]. \end{aligned} \quad (19)$$

For PPPs the second probability is the product of $m_{[1]}^\Xi(x) dx$ and $m_{[1]}^\Xi(y) dy$, a result that follows from

well-known independence properties of PPPs. In general, however, the second moment does not factor. The application of factorial moments in tracking applications is discussed in [2] but is outside the scope of the present paper.

3.5. Probability Generating Function of Canonical Number

The probability generating function (PGF) of N , denoted by $F^\Xi(x)$, is determined by evaluating the PGFL of Ξ for the constant function $h(s) \equiv x$. Substituting into (1) gives

$$\begin{aligned} F^\Xi(x) &\equiv G^\Xi[h]_{h(\cdot) \equiv x} \equiv G^\Xi[x] \\ &= \sum_{n=0}^{\infty} p_N^\Xi(n) x^n. \end{aligned} \quad (20)$$

In the signal processing literature, $F^\Xi(z^{-1})$ is called the z -transform of the sequence of probabilities $(p_N^\Xi(n) : n = 0, 1, \dots)$. The probability $p_N^\Xi(n)$ is

$$p_N^\Xi(n) = \frac{1}{n!} \frac{d^n F^\Xi}{dx^n}(0) \quad (21)$$

where the n th derivative with respect to x is the ordinary derivative evaluated at $x = 0$. The probability $p_N^\Xi(n)$ is $n!$ times the integral of the ordered pdf $p^\Xi(n, x_1, x_2, \dots, x_n)$ over all x_1, x_2, \dots, x_n . The first derivative of the PGF evaluated at $x = 1$ is

$$\frac{dF^\Xi}{dx}(1) = \sum_{n=0}^{\infty} p_N^\Xi(n) n x^{n-1} \Big|_{x=1} \equiv E^\Xi[N] \quad (22)$$

where $E^\Xi[N]$ is the expected number of points in a realization of Ξ .

4. BAYES POSTERIOR POINT PROCESS

In this section the conditional, or posterior, point process $\Sigma | \Upsilon$ is defined using Bayes method in terms of the bivariate process (Υ, Ξ) . The random variables are finite point processes, but this does not alter the Bayesian methodology. The PGFL of the Bayes posterior process $\Xi | \Upsilon$ and two summary statistics, namely, the intensity function and the distribution of the canonical number, are derived. Finally, Bayesian estimates are defined using the posterior point process and a specified loss function.

For tracking applications, Υ is the observation space and Ξ the multitarget state space. The points of a realization $\Upsilon = v$ are the measurements in a sensor scan. The joint pdf of the measurement and target processes is denoted by $p^{\Upsilon\Xi}(v, \xi)$, where $\Xi = \xi$ is a realization of the target process. The conditional pdf $p^{\Upsilon|\Xi}(v | \xi)$ is derived from physical models of the targets and the sensor likelihood function $p(y | s)$.

4.1. Bivariate PGFL

Let Υ be a finite point process with events $v = (m, \{y_1, \dots, y_m\}) \in \mathcal{E}(Y)$, where the space Y is in general unrelated to the space S . Extending the definition of the PGFL for Ξ to the joint process (Υ, Ξ) with events in the Cartesian product space $\mathcal{E}(Y) \times \mathcal{E}(S)$ gives the bivariate PGFL as the expectation of the product of random products $\prod_{i=1}^M g(y_i) \prod_{j=1}^N h(s_j)$, that is,

$$\begin{aligned} G^{\Upsilon\Xi}[g, h] &= \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} p_{MN}^{\Upsilon\Xi}(m, n) \\ &\quad \times \int_{Y^m} \int_{S^n} \left(\prod_{i=1}^m g(y_i) \right) \left(\prod_{j=1}^n h(s_j) \right) \\ &\quad \times p_{YX|MN}^{\Upsilon\Xi}(y_1, \dots, y_m, s_1, \dots, s_n | m, n) \\ &\quad \times dy_1 \cdots dy_m ds_1 \cdots ds_n \end{aligned} \quad (23)$$

where $p_{MN}^{\Upsilon\Xi}(\cdot)$ and $p_{YX|MN}^{\Upsilon\Xi}(\cdot)$ are the discrete and continuous pdfs associated with the joint process (Υ, Ξ) . If $m = 0$ or $n = 0$ in (23), the corresponding product is defined to be one. It is important to keep in mind that $g(\cdot)$ and $h(\cdot)$ are functions defined on Y and S , respectively.

Marginalizing the bivariate point process over one process yields the PGFL of other process. More formally,

$$G^{\Upsilon\Xi}[1, h] = G^\Xi[h] \quad \text{and} \quad G^{\Upsilon\Xi}[g, 1] = G^\Upsilon[g]. \quad (24)$$

To obtain the first expression, substitute $g(\cdot) = 1$ in (23), integrate over y_1, \dots, y_m , and sum over m . The other expression is obtained similarly.

4.2. PGFL of the Bayes Posterior Point Process

To write the PGFL of the Bayes posterior point process, note that the derivative of (23) with respect to impulses at the distinct points $\{y_1, \dots, y_m\} \subset Y$ evaluated for $g(\cdot) = 0$ is

$$\begin{aligned} &\frac{\partial^m G^{\Upsilon\Xi}}{\partial y_1 \cdots \partial y_m}[0, h] \\ &= m! \sum_{n=0}^{\infty} p_{MN}^{\Upsilon\Xi}(m, n) \int_{S^n} \left(\prod_{j=1}^n h(s_j) \right) \\ &\quad \times p_{YX|MN}^{\Upsilon\Xi}(y_1, \dots, y_m, s_1, \dots, s_n | m, n) ds_1 \cdots ds_n. \end{aligned} \quad (25)$$

Evaluating the derivative of $G^{\Upsilon\Xi}[g, 1] = G^\Upsilon[g]$ with respect to impulses at y_1, \dots, y_m for $g(\cdot) = 0$ gives

$$\frac{\partial^m G^{\Upsilon\Xi}}{\partial y_1 \cdots \partial y_m}[0, 1] = \frac{\partial^m G^\Upsilon}{\partial y_1 \cdots \partial y_m}[0] = m! p_{MY}^\Upsilon(m, y_1, \dots, y_m). \quad (26)$$

Dividing (25) by (26) yields

$$\begin{aligned}
& \frac{\frac{\partial^m G^{\Upsilon \Xi}}{\partial y_1 \cdots \partial y_m}[0, h]}{\frac{\partial^m G^{\Upsilon \Xi}}{\partial y_1 \cdots \partial y_m}[0, 1]} \\
&= \sum_{n=0}^{\infty} \int_{S^n} \left(\prod_{j=1}^n h(s_j) \right) \\
&\quad \times \frac{P_{MYNX}^{\Upsilon \Xi}(m, y_1, \dots, y_m, n, s_1, \dots, s_n)}{P_{MY}^{\Upsilon}(m, y_1, \dots, y_m)} ds_1 \cdots ds_n \\
&= \sum_{n=0}^{\infty} \int_{S^n} \left(\prod_{j=1}^n h(s_j) \right) \\
&\quad \times P_{NX|MY}^{\Xi}(n, s_1, \dots, s_n | m, y_1, \dots, y_m) ds_1 \cdots ds_n.
\end{aligned} \tag{27}$$

Substituting the Bayes factorization

$$\begin{aligned}
& P_{NX|MY}^{\Xi}(n, s_1, \dots, s_n | m, y_1, \dots, y_m) \\
&= P_{N|MY}^{\Xi}(n | m, y_1, \dots, y_m) P_{X|NMY}^{\Xi}(s_1, \dots, s_n | m, y_1, \dots, y_m)
\end{aligned}$$

into (27) and comparing the result with definition (1) shows that the ratio is the PGFL of the Bayes posterior process, that is,

$$G^{\Xi|\Upsilon}[h | y_1, \dots, y_m] = \frac{\frac{\partial^m G^{\Upsilon \Xi}}{\partial y_1 \cdots \partial y_m}[0, h]}{\frac{\partial^m G^{\Upsilon \Xi}}{\partial y_1 \cdots \partial y_m}[0, 1]}. \tag{28}$$

The PGFL (28) is valid for general finite point processes. The denominator of (28) is a constant that scales the PGFL of the numerator.

The probability structure of the Bayes posterior process is characterized by the functional derivatives of (28) with respect to impulses at the distinct points $\{x_1, \dots, x_n\} \subset S$. A specialized version of (28) for multitarget tracking applications is derived in [10, p. 757], where it is described as the PGFL ‘‘form of the multi-target corrector.’’

4.3. Summary Statistics of the Bayes Posterior Process

Since the event space $\mathcal{E}(S)$ is very large, it is useful to provide summary statistics of the posterior process $\Xi | \Upsilon$. Two statistics are of interest here. The first is the intensity function $f^{\Xi|\Upsilon}(x)$ of $\Xi | \Upsilon$. It is found by the evaluating at $h(s) = 1$ the functional derivative of (28) with respect to an impulse at $x \in S$:

$$f^{\Xi|\Upsilon}(x) = \frac{\frac{\partial^{m+1} G^{\Upsilon \Xi}}{\partial y_1 \cdots \partial y_m \partial x}[0, 1]}{\frac{\partial^m G^{\Upsilon \Xi}}{\partial y_1 \cdots \partial y_m}[0, 1]}, \quad x \in S. \tag{29}$$

The expression (29) holds for general finite point processes.

The other summary statistic is the distribution of $N^{\Xi|\Upsilon}$, the canonical number of points in the Bayes posterior process. The PGF of $N^{\Xi|\Upsilon}$ is the PGFL (28) evaluated for the constant function $h(s) = x$; that is,

$$F^{\Xi|\Upsilon}(x) = \frac{\frac{\partial^m G^{\Upsilon \Xi}}{\partial y_1 \cdots \partial y_m}[0, h] \Big|_{h(\cdot) \equiv x}}{\frac{\partial^m G^{\Upsilon \Xi}}{\partial y_1 \cdots \partial y_m}[0, 1]}. \tag{30}$$

The posterior pdf of the canonical number is, using (21),

$$\begin{aligned}
P_N^{\Xi|\Upsilon}(n) &= \frac{1}{n!} \frac{d^n}{dx^n} F^{\Xi|\Upsilon}(0) \\
&= \frac{1}{n!} \frac{d^n}{dx^n} \left(\frac{\frac{\partial^m G^{\Upsilon \Xi}}{\partial y_1 \cdots \partial y_m}[0, h] \Big|_{h(\cdot) \equiv x}}{\frac{\partial^m G^{\Upsilon \Xi}}{\partial y_1 \cdots \partial y_m}[0, 1]} \right)_{x=0}.
\end{aligned} \tag{31}$$

From (22), the expected number of targets is $E[N^{\Xi|\Upsilon}] = (d/dx)F^{\Xi|\Upsilon}(1)$.

4.4. Bayesian and Pseudo-MAP Estimators

A Bayesian estimate of Ξ is determined using a specified loss function $L(\zeta | \xi)$. This function gives the loss associated with choosing the estimate $\zeta \in \mathcal{E}(S)$ for Ξ when the true realization is $\xi \in \mathcal{E}(S)$. The Bayes loss of ζ is the expected loss, $R(\zeta) \equiv E[L(\zeta | \xi)]$, where the expectation (see (1)) is with respect to the density $p^{\Xi|\Upsilon}(\xi | \nu)$ of the Bayes posterior process $\Xi | \Upsilon$. The Bayes estimate $\hat{\xi}_{\text{Bayes}} \in \mathcal{E}(S)$ for Ξ minimizes the Bayes loss:

$$\hat{\xi}_{\text{Bayes}} = \arg \min_{\zeta \in \mathcal{E}(S)} R(\zeta). \tag{32}$$

The Bayes estimate depends on the choice of the loss function $L(\zeta | \xi)$.

In many problems, $L(\zeta | \xi)$ can be specified so that the Bayes estimate reduces to the maximum a posteriori (MAP) estimate, $\arg \max_{\xi} p^{\Xi|\Upsilon}(\xi | \nu)$. However, the MAP estimate is undefined for the posterior pdf $p^{\Xi|\Upsilon}(\xi | \nu)$. To see this, it is only necessary to observe that $p^{\Xi|\Upsilon}(\xi_1 | \nu)$ and $p^{\Xi|\Upsilon}(\xi_2 | \nu)$ have different units when the realizations ξ_1 and ξ_2 have different numbers of points.

Pseudo-MAP estimates can be defined using the posterior distribution of the canonical number and intensity functions, or other summary statistics. The consistency of such estimates must be verified on a case by case basis. These topics are outside the scope of the paper.

4.5. Branching Process Form of the Bivariate PGFL

The fundamental Bayesian paradigm of traditional single target tracking is: The observation process is Y , the object/target process is X , the conditional process $Y | X$ is the connection between them, and Bayes Theorem is used to inference the Bayesian process $X | Y$. The

same approach is adopted here, but with X replaced by Ξ and Y by Υ . This changes the mathematical details, but not the Bayesian paradigm.

Note that the Bayes posterior point process—its pdf, intensity, and canonical distribution—were obtained above without reference to the conditional measurement point process $\Upsilon | \Xi$. The conditional process is fundamental to the traditional understanding of how measurement processes are exploited to make inferences about target processes. This section follows the traditional Bayesian paradigm, but uses point process models. The approach to Bayesian tracking problems using point processes was first discussed in [9].

The bivariate PGFL is written in a branching process form, that is, as the composition of two functionals. This form lies at the root of many diverse applications, including the tracking applications discussed in the succeeding sections of the paper. Population and branching processes are an established part of the classic literature of probability [1] and [6].

From Bayes theorem,

$$p_{MY|NX}^{\Upsilon\Xi}(\cdot) \equiv p_{NX}^{\Xi}(\cdot)p_{MY|NX}^{\Upsilon|\Xi}(\cdot) = p_N^{\Xi}(\cdot)p_{X|N}^{\Xi}(\cdot)p_{MY|NX}^{\Upsilon|\Xi}(\cdot). \quad (33)$$

Substituting this into (23) and interchanging the order of sums and integrals (justified by absolute convergence) gives

$$G^{\Upsilon\Xi}[g, h] = \sum_{n=0}^{\infty} p_N^{\Xi}(n) \int_{S^n} \left(\prod_{i=1}^n h(s_i) \right) G^{\Upsilon|\Xi}[g | s_1, \dots, s_n] \times p_{X|N}^{\Xi}(s_1, \dots, s_n | n) ds_1 \cdots ds_n \quad (34)$$

where

$$G^{\Upsilon|\Xi}[g | s_1, \dots, s_n] = \sum_{m=0}^{\infty} \int_{Y^m} \left(\prod_{j=1}^m g(y_j) \right) \times p_{MY|NX}^{\Upsilon|\Xi}(m, y_1, \dots, y_m | n, s_1, \dots, s_n) dy_1 \cdots dy_m \quad (35)$$

is the PGFL of the conditional process $\Upsilon | \Xi$, as is seen by using the Bayes factorization $p_{MY|NX}^{\Upsilon|\Xi}(\cdot) = p_{M|NX}^{\Upsilon|\Xi}(\cdot) p_{Y|MNX}^{\Upsilon|\Xi}(\cdot)$.

The pdf of the conditional process $\Upsilon | \Xi$ is found by functional differentiation (cf. (15)):

$$p^{\Upsilon|\Xi}(v | n, s_1, \dots, s_n) = \frac{\partial^m}{\partial y_1 \cdots \partial y_m} G^{\Upsilon|\Xi}[0 | s_1, \dots, s_n] \quad (36)$$

where $v = (m, \{y_1, \dots, y_m\})$. Evaluating the functional derivatives (36) reveals the detailed structure of the likelihood function, including any enumerations that are inherent in the conditional process $\Upsilon | \Xi$. A target tracking example is given in the Appendix.

The expression (34) simplifies greatly if the conditional PGFL factors, that is, if it corresponds to the superposition of conditionally independent measurement processes. If it does, then

$$G^{\Upsilon|\Xi}[g | s_1, \dots, s_n] = \prod_{i=1}^n T[g | s_i] \quad (37)$$

where $T[g | s]$, $s \in S$, is a specified functional. Substituting (37) into (34) gives the branching process form of the bivariate PGFL:

$$G^{\Upsilon\Xi}[g, h] = G^{\Xi}[hT[g | \cdot]] \quad (38)$$

where $G^{\Xi}[h]$ denotes the PGFL of Ξ . This is the branching process form of the PGFL. It is central to multitarget tracking applications.

5. SINGLE SENSOR MULTITARGET TRACKING

As noted above, the PGFL characterizes the finite point process. Therefore, in applications, the formulation of the PGFL is of first importance. The filters in this section and the next are derived directly from the relevant PGFL.

The PGFL approach is applied to two measurement models for multitarget tracking. Both models are consistent with the “at most one measurement per target” rule. The PHD filter uses an exogenous clutter model, that is, clutter arises spontaneously in the sensor measurement space and is superposed with the target measurement processes. This is a natural model if the outputs of the sensor signal processor are thresholded to produce point measurements that are fed to a post-processor. It is a standard model that is widely accepted in the tracking literature. The exogenous clutter process is assumed to be a PPP on S with intensity $\lambda(x)$, $x \in S$.

In contrast, the iFilter uses an endogenous model in which all measurements are attributed to scatterers in the augmented target state space $S^+ \equiv S \cup \phi$ defined below in Section 5.2. A target is a scatterer whose state is in S , and a clutter measurement corresponds to a scatterer whose state in S is unknown, i.e., it is ϕ . This is a natural model for sensor signal processors when distinctions between scatterers are not drawn. It is relatively unused in the tracking community. The different models lead to remarkably similar filters.

The endogenous and exogenous models are mathematically compatible; that is, they can be used together in the same problem. What is mathematically possible, however, must also be justified in the application. This possibility is not explored further in this paper.

The PGFL of the superposition of mutually independent finite point processes is the product of their PGFLs (see [5, Prop. 9.4.1.IX]). This property makes the PGFL useful in many problems involving enumeration, since crucial questions often revolve around learning which process gave rise to which point in a superposition of points—this is the measurement to target assignment problem.

The target process Ξ is interpreted throughout the paper as the point process that is predicted to the current time from the previous time step. Prediction involves independent thinning (see [16, Sec. 2.8]) and Markovian target motion, neither of which alters the character of the target process model—if the target process is a PPP at the previous time step, the predicted process is a PPP. The essential differences between the PHD filter and the iFilter are due the different measurement models, not the predicted target process.

For concreteness, denote the target process at the previous time step by Ξ^- . It is defined on S for exogenous models and on S^+ for endogenous models. For the filters considered here, Ξ^- is assumed to be a PPP to close the Bayesian recursion. Denote its intensity function by $f^{\Xi^-}(\cdot)$. Target motion from the previous time step to the current one is assumed to be Markovian. For exogenous models, denote the transition (motion) model by $F(x | x^-)$, $x^-, x \in S$. Thus, $\int_S F(x | x^-) dx = 1$ for all $x^- \in S$. Let $d(x)$ denote the probability that a target at x does not survive to the next time step, and let $B(x)$ denote the intensity of a new target PPP at x in the current time step. The predicted process Ξ is the process Ξ^- after it is thinned by $d(x)$ and transformed by $F(x | x^-)$, and with new targets superimposed. The process Ξ is a PPP on S , and its intensity is

$$f^{\Xi}(x) = B(x) + \int_S F(x | x^-)(1 - d(x^-))f^{\Xi^-}(x^-)dx^- \quad (39)$$

The result can be derived in several ways (see, e.g., [16]), but none are given here.

The motion model for endogenous models is denoted by $\Psi(x | x^-)$, $x^-, x \in S^+$. Integrals over the augmented space are defined as in (56) below. Thus,

$$1 = \int_{S^+} \Psi(x | x^-)dx \equiv \Psi(\phi | x^-) + \int_S \Psi(x | x^-)dx$$

for all $x^- \in S^+$. The survival probability $d(x)$ is defined for all $x \in S^+$. For convenience, $d(x)$ is assumed to be the same as for the exogenous model for $x \in S$, but the probability $d(\phi)$ is new and must be specified. After thinning and transformation by $\Psi(x | x^-)$, the predicted process Ξ is a PPP on S^+ , and its intensity is

$$\begin{aligned} f^{\Xi}(x) &= \int_{S^+} \Psi(x | x^-)(1 - d(x^-))f^{\Xi^-}(x^-)dx^- \\ &\equiv \Psi(x | \phi)(1 - d(\phi))f^{\Xi^-}(\phi) \\ &\quad + \int_S \Psi(x | x^-)(1 - d(x^-))f^{\Xi^-}(x^-)dx^-, \quad x \in S^+. \end{aligned} \quad (40)$$

This result can be derived in the same manner as (39).

The transition model $\Psi(x | x^-)$ and intensity $f^{\Xi^-}(\phi)$ can be chosen to match (39) on the subspace S of S^+ . To do this, let $d(\phi) = 0$ and set $f^{\Xi^-}(\phi) = \mu + \lambda$,

where $\mu = \int_S B(x)dx$ and $\lambda = \int_Y \lambda(y)dy$, respectively, are the expected numbers of new targets and clutter (see Assumption 2 in Section 5.1 below) in the exogenous model. The transition function Ψ on S^+ is defined in terms of the parameters of the exogenous model via the partitioned matrix

$$\begin{aligned} &\left[\begin{array}{c|c} \Psi(x | x^-) & \Psi(\phi | x^-) \\ \hline \Psi(x | \phi) & \Psi(\phi | \phi) \end{array} \right] \\ &= \left[\begin{array}{c|c} F(x | x^-) & 0 \\ \hline B(x)/(\mu + \lambda) & \lambda/(\mu + \lambda) \end{array} \right] \quad \text{for } x, x^- \in S. \end{aligned} \quad (41)$$

For these choices, and complementary ones for the measurement likelihood function (see the paragraph following (63) below), (40) reduces to (39) on S .

5.1. Exogenous Clutter—the PHD Filter

The set of target states and the set of sensor measurements are modeled as finite point processes Ξ and Υ with outcomes $\xi \in \mathcal{E}(S)$ and $v \in \mathcal{E}(Y)$, respectively, where the target state space is S and sensor measurement space is Y . The PGFL of $\Upsilon | \Xi$ is obtained under three assumptions:

1. The target process Ξ is a PPP on S with intensity function $f^{\Xi}(s)$.
2. Conditioned on the event $\Xi = \xi = (n, \{s_1, \dots, s_n\})$, the measurement process is the superposition of n mutually independent, identical, target-originated measurement processes and a PPP clutter process on Y with intensity function $\lambda(y)$ that is independent of targets.
3. No target generates more than one measurement in the space Y .

The exogenous clutter model is part of the second assumption. The target-originated measurement processes in the second assumption are finite point processes on Y . It leads to the factorization (42) below. The third assumption says that the target-originated processes have at most one point.

Bivariate PGFL for the PHD Filter

Assumptions 1–3 lead directly to the factored form (44) of the PGFL of the conditional process. The measurement process Υ is the superposition of a clutter process $\Upsilon_{\text{clutter}}$ and an unknown number of identical target-originated measurement processes Υ_{target} . Conditioned on $\Xi = \xi = (n, \{s_1, \dots, s_n\})$, there are $N = n$ target-originated measurement processes. These n processes and the clutter process are conditionally independent, so the PGFL of the conditional process $\Upsilon | \Xi$ is the product of their individual PGFLs. Thus, for real-valued functions g defined on Y ,

$$G^{\Upsilon | \Xi}[g | s_1, \dots, s_n] = G^{\Upsilon_{\text{clutter}}}[g] \prod_{i=1}^n G^{\Upsilon_{\text{target}}}[g | s_i] \quad (42)$$

where the product is taken equal to one if $n = 0$.

The form of $G^{\Upsilon_{\text{target}}}[g | s]$ is derived from the assumption that a target at $s \in S$ generates at most one measurement. It is also assumed that a target at s is detected with probability $P^D(s)$ and not detected with probability $1 - P^D(s)$. Then, using (1) directly, the PGFL of target-originated measurement is a two term sum,

$$G^{\Upsilon_{\text{target}}}[g | s] = 1 - P^D(s) + P^D(s) \int_Y g(y)p(y | s)dy \quad (43)$$

where $p(y | s)$ is the (assumed known) sensor pdf of the point measurement $y \in Y$ given a target at $s \in S$. The clutter PGFL is a PPP on Y with intensity function $\lambda(y)$, so its PGFL is the same form as (2). The measurement set is the superposition of independent processes, by Assumption 2. Substituting (43) and the clutter PGFL into (42) gives

$$\begin{aligned} G^{\Upsilon|\Xi}[g | s_1, \dots, s_n] \\ &= \exp \left[- \int_Y \lambda(y)dy + \int_Y g(y)\lambda(y)dy \right] \\ &\quad \times \prod_{i=1}^n \left(1 - P^D(s_i) + P^D(s_i) \int_Y g(y)p(y | s_i)dy \right). \end{aligned} \quad (44)$$

In words, (44) says that the measurement set is the outcome of an insertion/deletion process, or “indel” process for short, because PPP clutter is randomly inserted and target-originated measurements are randomly deleted.

Explicit expressions for the pdf of the conditional process $\Upsilon | \Xi$ are not needed. These expressions are, however, very insightful because to write them down is to enumerate the measurement to target assignments. Examples are given in the Appendix.

Substituting (42) into (34) gives

$$\begin{aligned} G^{\Upsilon|\Xi}[g, h] &= G^{\Upsilon_{\text{clutter}}}[g] \sum_{n=0}^{\infty} p_N^{\Xi}(n) \int_{S^n} \prod_{i=1}^n (h(s_i)G^{\Upsilon_{\text{target}}}[g | s_i]) \\ &\quad \times P_{X|N}^{\Xi}(s_1, \dots, s_n | n) ds_1 \cdots ds_n. \end{aligned} \quad (45)$$

The sum-integral in (45) is equal to the PGFL of the target process Ξ evaluated at the function $h(s)G^{\Upsilon_{\text{target}}}[g | s]$. By assumption, the PGFL of Ξ is a PPP with intensity f^{Ξ} and its PGFL is given by (2). Substituting $h(s)G^{\Upsilon_{\text{target}}}[g | s]$ into (2), and then substituting the

PGFLs for target (43) and for PPP clutter, yields the PGFL of the joint process:

$$\begin{aligned} G^{\Upsilon|\Xi}[g, h] &= \exp \left[- \int_Y \lambda(y)dy + \int_Y g(y)\lambda(y)dy \right. \\ &\quad - \int_S f^{\Xi}(s)ds + \int_S h(s)f^{\Xi}(s)ds \\ &\quad - \int_S h(s)P^D(s)f^{\Xi}(s)ds \\ &\quad \left. + \int_S \int_Y g(y)h(s)p(y | s)P^D(s)f^{\Xi}(s)dy ds \right]. \end{aligned} \quad (46)$$

Except for clutter, this PGFL is a special case of the general result (38).

First Summary Statistic—the Target Intensity Function

The intensity function of the Bayes posterior process $\Xi | \Upsilon = v = (m, y_1, \dots, y_m)$ is obtained by substituting the functional derivatives of (46) with respect to impulses at the measurement points y_1, \dots, y_m into (29). The derivative of $G^{\Upsilon|\Xi}[g, h]$ with respect to an impulse at $y_1 \in Y$ is

$$\begin{aligned} \frac{\partial G^{\Upsilon|\Xi}}{\partial y_1}[g, h] \\ &= G^{\Upsilon|\Xi}[g, h] \left\{ \lambda(y_1) + \int_S h(s)p(y_1 | s)P^D(s)f^{\Xi}(s)ds \right\}. \end{aligned} \quad (47)$$

The term in braces in (47) does not depend on $g(y)$, so its functional derivative with respect to an impulse at $y \neq y_1$ is zero. Thus, for distinct impulses at y_1, \dots, y_m in Y ,

$$\begin{aligned} \frac{\partial^m G^{\Upsilon|\Xi}}{\partial y_1 \cdots \partial y_m}[g, h] &= \\ G^{\Upsilon|\Xi}[g, h] \prod_{i=1}^m \left(\lambda(y_i) + \int_S h(s)p(y_i | s)P^D(s)f^{\Xi}(s)ds \right). \end{aligned} \quad (48)$$

The functional derivative of (48) with respect to an impulse at $x \in S$ is

$$\begin{aligned} \frac{\partial^{m+1} G^{\Upsilon|\Xi}}{\partial x \partial y_1 \cdots \partial y_m}[g, h] &= G^{\Upsilon|\Xi}[g, h] f^{\Xi}(x) \left(1 - P^D(x) + P^D(x) \int_Y g(y)p(y | x)dy \right) \\ &\quad \times \prod_{i=1}^m \left(\lambda(y_i) + \int_S h(s)p(y_i | s)P^D(s)f^{\Xi}(s)ds \right) \\ &\quad + G^{\Upsilon|\Xi}[g, h] P^D(x) f^{\Xi}(x) \sum_{i=1}^m p(y_i | x) \prod_{k=1, k \neq i}^m \left(\lambda(y_k) + \int_S h(s)p(y_k | s)P^D(s)f^{\Xi}(s)ds \right). \end{aligned} \quad (49)$$

Substituting $g(y) = 0$ and $h(x) = 1$ gives the unconditional pdf

$$p^\Upsilon(v) = \frac{\partial^m G^{\Upsilon\Xi}}{\partial y_1 \cdots \partial y_m} [0, 1] \\ = G^{\Upsilon\Xi} [0, 1] \prod_{i=1}^m \left(\lambda(y_i) + \int_S p(y_i | s) P^D(s) f^\Xi(s) ds \right) \quad (50)$$

and

$$\frac{\partial^{m+1} G^{\Upsilon\Xi}}{\partial x \partial y_1 \cdots \partial y_m} [0, 1] = G^{\Upsilon\Xi} [0, 1] f^\Xi(x) \left\{ (1 - P^D(x)) \prod_{i=1}^m \left(\lambda(y_i) + \int_S p(y_i | s) P^D(s) f^\Xi(s) ds \right) \right. \\ \left. + P^D(x) \sum_{i=1}^m p(y_i | x) \prod_{k=1, k \neq i}^m \left(\lambda(y_k) + \int_S p(y_k | s) P^D(s) f^\Xi(s) ds \right) \right\}. \quad (51)$$

Substituting (51) and (50) into (29) gives the intensity function, that is, the PHD:

$$f_{\text{PHD}}^{\Xi|\Upsilon}(x) = f^\Xi(x) \left[1 - P^D(x) + P^D(x) \sum_{i=1}^m \frac{p(y_i | x)}{\lambda(y_i) + \int_S p(y_i | s) P^D(s) f^\Xi(s) ds} \right]. \quad (52)$$

The expected number of targets in S is

$$\hat{N}_{\text{PHD}}(S) = \int_S f_{\text{PHD}}^{\Xi|\Upsilon}(x) dx. \quad (53)$$

The number $\hat{N}_{\text{PHD}}(S)$ is an expectation over the grand canonical ensemble and is typically not an integer.

Second Summary Statistic—Distribution of Target Canonical Number

The Bayes posterior target process $\Xi | (\Upsilon = v)$ is a PPP only for $m = 0$. This follows from the form of the PGF of $N^{\Xi|\Upsilon}$,

$$F_{\text{PHD}}^{\Xi|\Upsilon}(x) = \exp \left[(x - 1) \int_S (1 - P^D(s)) f^\Xi(s) ds \right] \\ \times \prod_{i=1}^m \frac{\lambda(y_i) + x \int_S p(y_i | s) P^D(s) f^\Xi(s) ds}{\lambda(y_i) + \int_S p(y_i | s) P^D(s) f^\Xi(s) ds} \quad (54)$$

where the product is equal to one for $m = 0$. This expression is obtained by substituting (48) into the general result (30). The PGF (54) is the product of two PGFs. One is the PGF of the number of clutter points, which is Poisson distributed with mean $\int_S (1 - P^D(s)) f^\Xi(s) ds$. The other is the PGF of the number of heads, i.e., targets detected, in a coin tossing experiment. The experiment uses m non-identical coins, each tossed only once, where the probability of a target detection for the i th coin is

$$\text{Pr}[\text{target detection} | \text{coin } i] \\ = \frac{\int_S p(y_i | s) P^D(s) f^\Xi(s) ds}{\lambda(y_i) + \int_S p(y_i | s) P^D(s) f^\Xi(s) ds}. \quad (55)$$

The distribution $p_N^{\Xi|\Upsilon}(n)$ of canonical number is obtained by differentiating $F^{\Xi|\Upsilon}(x)$ with respect to x . The mean number of targets is the sum of the means of the factors in (54), and this is clearly identical to (53).

The PGF (54) of the Bayes canonical number distribution is an immediate consequence of the connection between PGFs and PGFLs. Nonetheless, it appears to be new to the PHD literature.

The PHD filter approximates the Bayes posterior process by a PPP with intensity (52). The PGF of the Bayes posterior distribution of canonical number is approximated by the Poisson distribution whose mean is (53). The mean canonical numbers of the Bayes posterior process and the PPP approximation are equal, but the distributions are mismatched.

Post-processing decision procedures estimate the actual number of targets, decide which measurements correspond to targets and which to clutter, and compute point estimates and areas of uncertainty (AOUs) for detected targets. The point estimates are interpreted as pseudo-MAP estimates of target states, as discussed in Section 3, and the AOUs are surrogates for error covariance matrices. These important topics are outside the scope of the present paper.

5.2. Endogenous Scattering—the iFilter

An endogenous measurement model is a model of the spatial distribution of scatterers. It makes no distinction between scatterers that are targets and those that are clutter; such distinctions are relegated to a post-processing classification decision procedure. The predicted target process Ξ is defined on S^+ , and its intensity is given by (40).

To compare the endogenous measurement model to the standard exogenous model, interpret a scatterer with state $s \in S$ as a target in the same state. Scatterers whose state is ϕ are clutter-generators in the exogenous model. Because Ξ is a PPP on S^+ , more than one point in the realization $\xi = (n, \{s_1, \dots, s_n\})$ can be equal to ϕ . Thus, the model accommodates a variable number of clutter

points by varying the number of scatterers with state ϕ . Assumptions 1–3 above are modified as follows:

1'. The scattering process Ξ is a PPP on $S^+ = S \cup \phi$ with intensity function $f^\Xi(s)$, $s \in S^+$, where the state ϕ is assigned to scatterers whose state $s \in S$ is unknown.

2'. Conditioned on the scattering event $\Xi = \xi = (n, \{s_1, \dots, s_n\})$, the measurement process is the superposition of n mutually independent identical scatterer-originated measurement processes.

3'. No scatterer generates more than one measurement in the space Y .

The third assumption says that a scatterer-originated process generates at most one measurement regardless of the scatterer state. It leads to the two term expression (57) below. The scatterer-originated measurement processes in the second assumption are finite point processes on Y . It leads to the factorization (58) below.

Bivariate PGFL for the iFilter

The Markovian transition model that determines the predicted intensity on S^+ is an essential ingredient of the ability of the iFilter to distinguish between scatterers in state $s \in S$ and those in state ϕ . Transition models on S^+ incorporate within themselves models for initiation and termination of tracks; however, the details are not discussed in this paper to focus attention on the PGFL aspects of the point process theory. It is also necessary that the temporal stability of measurements from scatterers in state $s \in S$ (i.e., the targets) be greater than that of measurements from scatterers in state ϕ (i.e., the clutter). This depends on the character of the data. For further discussion of iFilters and an application to real data, see [12] and [13].

The sensor likelihood function $p(y | s)$, detection probability $P^D(s)$, and intensity function $f^\Xi(s)$ are extended from S to S^+ . The density $p(y | \phi)$ is interpreted as the pdf of y given that it arises from a scatterer with state ϕ ; the probability of detection $P^D(\phi)$ is the probability that a scatterer with state ϕ generates a measurement; and the intensity $f^\Xi(\phi)$ is the expected number of scatterers with state ϕ . The number $f^\Xi(\phi)$ is dimensionless. Integrals over S^+ are defined for real-valued functions $h(s)$, $s \in S^+$, by

$$\int_{S^+} h(s) ds \equiv h(\phi) + \int_S h(s) ds. \quad (56)$$

This definition is used in the PGFL. Functional derivatives extend to the space S^+ by defining the Dirac delta function so that $\delta_\phi(\phi) = 1$ and $\delta_x(\phi) = \delta_\phi(s) = 0$ for $s, x \in S$. Repeated functional derivatives with respect to impulses at ϕ are used to show that the PGFL characterizes the point process; however, as will be seen, the iFilter requires only one such derivative. Further details are straightforward and are omitted.

Detected scatterers contribute measurements to the measurement set, but undetected scatterers do not. The

PGFL of a scatterer is

$$G^{\Upsilon \text{scatter}}[g | s] = 1 - P^D(s) + P^D(s) \int_Y g(y) p(y | s) dy, \quad s \in S^+ \quad (57)$$

an expression identical to (43) except that it holds on S^+ . The PGFL of the measurement set is, from the conditional independence assumptions,

$$G^{\Upsilon \Xi}[g | s_1, \dots, s_n] = \prod_{i=1}^n \left(1 - P^D(s_i) + P^D(s_i) \int_Y g(y) p(y | s_i) dy \right). \quad (58)$$

A separate clutter model is not used because such points are modeled as scatterers whose state is $\phi \in S^+$. The joint PGFL is obtained using the conditional process (58) in the same manner as before (see (45)). The result is

$$G^{\Upsilon \Xi}[g, h] = \exp \left[- \int_{S^+} f^\Xi(s) ds + \int_{S^+} h(s) f^\Xi(s) ds - \int_{S^+} h(s) P^D(s) f^\Xi(s) ds + \int_{S^+} \int_Y g(y) h(s) p(y | s) P^D(s) f^\Xi(s) dy ds \right]. \quad (59)$$

The PGFLs (59) and (46) fully characterize the similarities and differences between the scattering and clutter models, respectively.

First Summary Statistic—the Scatterer Intensity Function

The derivatives are

$$\frac{\partial^m G^{\Upsilon \Xi}}{\partial y_1 \cdots \partial y_m} [0, h] = G^{\Upsilon \Xi}[0, h] \prod_{i=1}^m \int_{S^+} h(s) p(y_i | s) P^D(s) f^\Xi(s) ds \quad (60)$$

and

$$\begin{aligned} \frac{\partial^{m+1} G^{\Upsilon \Xi}}{\partial x \partial y_1 \cdots \partial y_m} [0, h] &= G^{\Upsilon \Xi}[0, h] f^\Xi(x) (1 - P^D(x)) \\ &\times \prod_{i=1}^m \int_{S^+} h(s) p(y_i | s) P^D(s) f^\Xi(s) ds \\ &+ G^{\Upsilon \Xi}[0, h] P^D(x) f^\Xi(x) \sum_{i=1}^m p(y_i | x) \\ &\times \prod_{k=1, k \neq i}^m \int_{S^+} h(s) p(y_k | s) P^D(s) f^\Xi(s) ds. \end{aligned} \quad (61)$$

The intensity function of the iFilter is the ratio of (61) to (60) evaluated at $h(s) = 1$:

$$\begin{aligned} f_{\text{iFilter}}^{\Xi|\Upsilon}(x) &= f^{\Xi}(x) \left[1 - P^D(x) + P^D(x) \sum_{i=1}^m \frac{p(y_i | x)}{\int_{S^+} p(y_i | s) P^D(s) f^{\Xi}(s) ds} \right] \\ &= f^{\Xi}(x) \left[1 - P^D(x) + P^D(x) \sum_{i=1}^m \frac{p(y_i | x)}{\hat{\lambda}(y_i) + \int_S p(y_i | s) P^D(s) f^{\Xi}(s) ds} \right] \end{aligned} \quad (62)$$

where, for any $y \in Y$,

$$\hat{\lambda}(y) = p(y | \phi) P^D(\phi) f^{\Xi}(\phi) \quad (63)$$

is the estimated measurement intensity at $y \in Y$ due to scatterers with state ϕ . Since (62) holds for $x \in S^+$, the updated intensity $f_{\text{iFilter}}^{\Xi|\Upsilon}(\phi)$ is (62) evaluated for $x = \phi$.

The likelihood function can be chosen so that the posterior intensity $\hat{\lambda}(y)$ matches the specified exogenous clutter intensity $\lambda(y)$. In addition to the parameter choices made in (41), let the detection probability be $P^D(\phi) = 1$ and define $p(y | \phi) = \lambda(y)/\lambda$. From (41), the predicted intensity is $f^{\Xi}(\phi) = \Psi(\phi | \phi) f^{\Xi}(\phi) = \lambda$, so that $\hat{\lambda}(y) = \lambda(y)$.

Second Summary Statistic—the Scatterer Canonical Distribution

The canonical number is the number of scatterers in all of S^+ . Note that the count necessarily includes scatterers with state ϕ . The PGF of the canonical number is, using (60) and (30),

$$\begin{aligned} F_{\text{iFilter}}^{\Xi|\Upsilon}(x) &= x^m \exp \left[(x-1) \int_{S^+} (1 - P^D(s)) f^{\Xi}(s) ds \right] \\ &= x^m \exp \left[(x-1) (1 - P^D(\phi)) f^{\Xi}(\phi) \right] \\ &\quad \times \exp \left[(x-1) \int_S (1 - P^D(s)) f^{\Xi}(s) ds \right]. \end{aligned} \quad (64)$$

This PGF is the product of the PGFs of three mutually independent scattering processes. One is due to the endogenous measurement model and generates exactly m scatterers. The others are Poisson distributed and correspond to scatterers in state ϕ and to scatterers with states in S , with expected canonical numbers $(1 - P^D(\phi)) f^{\Xi}(\phi)$ and $\int_S (1 - P^D(s)) f^{\Xi}(s) ds$, respectively. The latter PGF is the first factor in (54).

The expected number of scatterers in the Bayes posterior process is

$$\hat{N}_{\text{iFilter}}(S^+) = \int_{S^+} f_{\text{iFilter}}^{\Xi|\Upsilon}(x) dx = f_{\text{iFilter}}^{\Xi|\Upsilon}(\phi) + \hat{N}_{\text{iFilter}}(S) \quad (65)$$

where

$$\hat{N}_{\text{iFilter}}(S) = \int_S f_{\text{iFilter}}^{\Xi|\Upsilon}(x) dx \quad (66)$$

is the expected number of scatterers with states in S . The iFilter estimate (66) is similar to the PHD estimate (53). The distribution of the number of scatterers in S (i.e., targets) is conditioned on the number of scatterers with state ϕ . This topic is the subject of on-going work [2] and is outside the scope of the present paper. Practical experience to date [12, 13] shows that the iFilter has excellent performance.

An alternative derivation of the iFilter can be based on the PGFLs of the detected and undetected scatterer processes. These processes are thinned versions of the parent process Ξ , where the thinning function is the probability of detection $P^D(s)$. Under the PPP assumption for Ξ , they are also mutually independent, not conditionally independent. The derivation is similar to the one just given and is omitted.

6. CONCLUDING REMARKS

Finite point process models are excellent models for sensor measurement sets in many traditional applications involving point targets whose measurements are superimposed with clutter. In contrast, they are only approximate models for multitarget state. Accepting the point process model for the multitarget state as a given, the PHD filter and iFilter are good applications of the PGFL approach.

The PGFL approach is seen to provide a common language to clarify the similarities and differences between the clutter and target models used in the PHD filter and the iFilter. Approximations and other issues seem to preclude using PGFLs directly as a basis for comparing tracking performance. In any event, such comparisons are outside the scope of the present paper.

Both the exogenous clutter model and the endogenous scattering model lead to enumerations of measurements to targets. Although the technical details differ somewhat between the PHD filter and iFilter, manipulating the required enumerations is facilitated by the PGFL approach.

By-passing explicit enumerations in the Bayes posterior process can sometimes obscure salient features of the problem, features that can make direct methods worthwhile. For example, the iFilter is derived in [15] for the scattering model by a direct Bayesian method without resorting to the PGFL. This alternative

derivation—not given here—illuminates several aspects of the PGFL method and how it works.

Enumerations of measurement to target assignments are encoded in the PGFL of the measurement-target process. An excellent example is the way in which the PGFL of the general Bayes posterior process is written as a ratio of functional derivatives of the joint PGFL.

The functional derivatives of the PGFL of the Bayes posterior process decode the probability structure. When the functional derivatives are of high computational complexity, the utility of the PGFL approach is severely limited. Examples of this limitation—not discussed in this paper—are extended targets and the target-centric multisensor PHD filter and iFilter.

Finally, the PGFL approach may suggest alternative problems of independent interest. One example is a traffic process [18] that counts sensor target detections, not the targets themselves. Its computational complexity is linear in the number of sensors.

APPENDIX. ASSIGNMENT ENUMERATION IN THE PHD FILTER

For $n = 0$, no targets are present and (44) reduces to the PGFL of clutter. When the sensor reports no measurements, $v = (0, \emptyset)$ and

$$\begin{aligned} p^{\Upsilon|\Xi}(v = (0, \emptyset) \mid \xi = (n, \{s_1, \dots, s_n\})) \\ = G^{\Upsilon|\Xi}[0 \mid \xi] = e^{-\int_Y \lambda(y) dy} \prod_{i=1}^n (1 - P^D(s_i)). \end{aligned} \quad (67)$$

This is the probability that the realizations of the clutter process and the n target processes contain no points.

Explicit expressions of $p^{\Upsilon|\Xi}(v \mid \xi)$ for events $v = (m, \{y_1, \dots, y_m\})$, $m \geq 1$, are found by functional differentiation of (44) with respect to impulses at the (distinct) points y_1, \dots, y_m . Functional differentiation masks enumerations of measurement to target. For $m = 1$,

$$\begin{aligned} p^{\Upsilon|\Xi}(v = (1, \{y_1\}) \mid \xi = (n, \{s_1, \dots, s_n\})) \\ = \frac{\partial G^{\Upsilon|\Xi}}{\partial y_1}[0 \mid \xi] \\ = e^{-\int_Y \lambda(y) dy} \lambda(y_1) \prod_{i=1}^n (1 - P^D(s_i)) \\ + e^{-\int_Y \lambda(y) dy} \sum_{i=1}^n P^D(s_i) p(y_1 \mid s_i) \prod_{k=1, k \neq i}^n (1 - P^D(s_k)). \end{aligned} \quad (68)$$

In words, (68) is the probability that y_1 is produced either by 1) the clutter process and all n target processes are undetected or by 2) exactly one of the n target processes and there is no clutter.

The functional derivatives of the PGFL for $m \geq 2$ are expressions that sum over all assignments of m measurements to n targets that are consistent with the

constraint of “at most one measurement per target per scan.” Details are omitted.


REFERENCES

- [1] K. B. Athreya and P. E. Ney *Branching Processes*. Berlin: Springer-Verlag, 1972 (Republished by Mineola, NY: Dover, 2004).
- [2] Ö. Bozdoğan, R. Streit, and M. Efe Palm intensity for track extraction. In preparation.
- [3] X. Chen, R. Tharamrasa, T. Kirubarajan, and M. Pelletier Integrated clutter estimation and target tracking using Poisson point process. SPIE vol. 7445, 2009.
- [4] D. J. Daley and D. Vere-Jones *An Introduction to the Theory of Point Processes, Vol. I: Elementary Theory and Methods*. NY: Springer, 1988 (Second Ed., 2003).
- [5] D. J. Daley and D. Vere-Jones *An Introduction to the Theory of Point Processes, Vol. II: General Theory and Structure*. NY: Springer, 1988 (Second Ed., 2008).
- [6] T. E. Harris *The Theory of Branching Processes*. Berlin: Springer-Verlag, 1963.
- [7] G. E. Kopec Formant tracking using hidden Markov models and vector quantization. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 34 (1986), 709–729.
- [8] R. P. S. Mahler A theoretical foundation for the Stein-Winter probability hypothesis density (PHD) multitarget tracking approach. In *Proceedings of the MSS National Symposium on Sensor and Data Fusion*, vol. 1, 2000, 99–117.
- [9] R. P. S. Mahler Multitarget Bayes filtering via first-order multitarget moments. *IEEE Transactions Aerospace and Electronic Systems*, 39 (2003), 1152–1178.
- [10] R. P. S. Mahler *Statistical Multisource-Multitarget Information Fusion*. Boston: Artech House, 2007.
- [11] J. E. Moyal The general theory of stochastic population processes. *Acta Mathematica*, 108 (1962), 1–31.
- [12] M. Schikora, W. Koch, R. Streit, and D. Cremers Sequential Monte Carlo method for the iFilter. In *Proceedings of the Fourteenth International Conference on Information Fusion*, June 2011.
- [13] M. Schikora, W. Koch, R. Streit, and D. Cremers A sequential Monte Carlo method for multi-target tracking with the intensity filter. Chapter 3 in: *Advances in Intelligent Signal Processing and Data Mining: Theory and Applications*, P. Georgieva, L. Mihaylova, L. Jain (Eds.): Springer, 2012.
- [14] M. C. Stein and C. L. Winter An additive theory of probabilistic evidence accrual. Los Alamos National Laboratory, Report LA-UR-93-3336, 1993.
- [15] R. L. Streit and L. D. Stone Bayes derivation of multitarget intensity filters. In *Proceedings of the Eleventh International Conference on Information Fusion*, June 2008.
- [16] R. L. Streit *Poisson Point Processes—Imaging, Tracking, and Sensing*. NY: Springer, 2010.

[17] R. L. Streit
PHD intensity filtering is one step of an EM algorithm for positron emission tomography.
In *Proceedings of the Twelfth International Conference on Information Fusion*, July 2009.

[18] R. L. Streit
Multisensor traffic mapping filters.
In *Proceedings of the Workshop on Sensor Data Fusion*, Sept. 4–6, 2012.

Roy Streit received a Ph.D. in mathematics in 1978 from the University of Rhode Island. He is currently a senior scientist at Metron in Reston, VA. His research interests include multi-target tracking, multi-sensor data fusion, distributed systems, medical imaging, and signal processing, as well as statistical methods for pharmacovigilance and business analytics. Prior to joining Metron in 2005, he was a senior scientist in the Senior Executive Service at the Naval Undersea Warfare Center in Newport, RI, working primarily on the development, evaluation and application of multi-sensor data fusion algorithms in support of submarine sonar and combat control automation. He was an exchange scientist at the Defence Science and Technology Organisation (DSTO) in Adelaide, Australia, from 1987–1989. In 1999, he received the Solberg Award from the American Society of Naval Engineers for contributions to naval engineering through personal research, and in 2001 the Department of the Navy Superior Civilian Achievement Award. From 1996 to 2005 he served on the Sonar Technology Panel (Panel 9) of The Technical Cooperation Program (TTCP), a multinational governmental organization supporting scientific information exchange between member nations. He was President of the International Society for Information Fusion (ISIF) in 2012, and he continues to serve on the ISIF Board of Directors.



Dr. Streit is a Fellow of the IEEE. He is the author of *Poisson Point Processes: Imaging, Tracking, and Sensing*, published by Springer in 2010. It was translated into Chinese and published by Science Press, Beijing, in 2013. He has published papers in over a dozen refereed technical journals, and given numerous invited and contributed papers at international conferences and workshops. He holds nine U.S. patents. He is a professor (adjunct) in the Department of Electrical and Computer Engineering at the University of Massachusetts–Dartmouth. He was a visiting scientist at Yale University from 1982–1984, and a visiting scholar at Stanford University from 1981–1982. He is currently an adjunct affiliate consultant with SUCCEED Educational Consultancy based in Rhode Island.

Establishment of Human Performance Baseline for Image Fusion Algorithms in the LWIR and SWIR Spectra

CHRISTOPHER HOWELL
STEVE MOYER

This research is complementary to research presented in “Establishment of Human Performance Baseline for Image Fusion Algorithms in the LWIR and MWIR Spectra” by Moyer and Howell in which we established a baseline performance candidate for image fusion comparison by investigating the impact of different display formats on the probability of identification, P(ID), performance of a human observer. We advance this line of research by measuring the inherent ability of the human observer to perform an identification task using source band imagery, long-wave (LW) infrared and short-wave (SW) infrared that was not algorithmically fused prior to human observation. A multi-part experiment was conducted where human observers were asked to identify displayed military targets using a standard set of tracked military vehicles. The observers performed the identification (ID) visual discrimination task using source band imagery concatenated or presented side-by-side on a single monitor, temporally interlaced source band imagery on a single monitor. Observers’ performances using source band imagery fused with the superposition fusion algorithm was also included as a reference because it is a well understood algorithm and shares an assumed similarity with the temporal interlaced display format. This research proposes that a forced choice human perception experiment is a more meaningful evaluation of an image fusion algorithm’s performance, specifically when the application of the algorithm is for dimensionality reduction to a single image designed for human observation. The results from this research identify a clear performance baseline when assessing human observer P(ID) performance employing an image fusion algorithm.

Manuscript received March 19, 2013; released for publication June 14, 2013.

Refereeing of this contribution was handled by Alexander Toet.

Authors’ address: U.S. Army REDECOM CERDEC Night Vision and Electronic Sensors Directorate, 10221 Burbeck Rd, Fort Belvoir, Va 22060-5806, USA, E-mail: (info@nvl.army.mil).

1557-6418/13/\$17.00 © 2013 JAIF

1. INTRODUCTION

The performance benchmark process for discriminating between image fusion algorithms presented by Howell [3] and again by Moyer and Howell [6] establishes a performance goal for any aggregate function merging the source band imagery under investigation. It was reported by Moyer and Howell, in “Establishment of Human Performance Baseline for Image Fusion Algorithms in the LWIR and MWIR Spectra,” that increased P(ID) performances could be realized dependent upon how the source band information was presented. In that work, human observers were asked to identify displayed targets using a standard set of tracked military vehicles. The observers performed the ID tasks using LW and MW source band imagery concatenated on a single monitor, presented side by side on a single monitor, temporally interlaced source band imagery on a single monitor, and source band imagery presented to each eye of the observer in parallel. The research presented in this paper explored the impact dissimilar source band information had on the observer’s ability to identify targets without the aid of image fusion algorithms. The spectral bands under consideration for this effort were the LW and SW spectral bands. It was hypothesized that the different display techniques using dissimilar source spectral band information would better allow the observers to choose the portions of the source band images they needed to perform the visual discrimination task of identification above that achievable using the fused superposition images. It was determined after comparing the observers P(ID) performances using the superposition fused images to the resultant P(ID)s’ using these display techniques, that the performances using the superposition fused images were well below that which was achieved by the observers benchmark source band performance. Allowing the observer to view spectral source band imagery in different display formats without the aid of image fusion algorithms, which we refer to as “self-fusion,” allows the experimenter to establish an absolute benchmark for discriminating between image fusion algorithm performances.

It was our intent to perform a mirror analysis of the research performed in Moyer and Howell [6] using LW and SW imagery; however the experiment where LW and SW source band imagery was presented to each eye of the observer in parallel could not be completed due to the effects of binocular rivalry [1] caused by the competing information presented in each eye independently. The remaining experiments performed using the LW and MW imagery were repeated in this work using LW and SW imagery.

This paper is outlined as follows: a background section describes some common approaches to image fusion along with some common image quality metrics and their shortcomings regarding predicting human task

performance; a section describing the imagery and experimental set-ups used in this study followed by a section showing the results of each experiment and a section discussing the results; followed by a summary of the conclusions.

2. BACKGROUND

Many military operations require soldiers to perform visual discrimination tasks, such as detection, recognition and identification (DRI) of targets. These tasks are conducted in a wide range of environments and in both daytime and nighttime settings. Information apparent in one spectral band might not be present in another. To this end, the military continuously seeks to improve its imaging capabilities for both day and night operations. As a result, many methods and practices have been employed to assess image fusion algorithm performances with hopes of improving DRI tasks performance [2, 5, and 9]. Combining multi-sensor data onto a single display or into a single image supports the need to provide each soldier or end user with as much relevant and high quality information in the most efficient manner possible.

In general, image fusion techniques can be categorized into three categories or levels: pixel level fusion, feature level fusion and decision level fusion [4]. Pixel level fusion requires an algorithm to first register the source band imagery before combining their information on a pixel-by-pixel basis. Feature level fusion requires fusion of features extracted from the images such as edges or textures to obtain new feature sets. Decision level fusion requires an initial judgment be made on extracted features of a target from multiple sensors and then renders a decision based on the aggregate result whether or not the correct target was identified. Several factors exist that complicate image fusion algorithm assessments, including but not limited to: a lack of proper registration between the source images; the non-linearity with which many image fusion algorithms operate on image data; the absence of an available reference image; and often large disparities exist between spectral bands being fused, types of backgrounds, the target sets and scenes used in comparison studies reported throughout the literature. Because of diverse image characteristics coupled with the lack of a standardized image database, it is a difficult task to identify common requirements and capabilities of image fusion algorithm performance. However, there are three fundamental requirements that should be achieved by any fusion algorithm: (1) the fusion algorithm should preserve as much of the salient information in the source images as possible; (2) the fusion algorithm should not generate artifacts that affect the human observer's ability to perform the task; and (3) the fusion algorithm must be tolerant of imperfections in the source imagery such as noise or improper registration.

Taking into account the fundamental requirements for all fusion algorithms, it seems reasonable then to approach the assessment of image fusion quality by: (1) evaluating the transfer of relevant information content from the spectral source band images to the fused images; (2) quantifying how much degradation or artifacts can exist before human performance is affected. Traditionally, image quality metrics are applied to fused imagery to discriminate between algorithms based on comparisons with other algorithms. The research in this paper places more emphasis on the "self-fusion" performance capability of the human. This "self-fusion" measurement provides an appropriate baseline to compare human task performance using image fusion algorithms. Understanding the human's performance capacity to exploit spectral source band images without the aid of a computer algorithm will ultimately contribute to understanding the relationship between measures of image quality and measures of task performance. Truly understanding the impact of "self-fusion" is necessary to make the best cost decisions and identify the best direction for future image fusion research.

3. DESCRIPTION OF IMAGES AND EXPERIMENTS

The target set used in this study was a standard set of tracked military vehicles referred to as the "8-target set." The "8-target set" was constructed, based on a history of research [8], so that certain vehicles were highly similar to other vehicles while at the same time subsets of vehicles still had distinct characteristics. As a review, the target set consists of the 2S3 and M109 self-propelled artillery pieces, the BMP, M113, and M2 armoured personnel carriers, and the M60, T-62 and T-72 main battle tanks. All observers were trained using the recognition of combatant vehicles (ROC-V) training package [7]. Prior to participating in the perception experiments each observer needed to obtain a 96% probability of correctly identifying the eight previously named vehicles at different aspects in both the reflective and emissive wavebands.

Figure 1 shows a sample of the imagery used in the experiments reported in this paper. The long-wave (LW) infrared and short-wave (SW) infrared spectral signatures for the targets in the scenes appear to be very different and one can immediately see significant differences exist between them. The source spectral bands shown in Figure 1 were selected specifically to test how human observers processed fused imagery from complementary sources.

Figure 2 shows a snapshot of one particular target imaged at its 3 different aspects and all the respective test ranges. The test ranges were distributed between 100 m and 2 km and all targets were imaged at each range. However, publishing specific information regarding the sensors used in this study is prohibited; as a result only relative ranges are reported throughout this



Fig. 1. Sample test imagery: Eight different targets; Two different spectral bands; One range; LW targets top row; SW targets bottom row.




Fig. 2. Single target; three different aspects; five different ranges.



Fig. 3. Example of a concatenated image used in side-by-side experiment, LW image on left and SW image on right.

paper. The interested reader should contact the corresponding author to request a copy of the experimental imagery.

3.1. Side-by-side Imagery

Images were concatenated so that each image rendered on the computer display was the height of each source band image but twice the width of each original image, as shown in Figure 3.

This experiment was designed as an 8-alternative forced choice (8-AFC) experiment to test the human observers' ability to identify targets in each source band image as a concatenated image and then a spectral combination of LW as the left image and SW as the right image. The only time a LW image appeared as the right image was when both images were LW. This was necessary to reduce the size of the experiment, and it was believed that any biases related to testing the SW imagery only on the right would have minimum impact.

The imagery was grouped by range into experimental cells. When the imagery was presented for the experiment, the experimental cells in the experiment were randomized and the images in each experimental cell were also randomized. This ensured that no observers saw the images in the same order.

3.2. Temporally Interlaced Imagery

Source band images were written as individual fields of a movie frame. The movie was played with a field rate of 30 Hz which produced a frame rate of 15 Hz. This simulated a progressive scanned display. Each movie was looped until the observer made a selection. As mentioned in the target set section these images were spatially registered to ensure that no additional blur was added when the movie was viewed. This experiment was designed as an 8-AFC experiment to test each source band image as a 2-field movie. A spectral combination of LW and SW was produced as the other

2-field movie. The movies were grouped by range into experimental cells. When the movies were presented for the experiment, the experimental cells were randomized and the movies in each experimental cell were also randomized. This ensured that no observers saw the same sequence of movies at any range or even the same sequence of ranges. While the movies were being played, there was discernible flickering between the images of the different spectral source bands.

3.3. Algorithm Fused Imagery

This experiment was also designed as an 8-AFC experiment to test each spectral source band image as an algorithmically fused image. The imagery was grouped by range into experimental cells. When the imagery was presented for the experiment, the experimental cells in the experiment were randomized and the images in each experimental cell were also randomized. This ensured that no observers saw the same sequence of images or ranges.

Each image was shown at its native format. However, each image was either a spectral source band image or an image fused using the super position algorithm with a ratio of 0.5. This fusion algorithm was chosen because it is well understood, easily implemented and shares an assumed similarity with the temporal interlaced display format. Additionally, in order to account for experimental variances regarding future image fusion studies, the superposition algorithm results presented in this work can be used as a normalization factor, thereby providing clear performance comparisons between future algorithm research and this research.

4. EXPERIMENTAL RESULTS

In each experiment, each observer's ID performance was calculated for the experiment. This performance probability was corrected for the probability of guessing

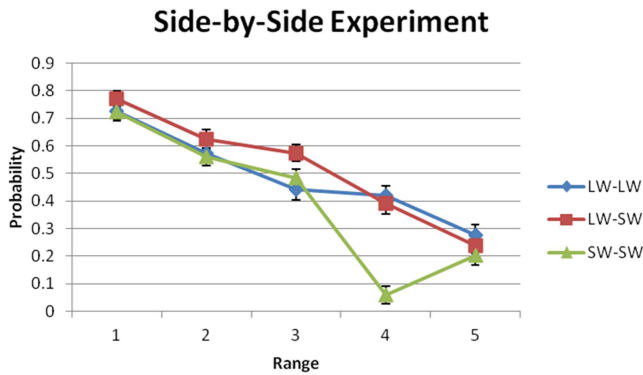


Fig. 4. Experimental results of concatenated (side-by-side) LW and SW spectral images. Error bars are one standard deviation from the mean normalized to the square root of the observer population.

the correct answer according to the following relationship

$$P_C = \frac{P_M - \left(\frac{1}{8}\right)}{\left(\frac{7}{8}\right)} \quad (1)$$

where P_M is the measured probability of the observer's response and $1/8$ is the probability of guessing the correct answer for an 8-AFC experiment. These probabilities were then ordered and two statistical tests were performed to test each observer's performance relative to the distribution of the ensemble. The first test was an inter-quartile distance test. If any observer's performance was less than three times the distance from the first quartile to the median, those observer's results were removed from the data set for the entire experiment. If any observer's performance was more than three times the distance from the median to the third quartile, those observers' results were removed from the data set for the entire experiment. The second statistical test was Chavenault's criterion. Since Chavenault's criterion assumes that measurements follow a Normal distribution, it is sensitive to the distribution of the data. An inter-quartile test based on the median value and the distance between the first and third inter-quartile to the median makes no assumptions on the distribution of the data and was found to be more useful initially to find outlying observer results. If an observer result was rejected by either test, that observer result was removed from any further analysis.

4.1. Side-by-side Experiment Results

Twenty-three observers participated in the experiment and five observers were removed because their overall results in the experiment were rejected as outliers by the previously mentioned statistical tests. The observer responses were then averaged over all images at each specific range for each waveband combination. These probabilities were then corrected for guessing using the same algorithm discussed earlier. The experimental results of the remaining eighteen observers are shown in Figure 4.

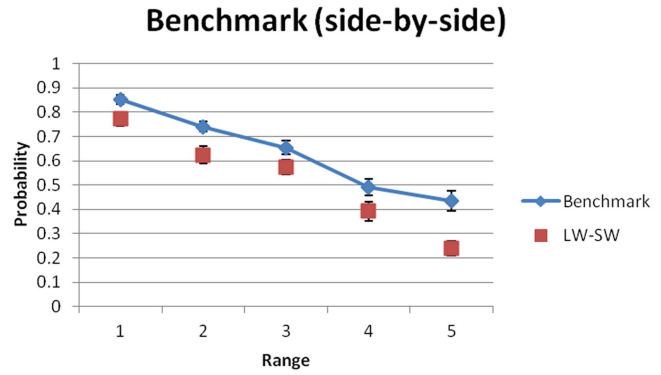



Fig. 5. Benchmark performance results of concatenated (side-by-side) LW and SW spectral images. Error bars are one standard deviation from the mean normalized to the square root of the observer population.

As may be seen in Figure 4, there are only 2 occurrences in which the average values differ by more than the standard error. LW-SW imagery outperformed both the LW-LW and SW-SW images at Range 3 and SW-SW underperformed both the LW-LW and LW-SW imagery at Range 4. The results using the LW-SW images show that the observers effectively used the available source band information as needed when making their decisions. This is evident by observing that the LW-SW curve tracks the better performing source band across range. This is an important characteristic that any image fusion algorithm should obtain if the intent is to judge how it affects task performance. This characteristic, the best combined spectral source band performance achievable, is referred to as benchmarking performance [3] and can be calculated by recording when the observer correctly identified the target using either spectral image. A performance benchmark indicates the optimal level of performance capability based on the fact that the information present in both source bands is the same across image fusion techniques; with differences being attributed to how the information is merged. The benchmark source band ID performance for the side-by-side LW and SW experiment is shown in Figure 5.

The benchmark source band performance was calculated from the LW-LW and SW-SW images. If the observer responded with the correct answer for a target and aspect in either spectral source band, then the image was graded as correct even if the observer recorded an incorrect answer in the other spectral band. A subtle but important point should be noted that range 4 produced the worst performance in the SW-SW performance with a 0.06 probability, yet the benchmark performance at range 4 is 0.06 greater than the better performing LW-LW source band. Therefore, the few images that were correctly identified in the SW-SW case were those vehicles and aspects that were incorrectly identified in the LW-LW source band. This subtlety of the data is a result that is missed if we assume that the observer is not allowed to select between different source bands

LW-LW Results Compared to LW-SW Results



SW-SW Results Compared to LW-SW Results




Fig. 6. Comparison of concatenated (side-by-side) LW-SW performance to individual source band performance.

and must choose only a source band image (LW or SW) or the fused image. A possible explanation for the degraded performance in the concatenated LW-SW imagery is that the observers may have cued exclusively onto one spectral image even though both were available. In order to test this explanation we can compare the LW-SW results to both the LW-LW and SW-SW results. These comparisons are shown as Figure 6.

As may be seen in Figure 6, LW-SW results are almost identical to the LW-LW results. In fact the correlation coefficient between these data sets is 0.9522, whereas the correlation coefficient between the LW-SW and SW-SW results is 0.8991. As in the LW-MW experiment, the LW-SW results are more like the LW-LW results than the SW-SW results, implying that the observers made more targeting decisions using the LWIR spectral image when both spectral images were present. This effect might be addressed in better training for this display format.

4.2. Temporally Interlaced Imagery Experiment Results

Twenty-three observers participated in the experiment and five observers were removed because their overall results in the experiment were rejected as outliers by either an inter-quartile distance test or Chave-nault's criterion. The experimental results of the remaining eighteen observers are shown in Figure 7.

Figure 7 shows there is one occurrence in which the averages were separated by more than the standard

Interlace Experiment




Fig. 7. Experimental results of LW and SW spectral 2-field movies. Error bars are one standard deviation from the mean normalized to the square root of the observer population.

Benchmark (Interlace)




Fig. 8. Benchmark performance results of LW and SW spectral 2-field movies. Error bars are one standard deviation from the mean normalized to the square root of the observer population.

error of the data. Range 4 LW-SW outperformed SW-SW and the LW-LW data outperformed the other two. For all other ranges, observers performed equally well on all movies regardless of spectral content. The results using the LW-SW images show that the observers effectively used the available source band information as needed when making their decisions. This is evident by observing that the LW-SW curve tracks the better performing source band across range. This is an important characteristic that any image fusion algorithm should obtain if the intent is to judge how it affects task performance. However, if either source band spectral image is available to the observer and viewed at will by the observer, the benchmark performance achievable for this experiment is shown in Figure 8.

As in the previous experiment, the calculated benchmark source band performance was calculated from the LW-LW and SW-SW movies. If the observer responded with the correct answer for a target and aspect in either spectral source band, then the movie was graded as correct even if the observer recorded an incorrect answer in the other spectral band. A possible explanation for the degraded performance in the LW-SW movie is that the spectral image least useful to the observer may have masked or diminished the spectral information in the

TABLE 1

Comparison of observer performance between the different display formats while viewing only LW source band imagery. Avg. is the average probability of identification P(ID) of all observers and \pm Error is the standard error for each P(ID).

LW	Range									
	1		2		3		4		5	
Display Type	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error
Side-by-Side	0.728	0.030	0.575	0.027	0.442	0.037	0.420	0.036	0.278	0.037
Temporal Interlace	0.778	0.026	0.644	0.030	0.553	0.033	0.466	0.035	0.257	0.024
Algorithm Fused	0.730	0.054	0.448	0.046	0.405	0.046	0.325	0.046		

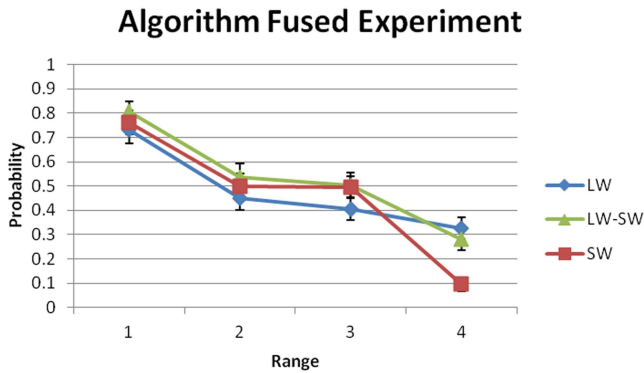


Fig. 9. Experimental results for the LW-SW superposition algorithm fused imagery. The error bars are one standard deviation from the mean normalized to the square root of the observer population.

resulting movie. This effect cannot be addressed with training as it is an artifact of the display format.

4.3. Algorithm Fused Image Experimental Results

Fifteen observers participated in this experiment and two of them were removed as outliers. The remaining thirteen observer responses were then averaged over all images at each specific range for each waveband combination. Range 5 was omitted from this experiment. The probabilities were corrected for guessing as described previously, Equation 1, and the results are shown in Figure 9.

Figure 9 shows the LW and SW images are never separated further than the standard error at ranges 1 and 2. However, at range 3 SW outperforms LW and at range 4 LW outperforms SW. The LW-SW superposition fused imagery performed as well as the best performing source band spectral images within the standard error of the data sets. This was not the case in the LW-MW superposition experiment, where the LW-MW superposition fused imagery underperformed both spectral source bands greater than can be accounted for by the standard error in the data sets. When the source band information was similar (LW-MW) the superposition algorithm degraded the source band information enough that task performance suffered; in comparison, when the source bands contain different information (LW-SW)




Fig. 10. Benchmark performance results for the LW-SW superposition algorithm fused imagery. The error bars are one standard deviation from the mean normalized to the square root of the observer population.

about the same scene the superposition algorithm retained enough source band information such that task performance did not suffer.

Figure 10 shows the calculated benchmark source band performance obtained using the LW and SW images. If the observer responded with the correct answer for a target and aspect in either spectral source band, then the image was graded as correct even if the observer recorded an incorrect answer in the other spectral band. A possible explanation for the degraded performance in the LW-SW algorithm case is that the superposition algorithm may have masked or diminished the spectral information provided by the spectral image most useful to the observer. This effect cannot be addressed with training as it is an artifact of the algorithm.

5. COMPARISON OF EXPERIMENTS

In order to draw wider ranging conclusions, it was important to assess if the display format contributed to differences in performance. As a first order analysis, a direct comparison of observer responses for each display format when viewing single spectral source band images will be done followed by a comparison of performance between the display formats with mixed spectral image content.

Table 1 shows a comparison between the average observer performances using only LW images at each of the five ranges for each display format.

TABLE 2

Comparison of observers' LW P(ID) performance between the various display formats for each range which exceeded the standard error for the data set. (+) column exceeds row, (–) row exceeds column, (X) no comparison.

LW	Side-by-Side	Temporal Interlace
Temporal Interlace	–0.012 R2 –0.041 R3	X
Algorithm	+0.054 R2 +0.013 R4	+0.120 R2 +0.069 R3 +0.060 R4

not accounted for by the error associated with calculating the mean value and the largest value was 0.120 greater than the error associated with calculating the average value. However, 5 of the 7 comparisons involved the temporal interlace format outperforming the other two formats.

Table 3 shows a comparison between the average observer performances using only SW images at each of the five ranges for each display format.

Table 4 compares the observers' SW P(ID) performances between the various display formats for each

TABLE 3

Comparison of observer performance between the different display formats while viewing only SW source band imagery. Avg. is the average probability of identification P(ID) of all observers and \pm Error is the standard error for each P(ID).

SW	Range									
	1		2		3		4		5	
Display Type	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error
Side-by-Side	0.725	0.032	0.561	0.031	0.484	0.033	0.061	0.032	0.204	0.035
Temporal Interlace	0.783	0.032	0.616	0.045	0.585	0.025	0.132	0.029	0.220	0.022
Algorithm Fused	0.762	0.048	0.500	0.053	0.496	0.044	0.095	0.027		

TABLE 4

Comparison of observers' SW P(ID) performance between the various display formats for each range which exceeded the standard error for the data set. (+) column exceeds row, (–) row exceeds column, (X) no comparison, (ND) no statistically significant difference for all ranges.

SW	Side-by-Side	Temporal Interlace
Temporal Interlace	–0.043 R3 –0.010 R4	X
Algorithm	ND	+0.018 R2 +0.020 R3

range. The format in the column is being compared to the format or algorithm listed in the row. Hence reading the first column first row shows that SW Side-by-Side underperformed the temporal interlace format at ranges 3 and 4 greater than what is accounted for in the data set error.

Of the 13 paired comparisons made between the different experiments, only 4 comparisons produced differences not accounted for by the error associated with calculating the mean value and the largest value was 0.043 greater than the error associated with calculating the average value. However, all comparisons involved

TABLE 5

Comparison of observer's benchmark source band performance between the different display formats.

Benchmark	Range									
	1		2		3		4		5	
Display Type	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error
Side-by-Side	0.852	0.020	0.741	0.019	0.654	0.029	0.492	0.034	0.434	0.042
Temporal Interlace	0.892	0.021	0.793	0.027	0.725	0.027	0.561	0.035	0.447	0.024
Algorithm Fused	0.885	0.025	0.715	0.042	0.694	0.032	0.531	0.047		

Table 2 compares the observers' LW P(ID) performance between the various display formats for each range. The format in the column is being compared to the format or algorithm listed in the row. Hence reading the first column first row shows that LW Side-by-Side underperformed the temporal interlace format at ranges 2 and 3 (labeled as R2 and R3 respectively) greater than what is accounted for in the data set error.

Of the 13 paired comparisons made between the different experiments, 7 comparisons produced differences

the temporal interlace format outperforming the other two formats.

An analysis was conducted to further investigate any effect display format may have had on the observer's performance for their benchmark performance given either source spectral band imagery; Table 5 shows these comparisons.

Table 6 compares the observers' benchmark P(ID) performance between the various display formats for each range. The format in the column is being com-

TABLE 6

Comparison of observers' benchmark source band P(ID) performance between the various display formats for each range which exceeded the standard error for the data set. (+) column exceeds row, (-) row exceeds column, (X) no comparison, (ND) no statistically significant difference for all ranges.

Benchmark	Side-by-Side	Temporal Interlace
Temporal Interlace	-0.006 R2 -0.015 R3	X
Algorithm	ND	+0.009 R2

TABLE 8

Comparison of observer's performance between the different display formats using the LW-SW spectral source bands together for each range which exceeded the standard error of the data set. (+) column exceeds row, (-) row exceeds column, (X) no comparison, (ND) no statistically significant difference for all ranges.

Fused	Side-by-Side	Temporal Interlace
Temporal Interlace	ND	X
Algorithm	+0.025 R4	+0.052 R2 +0.004 R4

TABLE 7

Comparison of observer's performance between the different display formats using the two spectral source bands together. Avg. is the average probability of identification P(ID) of all observers and \pm Error is the standard error for each P(ID).

Display Type	Range									
	1		2		3		4		5	
	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error	Avg.	\pm Error
Side-by-Side	0.772	0.030	0.625	0.036	0.574	0.030	0.392	0.039	0.239	0.030
Temporal Interlace	0.823	0.028	0.680	0.036	0.563	0.044	0.370	0.038	0.204	0.031
Algorithm Fused	0.806	0.044	0.536	0.056	0.504	0.051	0.282	0.046		

pared to the format or algorithm listed in the row. Hence reading the first column first row shows that Benchmark Side-by-Side underperformed the temporal interlace format at ranges 2 and 3 greater than what is accounted for in the data set error.

Of the 13 paired comparisons made between the different experiments, only 3 comparisons produced differences not accounted for by the error associated with calculating the mean value and the largest value was 0.015 greater than the error associated with calculating the average value. However, all comparisons involved the temporal interlace format outperforming the other two formats. Although not nearly as significantly as in either source band comparison. This result provides us confidence that the benchmark concept is valid and should be reproducible regardless of display format.

Table 7 shows a comparison of observer performance between the different display formats using only LW-SW images at each of the five ranges.

Table 8 compares the observers' "self-fused" P(ID) performance between the various display formats and superposition fused P(ID) performance for each range. The format in the column is being compared to the format or algorithm listed in the row. Hence reading the first column second row shows that fused Side-by-Side outperformed the fusion algorithm format at range 4 greater than what is accounted for in the data set error.

Of the 13 paired comparisons made between the different experiments, only 3 comparisons produced differences not accounted for by the error associated with calculating the mean value and the largest value was 0.052 greater than the error associated with calculating the average value. However, all comparisons involved

the fusion algorithm underperforming the other two formats.

6. DISCUSSION AND RECOMMENDATIONS

The research presented in this paper utilized a standard set of multi-spectral images suitable for image fusion and image fusion algorithm performance evaluations. The results obtained from using both the display formats and the superposition algorithm correlated with the best performing individual source band. However, test results show the superposition fusion algorithm underperformed the temporal interlace format. In fact, temporal interlacing the imagery allowed for task performance closest to the observers' benchmark. Overall experimental results show that observers P(ID) performances using the superposition fused images are well below that which was achieved by the observers' benchmark source band performances. It is therefore recommended that superposition fusion not be used as a baseline when assessing image fusion P(ID) performance.

Benchmark performances were measured for a variety of "self-fusion" display techniques. Considering the diverse observer pool and small errors associated with the resultant data, the comparisons reported in this paper show that benchmark performances were relatively unaffected due to these changes in the display format. As a result, we are further recommending that human observer performance using fusion algorithms to fuse together the LW and SW spectral band imagery needs to achieve the optimal values shown in Table 9.

These recommendations are based on the measurements made in the experiments in which the human observer was viewing only single source band imagery.

TABLE 9

Recommended optimal human performance for fusion algorithms when fusing this LW and SW imagery.

	Range 1	Range 2	Range 3	Range 4	Range 5
Average	0.88	0.75	0.69	0.53	0.44
\pm Error	0.02	0.03	0.03	0.04	0.03

TABLE 10

Recommended optimal human performance for fusion algorithms when fusing this LW and MW imagery.

	Range 1	Range 2	Range 3	Range 4	Range 5
Average	0.88	0.66	0.50	0.33	0.38
\pm Error	0.03	0.03	0.03	0.03	0.03

That is, the image was LW or SW but not both simultaneously. For comparison, the recommended human performance given in Moyer and Howell using fusion algorithms to fuse together the LW and MW spectral band imagery needs to achieve the optimal values shown in Table 10.


It should be noted that the LW imagery used within the LW-MW tests had an average contrast of 0.32 while the LW imagery used within the LW-SW tests had an average contrast of 0.21. This is important if comparisons between the LW imagery used in the LW-MW and LW-SW experiments are conducted. The values in Table 9 and Table 10 are greater than any results achieved using the original bands exclusively. It is therefore reasonable with regards to benchmarking performance to expect this type of performance if the human observer is provided both source bands when making the targeting decision.

These benchmark performances exceeded the superposition algorithm performance for both the LW-MW and LW-SW experiments. By extension, if a fusion algorithm achieves an image quality metric similar to the superposition algorithm, then we can expect human performance using that image fusion algorithm to be

less than the human performance achievable if the observers had the ability to select between the original source band images. This research effectively bridges the knowledge gap between “self-fusion” and algorithmically fused performance assessments and has also identified a standard set of source band imagery suitable for image fusion and image fusion algorithm performance assessment.

REFERENCES

- [1] R. Blake
A Primer on Binocular Rivalry, Including Current Controversies.
Brain and Mind, 2, (2001), 5–38.
- [2] K. Byrd, H. Szu, and M. Chouika
A subspace learning approach to evaluating the performance of image fusion algorithms.
Proceedings of SPIE, 7703, (2007).
- [3] C. Howell
Benchmarking image fusion algorithm performance.
Proceedings of SPIE, 8355, (2012).
- [4] T. Hui and W. Binbin
Discussion and Analyze on Image Fusion Technology
International Conference on Machine Vision, (2009).
- [5] D. Kim, S. Choi, and K. Sohn
Depth adjustment for stereoscopic images and subjective preference evaluation
Journal of Electronic Imaging, 20(3), (2011).
- [6] S. Moyer and C. Howell
Establishment of Human Performance Baseline for Image Fusion Algorithms in the LWIR and MWIR Spectra.
Journal of Advances in Information Fusion, (2013), *submitted for publication*.
- [7] J. O’Connor
Infrared imagery acquisition process supporting simulation and real image training
Proceedings of SPIE, Modeling and Simulation for Defense Systems and Applications VII, 8403, 3 (May 2012).
- [8] J. O’Connor, P. O’Shea, J. Palmer, and D. Deaver
Standard target sets for field sensor performance measurements
Proceedings of SPIE, 6207, (2006).
- [9] A. Toet and M. Hogervorst
Progress in color night vision
Optical Engineering, 51(1), (2012).



Christopher Howell received his M.S. and Ph.D. degrees in electrical engineering from the University of Memphis, Memphis, TN in 2007 and 2010 respectively. He is currently working as an electronics engineer for the U.S. Army's Night Vision and Electronics Sensors Directorate. His current research interest involves investigating the effects of image fusion on human visual task performance and image quality assessment. Other research interests include model development for image fusion systems and target tracking.



Dr. Steve Moyer is currently working at the U.S. Army Night Vision Electronic Sensors Directorate (NVESD) as the Lead Researcher modeling EO/IR sensor performance in detecting hastily emplaced explosives from a moving ground vehicle. He has a B.S.E.E. from the Pennsylvania State University, a Master of Science from Georgia Institute of Technology, and completed his Ph.D. in August 2006 also from the Georgia Institute of Technology. Previously, Dr. Moyer was the Lead Researcher conducting experiments to further the application of sensor models, such as NVThermIP, to tasks associated with the urban battlespace. Most recently Dr. Moyer has been working on characterizing human performance in searching for and detecting small easily concealed explosives from moving vehicles using infrared sensors.

Testing Under Communication Constraints

SORA CHOI
BALAKUMAR BALASINGAM
PETER WILLETT

The problem of fault diagnosis with communication constraints is considered. Most approaches to fault diagnosis have been focused on detecting and isolating a fault under cost constraints such as economic factors and computational time. But in some systems, such as remote monitoring (e.g., satellite, sensor field) systems, there is a communication constraint between the system being monitored and the monitoring facility. In such circumstances it is desirable to isolate the faulty component with as few interactions as possible. The key consideration is that multiple tests are chosen at each stage in such a way that the tests within the chosen group complement each other. To this end we propose two algorithms for fault diagnosis under communications constraints. Their performances are analyzed in terms of the average number of testing stages as well as in terms of the required computational complexity.

Manuscript received August 2, 2012; revised June 8, 2013; released for publication October 13, 2013.

Refereeing of this contribution was handled by Mujdat Cetin.

This work acknowledges support from the U.S. Office of Naval Research under Grant N00014-10-10412, and from Qualtech Systems under a NASA STTR.

Authors' address: Department of Electrical and Computer Engineering, University of Connecticut, 371 Fairfield Way, U-2157, Storrs, Connecticut 06269 USA, E-mail: ({soc07002, willett, bala}@enr.uconn.edu).

1557-6418/13/\$17.00 © 2013 JAIF

1. INTRODUCTION

Fault diagnosis is the process of detecting and isolating component failure in a system via reports from a suite of sensors, each of which monitors a subset of the components. Since the life-cycle maintenance cost of large integrated systems such as an aircraft or the space shuttle can, due to the large number of failure states and the need to rectify these failures quickly [4], [12], exceed the purchase cost, it has been recognized that testability must be built into the manufacturing process. Owing to the advent of intelligent sensors, onboard tests are available to the diagnostic/fusion center or operators, but the computational burden of processing test results in large-scale systems is still a factor.

Various techniques for computationally efficient test sequencing to identify component failure have been developed: test sequencing for single fault diagnosis [14]–[16], [21], dynamic single fault diagnosis [5], [24], multiple fault diagnosis [10], [18], [19], dynamic multiple fault diagnosis [17], [20] and test sequencing for complex systems [2], [3], to name a few. As another approach, some fault detection schemes for networked control systems use residual generation and evaluation without utilizing built-in smart sensors to detect component failures [11], [22], [23], [25], [26].

To date, the purpose of test sequencing (see [2], [3], [14]–[16], [19], [21]) has usually been to find an optimal or suboptimal solution minimizing a “testing cost” that include economic factors, testing time, etc. In general, a sequential testing algorithm repeats the procedure (called a stage) consisting of deciding to test, ordering to perform tests, and updating the state of the system using the received test results, until the failed component is identified. Moreover, at each state, communication between the system and the diagnostic (fusion) center is required for the sensors to transmit the test results and for the center to request the performance of tests.

There are, sometimes, systems placed in distributed or remote configurations, causing unusually long delays or restriction in communication of instructions from or results to a monitoring facility. One application might be remotely to determine and diagnose the health of a hard-to-reach system, for example a space vehicle or a craft in a deep sea (please see an illustration in Fig. 1). In this situation the monitoring facility should be able to diagnose the remote system using limited communication with the remote or distributed system. Under these circumstances the *number of instances of communication* becomes a primary constraint and the testing cost, while still an issue, becomes secondary. Our goal in this paper is hence to minimize the number of communication stages between the monitoring facility




Fig. 1. Examples of a systems with limited communication between the system to be monitored and the monitoring facility. (a) A system with significant remote-ness, in which latency and communication costs are a concern. (b) A distributed surveillance system in which covertness is key.

and the remote system while keeping the testing cost within a specified level.

The optimal solution of fault diagnosis can be posed via generating a binary decision tree with DP (dynamic programming) and AND/OR graphs. Finding the optimal solution, however, is known to be an NP-complete problem for single fault diagnosis [7]–[9], [13] even when using just a single sensor at each stage, and (worse) an NP-hard problem for multiple fault diagnosis [1], meaning that it cannot be solved by an algorithm of polynomial complexity. Thus, a heuristic approach for test sequencing using multiple sensors at each stage is suggested. This involves:

- 1) selecting the sensors that will maximize the information gain (IG);
- 2) performing the test using those selected sensors;
- 3) updating conditional likelihood probability of all components' failure states, depending on the test outcomes;
- 4) pruning from consideration components that are unlikely to be faulty; and
- 5) repeating this procedure until the fault is isolated up to a certain specified probability of assurance.

We will discuss the IG heuristic later. Note that there are test inaccuracies due to unreliable sensors, electromagnetic interference, environmental conditions, and so on. Imperfect tests, for our purposes here, introduce two

uncertainties to the diagnosis process: missed detections and false alarms. Under the former, of course, a component may fail yet the test that covers it can show “pass”; and vice versa for a false alarm. Consequently, even after collection of an arbitrarily amount of test signature evidence, one is never *certain*, just sure enough up to a given probability level.¹

An important issue is that as we seek to reduce the number of iterations by selecting multiple sensors at a time, the computation—selecting the set of sensors that will maximize the information gain—increases rapidly both with number of tests and number of faults, and in many practical systems both are very large. We propose two algorithms for selecting multiple sensors at a time that maximize the information gain at an affordable computational complexity. The first algorithm introduces several thresholds in order to eliminate sensors that are less informative, so that fewer sensors form the candidate set for the maximization of information gain. The second approach populates the candidate set one-by-one, based on the correlation between the information state and the elements of the reachability matrix in addition to information gain. For simplicity, we assume that there is a single component failure during fault isolation, although the methods presented in this paper can be extended to multi-fault case by applying one of the multi-fault diagnosis techniques [10], [17]–[20] in order to mitigate (but not really avoid) the significant increase in resulting computational complexity.

The paper is organized as follows. In Section 2 we formulate the problem. In Section 3 two heuristic test sequencing algorithms are introduced. Section 4 presents the simulation results of the proposed algorithms and Section 5 concludes the paper.

2. PROBLEM FORMULATION

We consider the problem of single fault diagnosis: there are a set of components that might fail and a set of sensors each monitoring a subset of those components. The system is described in detail as given below.

- 1) A (finite) set of m possibly-faulty components $F = \{f_0, f_1, \dots, f_m\}$ (loosely: “faults”) is given, where f_0 denotes the no-fault condition and f_i denotes the i th faulty component. The state of faulty component f_i is expressed by x_i , where $x_i = 1$ if fault f_i occurs, otherwise $x_i = 0$.
- 2) A (finite) set of n binary sensors $S = \{s_1, s_2, \dots, s_n\}$ is given, where sensor s_j monitors a known subset of faulty components and costs an amount $c_j (> 0)$ to apply.

¹And indeed, depending on the test coverage (the R matrix to be defined shortly), it might *never* be possible even with perfect tests to isolate a fault more tightly than to a “ambiguity set.”

3) The reachability matrix² $R = [r_{ij}]$ represents the relationship between the faulty components and the sensors: $r_{ij} = 1$ if sensor s_j monitors faulty component i , and, otherwise, $r_{ij} = 0$. In addition, $r_{0j} = 0$ for all j .

4) Test (sensor) s_j has associated with it a probability of false alarm

$$P_{f_j} = \text{Prob}(s_j = 1 \mid \text{no component monitored by } s_j \text{ has failed}), \quad (1)$$

and probability of detection

$$P_{d_j} = \text{Prob}(s_j = 1 \mid \text{at least one component monitored by } s_j \text{ has failed}). \quad (2)$$

Since only one fault is being considered, the probabilities of detection and false alarm can be combined with the reachability matrix R to as the likelihood matrix $D = [d_{ij}]$, in which

$$d_{i,j} = \text{Prob}(s_j = 1 \mid x_i = 1) = r_{i,j}P_{d_j} + (1 - r_{i,j})P_{f_j}. \quad (3)$$

The element $d_{i,j}$ is the likelihood that the sensor j registers a “fail” despite the relation with a failure of component i , considering two cases: The first is that component i is failed and sensor s_j detects that there is a failure, when s_j is monitoring component i . The second is that sensor s_j is not monitoring component i , but s_j shows the result that there is a component failure as false alarm.

Let us address the single fault diagnosis problem under the following assumptions:

- The false-alarm and missed-detection probabilities of sensors are known and do not change with repeated testing.
- There is at most one component failure, which does not change over the course of (repeated) testing.
- It is possible that the system be fault-free.
- Each sensor’s missed-detection/false-alarm process is independent of those of other sensors.
- Outcomes of sensors are binary, meaning that there are two outcomes: pass (0) and fail (1).

Let us denote the complete set of (possibly multiple) tests applied at the k th stage as $S_c(k) = \{s_{j_q} \in S \mid j_q \in J(k)\}$ where $J(k) = \{j_q\}$ is the set of indices of the applied sensors. The outcome of the tests at the k th stage is denoted as $O(k) = \{o_q(k)\}$ where $o_q(k)$ is the result of sensor s_{j_q} . Thus, the past information available for sensor selection at the $(k + 1)$ th stage is

$$I_k = \{(S_c(l), O(l))\}_{l=1}^k. \quad (4)$$

With the past information I_k , the conditional failure probabilities $\pi_i(k + 1) = p(x_i = 1 \mid I_k)$ also known as the *information state* is updated from its previous state $\pi_i(k)$

²This is sometimes call the **D**-matrix, invoking variously test **depend**ency or **diagnosis**. Here we use “D” for the test reliabilities.

based on Bayes’ rule as

$$\begin{aligned} \pi_i(k + 1) &= p(x_i = 1 \mid I_k) \\ &= \frac{p((S_c(k), O(k)) \mid x_i = 1, I_{k-1})p(x_i = 1 \mid I_{k-1})}{p((S_c(k), O(k)) \mid I_{k-1})} \\ &= \frac{1}{c} \prod_{j_q \in \mathbf{J}(k)} p((s_{j_q}(k), o_q(k)) \mid x_i = 1)\pi_i(k) \\ &= \frac{1}{c} \prod_{j_q \in \mathbf{J}(k)} [o_q(k)d_{i,j_q} + (1 - o_q(k))(1 - d_{i,j_q})]\pi_i(k) \end{aligned} \quad (5)$$

where the normalization factor is

$$c = \sum_{i=0}^m \prod_{j_q \in \mathbf{J}(k)} [o_q(k)d_{i,j_q} + (1 - o_q(k))(1 - d_{i,j_q})]\pi_i(k).$$

In addition, the prior failure probability $\pi_i(1)$ of component i is assumed to be known, and the probability of a healthy system $\pi_0(1)$ satisfies the following:

$$\pi_0(1) = \prod_{i=1}^m \text{Prob}(x_i = 0) = \prod_{i=1}^m (1 - \pi_i(1)). \quad (6)$$

The test sequencing algorithm with imperfect tests can never, except in trivial cases, identify the faulty component deterministically, but is assumed content with a (pre-specified) level of certainty α . We have:

Stopping rule: The algorithm stops when any information state reaches a level of certainty α , i.e.,

$$\pi_i(k) > \alpha. \quad (7)$$

Pruning criterion: If $\pi_i(k)$ satisfies

$$\pi_i(k) \leq \beta\pi_i(0), \quad (8)$$

where threshold β is given, then it will be decided that the component i is not a faulty one and set $\pi_i(k) = 0$.

The *stopping rule* and *pruning criterion* in the algorithms will be described in later sections.

3. TEST SEQUENCING USING INFORMATION GAIN

Given the current information state $\{\pi_i(k)\}$ at the k th stage, the information gain achieved by testing with a set of sensors $S_c(k)$, i.e. the mutual information between the sensors and the information state, is written as

$$\text{IG}(\{\pi_i(k)\}, S_c(k)) = H(\{\pi_i(k)\}) - H(\{\pi_i(k)\} \mid S_c(k)) \quad (9)$$

where

$$H(\{\pi_i(k)\}) = - \sum_{i=0}^m \pi_i(k) \log \pi_i(k)$$

is the uncertainty (entropy) associated with the information state $\{\pi_i(k)\}_{i=1}^m$.

After performing some algebraic manipulations the information gain in (9) is written as follows (please see

Appendix for details):

$$\begin{aligned}
\text{IG}(\{\pi_i(k)\}, S_c(k)) = & \\
& - \sum_i \pi_i(k) \log \pi_i(k) \\
& + \sum_{\tilde{\mathbf{o}}_t \in \tilde{\mathcal{O}}} \sum_{i=0}^m \pi_i(k) \prod_{j \in J} A(t, i, j) \cdot \log \left\{ \pi_i(k) \prod_{q \in J} A(t, i, q) \right\} \\
& - \sum_{\tilde{\mathbf{o}}_t \in \tilde{\mathcal{O}}} \log \left\{ \sum_{p=0}^m \pi_p(k) \prod_{q \in J} A(t, p, q) \right\} \cdot \sum_{i=0}^m \pi_i(k) \prod_{j \in J} A(t, i, j)
\end{aligned} \quad (10)$$

where

$$A(t, i, j) := (\tilde{\mathbf{o}}_t(j) d_{i,j} + (1 - \tilde{\mathbf{o}}_t(j))(1 - d_{i,j})). \quad (11)$$

Here, $\tilde{\mathbf{o}}_t$ denotes a vector whose element is a possible outcome (0 or 1) of a sensor in set $S_c(k)$, the set $\tilde{\mathcal{O}} = \{\tilde{\mathbf{o}}_t\}$ consists of all possible vectors that can be generated by the sensors in set $S_c(k)$, and J is the set of indices of sensors in the set $S_c(k)$.

If the objective is only to minimize testing cost by using information heuristics, the sensor can be selected simply by maximizing the information gain per unit cost. That is, sensors $S_c(k)$, at the k th stage, can be selected based on the *Selection rule minimizing testing cost*:

$$S_c(k) = \arg \max_{\tilde{S}_c(k) \subset S} \frac{\text{IG}(\{\pi_i(k)\}, \tilde{S}_c(k))}{\sum_{j \in J(k)} c_j} \quad (12)$$

Our goal in this paper is to minimize the number of stages required to locate a fault, while limiting the cost spent at each stage to C_T . Hence, instead of performing one test at each stage, we propose to perform (possibly) several tests, where one should select the set of sensors $S_c(k)$ having the most information about the faulty component at each stage considering all possible combinations of sensors. This can be achieved via the following *Selection rule minimizing the number of stages*:

$$S_c(k) = \arg \max_{\tilde{S}_c(k) \subset S} \text{IG}(\{\pi_i(k)\}, \tilde{S}_c(k)) \quad (13)$$

$$\text{IG}(\{\pi_i(k)\}, \tilde{S}_c(k))$$

$$= -n\pi(k) \log \pi(k) - \pi_0(k) \log \pi_0(k)$$

$$+ \pi(k) \prod_{j \in J} A(t, i, j) \sum_{\tilde{\mathbf{o}}_t \in \tilde{\mathcal{O}}} \sum_{i=1}^m \left(\log \left\{ \prod_{q \in J} A(t, i, q) \right\} \log \left\{ \sum_{p=1}^m \prod_{q \in J} A(t, p, q) \right\} - \log \left\{ \pi_0(k) \prod_{q \in J} A(t, 0, q) \right\} \right) \quad (15)$$

and (13) reduces to

$$S_c(k) = \arg \max_{\tilde{S}_c(k)} \{\text{IG}(\{\pi_i(k)\}, \tilde{S}_c(k))\}$$

$$= \arg \max_{\tilde{S}_c(k)} \left\{ \prod_{j \in J} A(t, i, j) \sum_{\tilde{\mathbf{o}}_t \in \tilde{\mathcal{O}}} \sum_{i=1}^m \left(\log \left\{ \prod_{q \in J} A(t, i, q) \right\} - \log \left\{ \sum_{p=1}^m \prod_{q \in J} A(t, p, q) \right\} - \log \left\{ \pi_0(k) \prod_{q \in J} A(t, 0, q) \right\} \right) \right\}. \quad (16)$$

subject to

$$\sum_{j \in J(k)} c_j \leq C_T$$

where C_T is the cost threshold per stage.

If we countenance the use of an exhaustive search to select a set $S_c(k)$ according to the selection rule minimizing the number of stages, a test sequencing heuristic algorithm using the information gain will follow Algorithm 1.

ALGORITHM 1 *Exhaustive Search—ExS(N)*

- 1) After obtaining all possible combinations of sensors in the set S satisfying the cost constraint C_T , select $S_c(k)$ based on (13).
- 2) Obtain test outcomes of $S_c(k)$ and update information states using (5).
- 3) Apply the pruning criterion in (8): after pruning, normalize information states.
- 4) Repeat steps 1)–3) until the stopping rule in (7) is satisfied.

Algorithm 1 is computationally exhaustive due to its first step. For testing at most n_c sensors at a time, the number of all possible combinations of tests is

$$\sum_{n=1}^{n_c} \binom{n}{m}. \quad (14)$$

To give some perspective, if there are $m = 100$ sensors, considering pairs of tests jointly results in 4950 combinations and testing 3 tests at a time results in 161700 combinations. Thus, in what follows, we focus on finding suboptimal ways to find $S_c(k)$ that have good heuristics.

Before discussing suboptimal algorithms, we point out a useful property of the information gain in (10). It is observed that when the information states $\{\pi(k)\}$ are uniform, except for $\pi_0(k)$ —because $\pi_0(k)$ corresponds to the fault-free state—the information gain in (10) reduces to

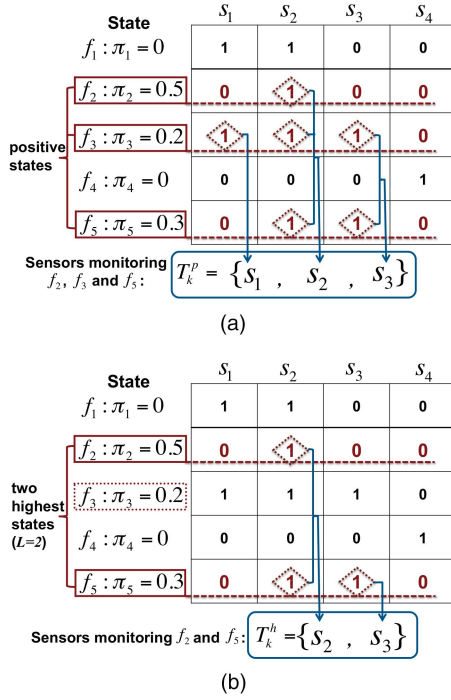


Fig. 2. The relationship between information state and set T_0 or T_1 . (a) Set T_0 . (b) Set T_1 : $L = 2$.

If the information state is uniform except for π_0 , one can select the set of sensors maximizing information gain simply by appeal to the limit of cost—even before beginning to monitor the system according to (16)—which means that the set $S_c(k)$ can be chosen off-line. For example, in the case that the prior probability of fault elements be uniform, one can start to monitor the system with $S_c(1)$ selected by off-line.

Now, we propose two suboptimal algorithms to choose sensors. The first, described in Section 3.1, proposes to decrease the size of set from which the sensors are chosen based on the current information state. The second, described in Section 3.2, proposes to select sensors sequentially based on information gain and current information states, after the first sensor is chosen via correlation between information states and sensors.

3.1. Suboptimal algorithm based on Exhaustive Search (ES)

It should be noted that the application of the pruning rule in (8) at each stage reduces the number of possible faulty components up to an uncertainty, and similarly the size of set $T_0(k)$ where

$$T_0(k) = \{s_{i_q} \in S \mid R(j, i_q) = 1, \text{ for any } j \text{ s.t. } \pi_j(k) > 0\}. \quad (17)$$

As a result there is no need to test those sensors that are monitoring *only* those excluded components, as a result of pruning. Thus, at every stage k , as pruning is performed to eliminate some of components having little possibility of a fault the new set of sensors $T_0(k)$ satisfies $T_0(k) \subset T_0(k-1) \subset S$.

If the size of $T_0(k)$ is too large—possible in the early stages of testing—one can consider a new set $T_1(k) \subset T_0(k)$ whose elements are those sensors monitoring components having high probabilities of failure, as follows. Given a threshold L

$$T_1(k) = \{s_{i_q} \mid R(j_p, i_q) = 1, \text{ for any } p = 1, 2, \dots, L\} \quad (18)$$

where $\pi_{j_1}(k) \geq \pi_{j_2}(k) \geq \dots \geq \pi_{j_L}(k) > 0$ and $L \geq \tilde{l}$. If \tilde{l} is less than L , $T_0(k) = T_1(k)$. By considering all possible combinations of sensors in $T_1(k)$ within the cost limit, one selects $S_c(k)$.

It may be that the number of sensors used at the current stage becomes unacceptable due to a large allowable cost limit: the number of combinations of sensors become too high. A similar situation is that there are many sensors monitoring each component, which is tantamount to each row of the R matrix having many ones—that is, the tests, taken individually, are not specific. In such cases $|T_1(k)|$ will not be significantly decreased even by using a smaller value for L . Both of these cases require another threshold (denoted N) limiting the number of sensors chosen to perform testing at each stage, i.e. $|S_c|$ is limited to N . Hence, if the size of $T_1(k)$ is larger than the given threshold M , the number of sensors used at each stage will be limited by N , where M is decided based on cost. In summary, the upper bound in $|S_c(k)|$ is

$$|S_c(k)| \leq \begin{cases} |T_1(k)| & \text{if } |T_1(k)| \leq M \\ N & \text{if } |T_1(k)| > M \end{cases}. \quad (19)$$

REMARK 1: If there are too many information states having the same probability, especially in the early stages, it becomes problematic to decide $T_1(k)$ just by using threshold L , since ordering becomes arbitrary. In this case we change the threshold: L is increased until all elements having the same probability are included in set $T_1(k)$ whereas threshold N is decreased depending on the size of $T_1(k)$.

A summary of the proposed fault diagnosis based on exhaustive search is summarized in Algorithm 2. An example of the sets T_0 and T_1 is illustrated in Fig. 2.

ALGORITHM 2 Exhaustive Search—ES(L, M, N)

- 1) If the prior probability is uniform, select the combination of sensors chosen off-line as $S_c(1)$. Otherwise, go to the next step.
- 2) Get $T_1(k)$ by using information states $\{\pi_i(k)\}$. If there are many information states having the same value, L is increased until all sensors monitoring those elements are included in $T_1(k)$.
- 3) If the number of elements in $T_1(k)$ is higher than M , obtain all possible combinations of at most N sensors satisfying the cost limit C_T . Otherwise, obtain all possible combination of sensors satisfying the limit.
- 4) Select $S_c(k)$ from the obtained combinations of sensors based on (18) and obtain test outcomes of $S_c(k)$.

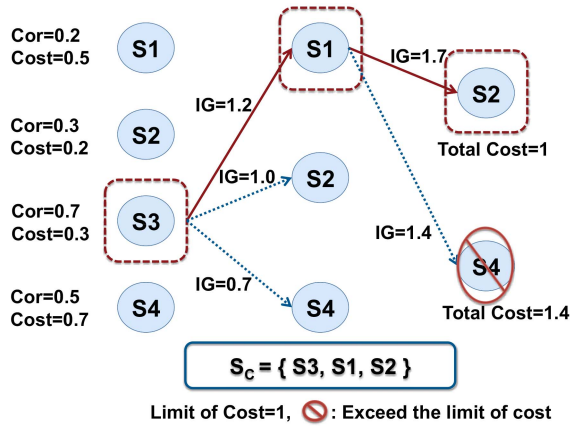


Fig. 3. An example of selecting $S_c(k)$ according to the CS approach.

- 5) Update information states using (5) and obtain information states $\{\pi(k+1)\}$ by applying the pruning criterion in (8).
- 6) Repeat steps 2)–5) until the stopping rule in (7) is satisfied.

3.2. Algorithm based on Correlative Search (CS)

While in the previous algorithm the sensors used at a single stage are selected jointly, this algorithm chooses these sensors one-by-one, which means that each sensor is added to the set of sensors chosen before as shown in Fig. 3. Before discussing how to choose sensors, let us define the correlation between matrix R and information state as follows. For each sensor j ,

$$\text{Cor}(j) = \sum_i r_{ij} \pi_i r_{ij}$$

The first sensor s_{j_1} is chosen based on correlation as follows:

$$s_{j_1} = \arg \max_{s_j \in \tilde{T}_1} \text{Cor}(j) \quad (20)$$

where

$$\tilde{T}_1 = \{s_q \in T_0 \mid c_q \leq C_T\}.$$

The second sensor is the one having the highest information gain calculated with the first sensor together:

$$s_{j_2} = \arg \max_{s_q \in \tilde{T}_2} \text{IG}(\{\pi_i\} \mid \{s_{j_1}, s_q\}) \quad (21)$$

where

$$\tilde{T}_2 = \{s_q \in T_0 \mid c_q \leq (C_T - c_{j_1})\} - \{s_{j_1}\}.$$

Assuming set \tilde{T}_p is nonempty, the next sensor is selected in the same way as the second sensor:

$$s_{j_p} = \arg \max_{s_q \in \tilde{T}_p} \text{IG}(\{\pi_i\} \mid \{s_{j_1}, s_{j_2}, \dots, s_{j_{p-1}}, s_q\}), \quad (22)$$

where

$$\tilde{T}_p = \{s_q \in T_0 \mid c_q \leq (C_T - \sum_{a=1}^{p-1} c_{j_a})\} - \{s_{j_1}, s_{j_2}, \dots, s_{j_{p-1}}\}.$$

Similar to the previous algorithm a threshold N is used to limit the maximum number of sensors in $S_c(k)$, i.e., such that $|S_c(k)| < N$.

REMARK 2: If $\text{IG}(\{\pi_i\} \mid \{s_{j_1}, s_{j_2}, \dots, s_{j_{p-1}}\}) = \text{IG}(\{\pi_i\} \mid \{s_{j_1}, s_{j_2}, \dots, s_{j_p}\})$, which means that the information gain does not increase by adding more sensors, no more sensors will be added.

REMARK 3: It should be noted that the set T_0 can be replaced by set T_1 before correlative selection.

A summary of the proposed fault diagnosis based on correlative search is in Algorithm 3. An example of selecting the set $S_c(k)$ is shown in Fig. 3.

ALGORITHM 3 Correlative Search—CS(N)

- 1) If the prior probability is uniform, select the combination of sensors chosen off-line as $S_c(1)$. Otherwise, go to the next step.
- 2) After getting $T_0(k)$ by using information states $\{\pi_i(k)\}$, select the first sensor s_{j_1} using (20).
- 3) Repeat: select p th sensor s_{j_p} using (22) until one of the following is satisfied:
 - p reaches N
 - there are no more sensors to add.
 - the information gain between information states and sensors is not increased by adding another (any other) sensor.
- 4) Set $S = \tilde{T}_p$ and select $S_c(k)$ based on (13).
- 5) Obtain the sensor outcomes of $S_c(k)$ and update information states using (5) and obtain information states $\{\pi(k+1)\}$ by applying the pruning rule in (8).
- 6) Repeat steps 2)–5) until the stopping rule (7) is satisfied.

4. SIMULATION RESULTS

In this section the performance of the proposed methods are analyzed through simulations. The following three algorithms are considered in the simulations.

- ExS(N) for $N = 1$ and 2.
- ES(L, M, N) for $(L, M, N) = (20, 10, 4)$ and $(50, 30, 2)$.
- CS(N) for $N = 2, 3$ and 4.

4.1. Randomly generated R matrix

We have generated two different R matrices each with 100 rows and 70 columns, i.e., there are 100 possible faults and 70 sensors monitoring those faults in the simulated system. Each element of the first R matrix, denoted R_2 , is generated as Bernoulli with success probability 0.2. In a similar fashion each element of the second R matrix, denoted as R_8 , used success probability 0.8. The cost of the test by each sensor is assigned randomly following a uniform distribution between 0.5 and 1. The maximum cost allowed to be spent at each stage is 3. The stopping rule is defined using a level of uncertainty $\alpha = 0.99$ and the pruning criterion is defined with threshold $\beta = 0.005$.

Now, we define some performance metrics in order to compare various algorithms. The *average number of stages* \bar{k} of a certain algorithm is obtained by repeating the algorithm over several Monte-Carlo runs and averaging the number of stages it took each time to locate the fault, i.e.,

$$\bar{k} = \frac{1}{N_m} \sum_{r=1}^{N_m} k_r \quad (23)$$

where k_r is the number of stages it took to locate the fault at the r th run and N_m is the number of Monte-Carlo runs. In all the simulations $N_m = 1000$ Monte-Carlo runs are used.

Figure 4 shows the average number of stages \bar{k} of all the algorithms as the probability of detection P_d increases from 0.8 to 1. In particular, Fig. 4(a) shows the comparison of \bar{k} vs. P_d of ExS and ES on different R matrices. The same is shown in Fig. 4(b) for ExS and CS. The summary of Fig. 4—comparison of all three algorithms—is shown in Fig. 4(c) which shows that the exhaustive search (ES) algorithm performs as well as the exhaustive search ExS(2), whereas the correlative search CS(N) outperforms the exhaustive search ExS(2) with increasing N .

Figure 5 repeats the simulation analysis shown in Fig. 4 for various false alarm rates while fixing P_d at 0.99. Similar conclusion is arrived from Fig. 5 as well, where it can be noticed in the summary Fig. 5(c) that the ES algorithms outperform ExS(1), and perform essentially as well as ExS(2). On the other hand, CS(N) outperforms the others with increasing N .

4.2. Real System

Our algorithms are applied to a real system, the so-called ‘‘Documatch,’’ which is an R -matrix representation of the Pitney Bowes Integrated Mail System that takes an original document from a Microsoft Windows based personal computer and turns it into a finished and properly-addressed mail item in a sealed envelope (see, for example, [6]). The R matrix of this system, denoted hereafter as R_d , has 258 components and 179 sensors.

First, let us compare the simulated R matrices used in earlier simulations and the R_d of the Documatch system. Out of the simulated R matrices used in this section, we select R_2 for comparison with R_d . For each of these matrices we counted the *number of components monitored by each sensor* and the *number of sensors monitoring each component*. The result is summarized in Fig. 6.

The performance comparison of all three algorithms in terms of \bar{k} vs. P_d is shown in Fig. 7. Due to the size of the R_d matrix, the exhaustive search is performed only for $N = 1$, i.e., we do not have ExS(2) in the simulations because of the time required to complete the simulation. The figure confirms the earlier conclusions arrived through the simulated R matrices.

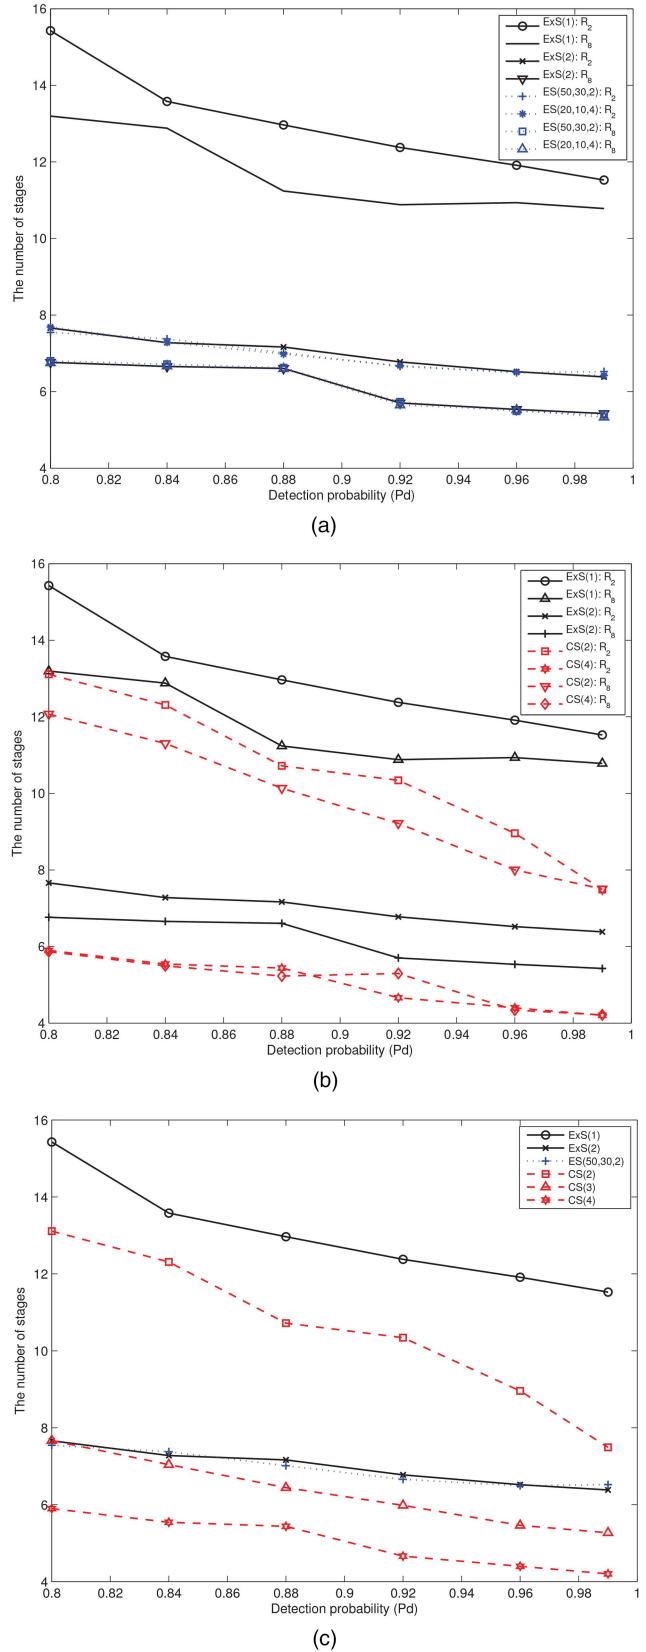


Fig. 4. Comparison of algorithms, in terms of the average number of stages \bar{k} vs. probability of detection P_d . The false alarm rate is fixed at $P_f = 0.01$ in all the figures. (a) Comparison of ExS and ES on matrices R_2 and R_8 . (b) Comparison of ExS and CS on matrices R_2 and R_8 . (c) Comparison of all algorithms, ExS, ES and CS, on matrix R_2 .

4.3. Computational complexity analysis

The computational cost arises mainly based on how many combinations of sensors there are, i.e., the size of the candidate set for calculation of information gain. For example, if the system has 100 sensors, with the exhaustive search in ExS(2), we need to calculate information gain for $\sum_{p=1}^2 100!/((100-p)!p!) = 5050$ combinations of sensors at each stage, as 2 sensors are allowed to be tested at each stage and all possible combinations of two sensors can be used to test in the respect of the limit of cost. It means that $\tilde{S}_c(k) \in \{\text{all possible combinations of two sensors}\}$ and $|\{\tilde{S}_c(k)\}| = 5050$. If it takes 2 stages to isolate the faulty element, i.e. $k_r = 2$ for a particular run r , the total number of combinations (i.e. $n_t := \sum_{k=1}^{k_r} |\{\tilde{S}_c(k)\}|$) is 10100. Let us consider another example: Assume that only one sensor is allowed to test at each stage and it takes 110 stages to isolate the faulty component; in this case $n_t = 11000$. It should be noted that, in a real situation, due to pruning the number of these combinations $|\{\tilde{S}_c(k)\}|$ varies for each stage. In summary, n_t is the accumulated size of candidate set of $S_c(k)$ over all stages until the faulty component is isolated for a particular run. The average of n_t^r is defined as N_t , where

$$N_t = \frac{1}{N_m} \sum_{r=1}^{N_t} n_t^r, \quad (24)$$

in which

$$n_t^r = \sum_{k=1}^{k_r} |\{\tilde{S}_c^r(k)\}|. \quad (25)$$

The comparison of all the algorithms in terms of N_t vs. P_d is shown in Fig. 8. It shows that CS(N) requires much reduced computation compared to ExS(2) for $N = 2, 3$, and 4. Further it must be emphasized that N_t reduces with increasing N for CS whereas N_t increases with increasing N for ExS(N). At this point it must be re-emphasized that the CS(N) algorithm also reduces the average number of iterations required to isolate the faulty element with increasing number of N (see Figs. 4 and 5) so that N_t (or computations) is reduced even with increasing N . Fig. 9 shows the comparison of all the algorithms in terms of N_t vs. P_f for fixed P_d . It shows that CS(N) requires a similar amount of—and at times lower—computation compared to ExS(1).

The comparison of N_t for the Documatch system is shown in Fig. 10. Due to the size of the system the computational complexity of ExS(2) was not computed. The figure confirms the results of the experiments performed on the R_2 and R_8 matrices. It confirms that CS(N) requires slightly less computation with increasing N than ExS(1), especially with higher probability of detection P_d and lower false alarm rate P_f .

5. CONCLUSIONS

Two algorithms for single fault diagnosis under communication constraints are presented. Both of the al-

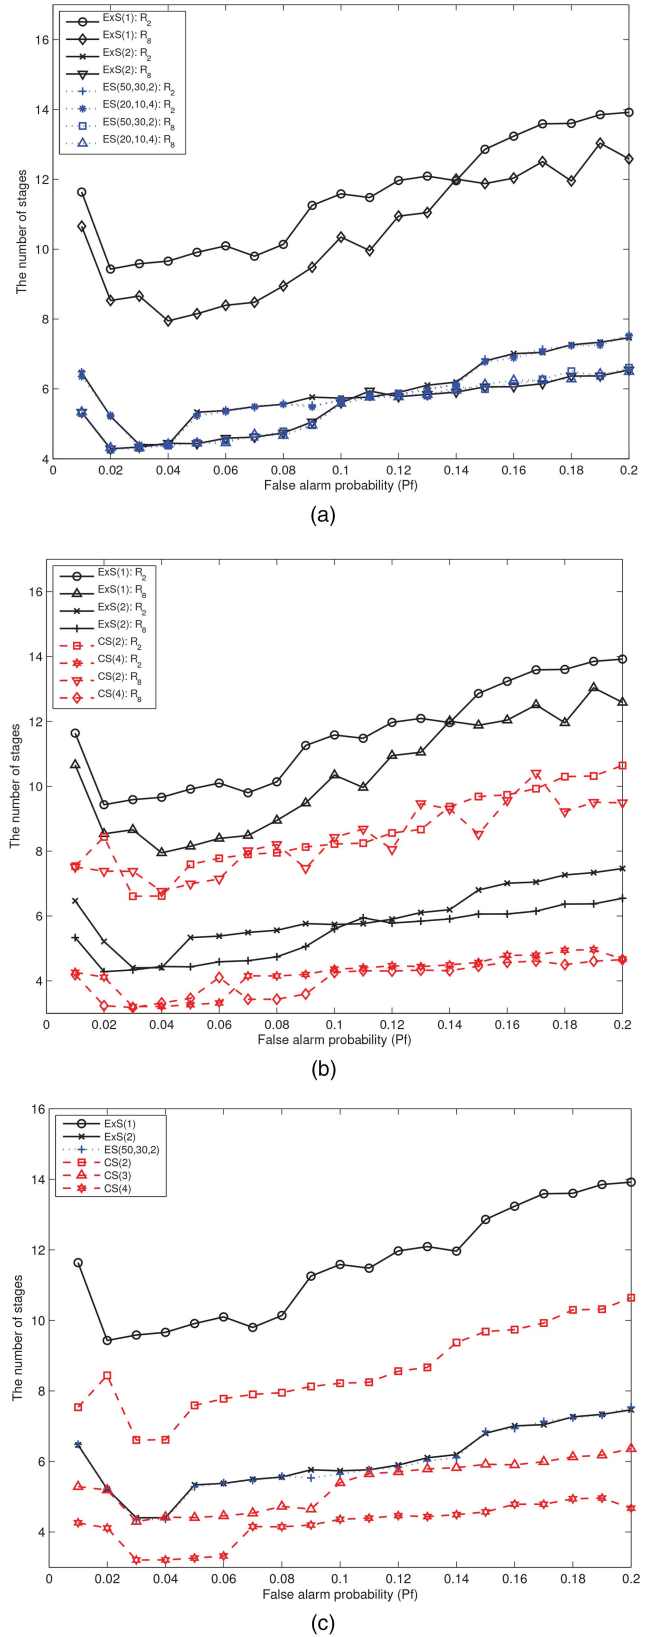


Fig. 5. Comparison of algorithms, in terms of the average number of stages k vs. probability of false-alarm P_f . The probability of detection is fixed at $P_d = 0.99$ in all the figures. (a) Comparison of ExS and ES on matrices R_2 and R_8 . (b) Comparison of ExS and CS on matrices R_2 and R_8 . (c) Comparison of all algorithms, ExS, ES and CS, on matrix R_2 .

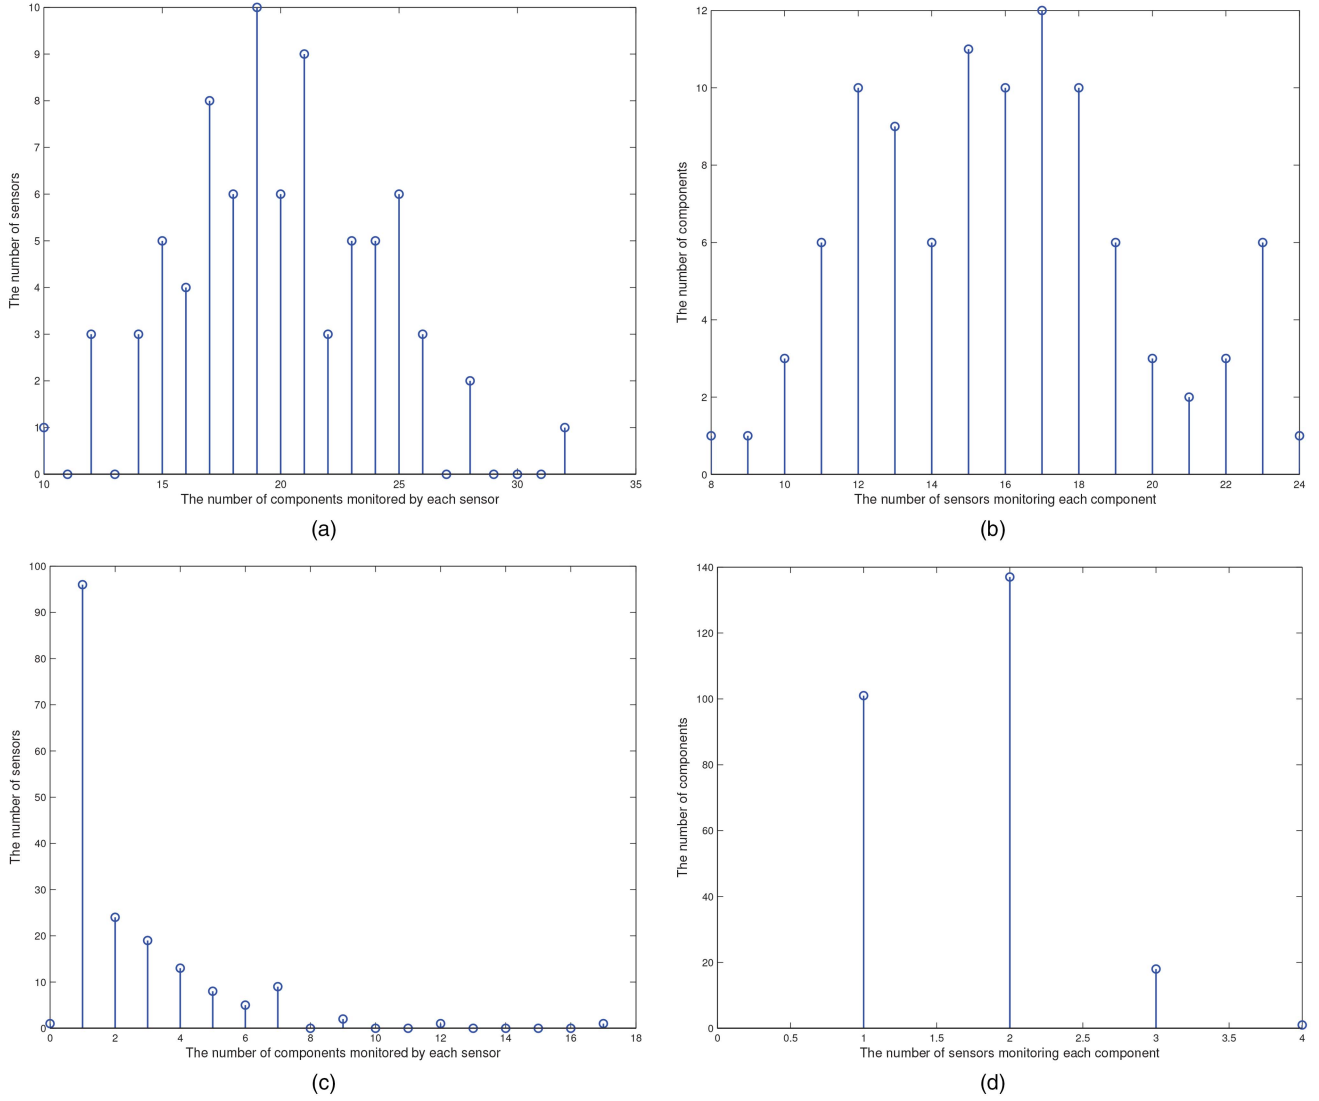


Fig. 6. Summary of comparisons: Simulated R_2 matrix vs. practical Documatch R_d matrix. (a) Simulated R_2 matrix. (b) Simulated R_2 matrix. (c) Documatch R_d matrix. (d) Documatch R_d matrix.

gorithms are concerned with the selection of (possibly) several sensors at a time in order to reduce the number of iterations/communication. The algorithms differ in their approach in selecting multiple sensors that maximize the information gain. The first algorithm, termed the exhaustive search (ES) method, introduces several thresholds in order to eliminate sensors that are less informative, so that fewer sensors form the candidate set for the maximization of information gain. The second approach, referred to as correlative search (CS), selects the candidate set one-by-one based on the correlation between the information state and the elements of the reachability matrix. Both of the proposed approaches demonstrate their ability to reduce the number of iterations in fault diagnosis.

APPENDIX

Given a subset S_c of S , $\tilde{\mathbf{o}}_t$ denotes a vector of outcomes of the sensors in a set S_c , and the set $\tilde{O} = \{\tilde{\mathbf{o}}_t\}$




Fig. 7. Comparison of all algorithms in terms of \bar{k} vs. P_d on Documatch system for $P_f = 0.01$.



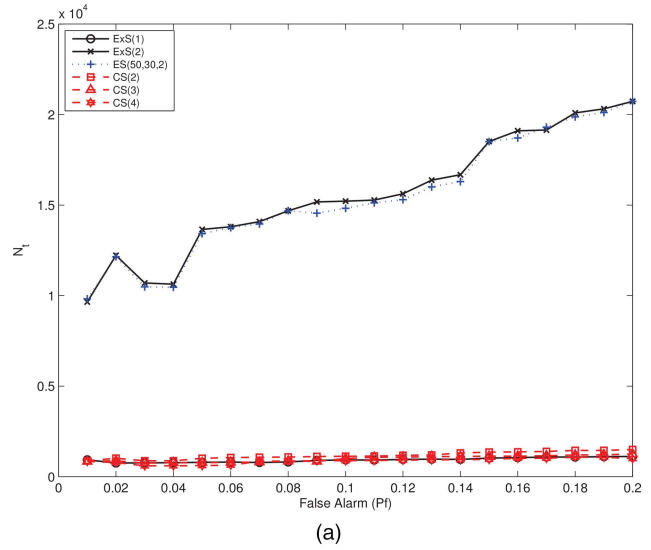



Fig. 8. Comparison of all algorithms on R_2 for varying P_d with $P_f = 0.01$. (a) Comparison of N_t for various algorithms. (b) Closer view of the bottom portion of Fig. 8(a) above.

Fig. 9. Comparison of all algorithms on R_2 for varying P_f with $P_d = 0.99$. (a) Comparison of N_t for various algorithms. (b) Closer view of the bottom portion of Fig. 9(a) above.

consists of all possible vector which can be generated by the sensors. And J is the set of the indices of sensors in the set S_c . The information gain is defined with information state as

$$IG(\{\pi_i(k)\}, \mathbf{S}(k)) = H(\{\pi_i(k)\}) - H(\{\pi_i(k)\} | S_c(k)) \quad (26)$$

where

$$H(\{\pi_i(k)\}) = - \sum_i \pi_i(k) \log \pi_i(k). \quad (27)$$

The following is obtained:

$$\begin{aligned} IG(\{\pi_i(k)\}, S_c(k)) &= - \sum_i \pi_i(k) \log \pi_i(k) \\ &+ \sum_{\tilde{\mathbf{o}}_t \in \tilde{\mathcal{O}}} \sum_{i=0}^m \pi_i(k) \prod_{j \in J} (\tilde{\mathbf{o}}_t(j) d_{i,j} + (1 - \tilde{\mathbf{o}}_t(j))(1 - d_{i,j})) \end{aligned}$$

$$\begin{aligned} &\times \log \left\{ \pi_i(k) \prod_{q \in J} (\tilde{\mathbf{o}}_t(q) d_{i,q} + (1 - \tilde{\mathbf{o}}_t(q))(1 - d_{i,q})) \right\} \\ &- \sum_{\tilde{\mathbf{o}}_t \in \tilde{\mathcal{O}}} \sum_{i=0}^m \pi_i(k) \prod_{j \in J} (\tilde{\mathbf{o}}_t(j) d_{i,j} + (1 - \tilde{\mathbf{o}}_t(j))(1 - d_{i,j})) \\ &\times \log \left\{ \sum_{p=0}^m \prod_{q \in J} \pi_i(k) (\tilde{\mathbf{o}}_t(q) d_{p,q} + (1 - \tilde{\mathbf{o}}_t(q))(1 - d_{p,q})) \right\}. \end{aligned} \quad (28)$$

PROOF The conditional entropy of information state is described by

$$H(\{\pi_i(k)\} | S_c) = H(\{\pi_i(k)\} | S_c, I_{k-1}) \quad (29)$$

$$= \sum_{\tilde{\mathbf{o}}_t \in \tilde{\mathcal{O}}} \text{Prob}(\tilde{\mathbf{o}}_t | I_{k-1}) H(\{\pi_i(k)\} | \tilde{\mathbf{o}}_t, I_{k-1}). \quad (30)$$

The entropy $H(\{\pi_i(k)\} | \tilde{\mathbf{o}}_t, I_{k-1})$ is given as

$$\begin{aligned}
& H(\{\pi_i(k)\} | \tilde{\mathbf{o}}_t, I_{k-1}) \\
&= - \sum_{i=0}^m \text{Prob}(\pi_i(k) | \tilde{\mathbf{o}}_t, I_{k-1}) \log(\text{Prob}(\pi_i(k) | \tilde{\mathbf{o}}_t, I_{k-1})) \\
&= - \sum_{i=0}^m \frac{\pi_i(k)}{\text{Prob}(\tilde{\mathbf{o}}_t | I_k)} \prod_{j \in J} (\tilde{\mathbf{o}}_t(j) d_{ij} + (1 - \tilde{\mathbf{o}}_t(j))(1 - d_{ij})) \\
&\quad \times \log \left\{ \frac{\pi_i(k)}{\text{Prob}(\tilde{\mathbf{o}}_t | I_k)} \prod_{j \in J} (\tilde{\mathbf{o}}_t(j) d_{ij} + (1 - \tilde{\mathbf{o}}_t(j))(1 - d_{ij})) \right\}
\end{aligned} \tag{31}$$

where

$$\begin{aligned}
& \text{Prob}(\pi_i(k) | \tilde{\mathbf{o}}_t, I_{k-1}) \\
&= \text{Prob}(\tilde{\mathbf{o}}_t | \pi_i(k), I_{k-1}) \frac{\text{Prob}(\pi_i(k) | I_{k-1})}{\text{Prob}(\tilde{\mathbf{o}}_t | I_{k-1})} \\
&= \frac{\pi_i(k)}{\text{Prob}(\tilde{\mathbf{o}}_t | I_{k-1})} \prod_{j \in J} (\tilde{\mathbf{o}}_t(j) d_{ij} + (1 - \tilde{\mathbf{o}}_t(j))(1 - d_{ij}))
\end{aligned} \tag{32}$$

and

$$\begin{aligned}
& \text{Prob}(\tilde{\mathbf{o}}_t | I_{k-1}) \\
&= \sum_{i=0}^m \prod_{j \in J} \text{Prob}(\tilde{\mathbf{o}}_t(j) | x_i = 1, I_{k-1}) \text{Prob}(x_i = 1 | I_{k-1}) \\
&= \sum_{i=0}^m \prod_{j \in J} \pi_i(k) (\tilde{\mathbf{o}}_t(j) d_{ij} + (1 - \tilde{\mathbf{o}}_t(j))(1 - d_{ij})).
\end{aligned} \tag{33}$$

Thus,

$$\begin{aligned}
& H(\{\pi_i(k)\} | S_c) \\
&= \sum_{\tilde{\mathbf{o}}_t \in \tilde{\mathcal{O}}} \text{Prob}(\tilde{\mathbf{o}}_t | I_{k-1}) H(\{\pi_i(k)\} | \tilde{\mathbf{o}}_t, I_{k-1}) \\
&= - \sum_{\tilde{\mathbf{o}}_t \in \tilde{\mathcal{O}}} \sum_{i=0}^m \pi_i(k) \prod_{j \in J} (\tilde{\mathbf{o}}_t(j) d_{ij} + (1 - \tilde{\mathbf{o}}_t(j))(1 - d_{ij})) \\
&\quad \times \log \left\{ \frac{\pi_i(k)}{\text{Prob}(\tilde{\mathbf{o}}_t | I_k)} \prod_{j \in J} (\tilde{\mathbf{o}}_t(j) d_{ij} + (1 - \tilde{\mathbf{o}}_t(j))(1 - d_{ij})) \right\} \\
&\quad + \sum_{\tilde{\mathbf{o}}_t \in \tilde{\mathcal{O}}} \sum_{i=0}^m \pi_i(k) \prod_{j \in J} (\tilde{\mathbf{o}}_t(j) d_{ij} + (1 - \tilde{\mathbf{o}}_t(j))(1 - d_{ij})) \\
&\quad \times \log \left\{ \sum_{p=0}^m \prod_{q \in J} \pi_i(k) (\tilde{\mathbf{o}}_t(q) d_{p,q} + (1 - \tilde{\mathbf{o}}_t(q))(1 - d_{p,q})) \right\}
\end{aligned} \tag{34}$$

and hence the information gain is as in (28).

REFERENCES

- [1] D. Blough and A. Pelc. Complexity of fault diagnosis in comparison models. *IEEE Transactions on Computers*, 41(3):318–324, Mar. 1992.





Fig. 10. Comparison of all algorithms on Documatch for varying P_d with $P_f = 0.01$. (a) Comparison of N_t for various algorithms. (b) Closer view of the bottom portion of Fig. 10(a) above.

- [2] R. Boumen, I. S. M. de Jong, J. W. H. Vermunt, J. M. van de Mortel-Fronczak, and J. E. Rooda. Test sequencing in complex manufacturing systems. *IEEE Transactions on Systems, Man, and Cybernetics: Part A*, 38(1):25–37, Jan. 2008.
- [3] R. Boumen, S. Ruan, I. S. M. de Jong, J. M. van de Mortel-Fronczak, J. E. Rooda, and K. R. Pattipati. Hierarchical test sequencing for complex systems. *IEEE Transactions on Systems, Man, and Cybernetics: Part A*, 39(3):640–649, May 2009.
- [4] J. M. Christensen and J. M. Howard, editors. *Human Detection and Diagnosis of System Failure*, chapter Field experience in maintenance. New York: Plenum, 1981.
- [5] O. Erdinc, C. Brideau, P. Willett, and T. Kirubarajan. Fast diagnosis with sensors of uncertain quality. *IEEE Transactions on Systems, Man, and Cybernetics: Part B*, 38(4):1157–1165, 2008.
- [6] O. Erdinc, C. Brideau, P. Willett, and T. Kirubarajan. The problem of test latency in machine diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics—Part A*, 38(1):88–92, 2008.
- [7] M. R. Garey. Optimal binary identification procedures. *SIAM Journal on Applied Mathematics*, 23(2):173–186, 1972.
- [8] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. New York: W. H. Freeman and Company, 1979.

- [9] L. Hyafil and R. L. Rivest. Constructing optimal binary decision trees in NP-complete. *Information Processing Letters*, 5(1):15–17, 1976.
- [10] T. Le and C. Hadjicostis. Max-product algorithms for the generalized multiple-fault diagnosis problem. *IEEE Transactions on Systems, Man, and Cybernetics: Part B*, 37(6):1607–1621, Dec. 2007.
- [11] Z. Mao, B. Jiang, and P. Shi. H_∞ fault detection filter design for networked control systems modelled by discrete Markovian jump systems. In *2010 IEEE International conference on Information and Automation*, Jun. 2010.
- [12] M. Mastrianni. Virtual integrated test and diagnosis. *Sikorsky Pres. ARPA*, Nov. 9 1994.
- [13] B. M. E. Moret. Decision trees and diagrams. *Computing Surveys*, 14(4):593–623, 1982.
- [14] K. R. Pattipati and M. G. Alexandridis. Application of heuristic search and information theory to sequential fault diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(4):872–887, 1990.
- [15] V. Raghavan, M. Shakeri, and K. R. Pattipati. Optimal and near-optimal test sequencing algorithms with realistic test models. *IEEE Transactions on Systems, Man, and Cybernetics: Part A*, 29(1):11–26, 1999.
- [16] V. Raghavan, M. Shakeri, and K. R. Pattipati. Test sequencing algorithms with unreliable tests. *IEEE Transactions on Systems, Man, and Cybernetics: Part A*, 29(4):347–357, 1999.
- [17] S. Ruan, Y. Zhou, F. Yu, K. R. Pattipati, P. Willett, and A. Patterson-Hine. Dynamic multiple-fault diagnosis with imperfect tests. *IEEE Transactions on Systems, Man, and Cybernetics: Part A*, 39(6):1224–1236, Nov. 2009.
- [18] M. Shakeri, K. R. Pattipati, V. Raghavan, and A. Patterson-Hine. Optimal and near-optimal algorithms for multiple fault diagnosis with unreliable tests. *IEEE Transactions on Systems, Man, and Cybernetics: Part C*, 28(3):431–440, 1998.
- [19] M. Shakeri, V. Raghavan, K. R. Pattipati, and A. Patterson-Hine. Sequential testing algorithms for multiple fault diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics: Part A*, 30(1):1–14, 2000.
- [20] S. Singh, A. Kodali, K. Choi, K. Pattipati, S. M. Namburu, S. C. Sean, D. V. Prokhorov, and L. Qiao. Dynamic multiple fault diagnosis: Mathematical formulations and solution techniques. *IEEE Transactions on Systems, Man, and Cybernetics: Part A*, 39(1):160–176, 2009.
- [21] F. Tu and K. R. Pattipati. Rollout strategy for sequential fault diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics: Part A*, (1):86–99, 2003.
- [22] Y. Wang, S. Ding, H. Ye, L. Wei, P. Zhang, and G. Wang. Fault detection of networked control systems with packet based periodic communication. *International Journal of Adaptive Control and Signal Processing*, 23(8):682–698, 2009.
- [23] Y. Wang, H. Ye, S. Ding, Y. Cheng, P. Zhang, and G. Wang. Fault detection of networked control systems with limited communication. *International Journal of Control*, 82(7):1344–1356, 2009.
- [24] J. Ying, T. Kirubarajan, K. R. Pattipati, and A. Patterson-Hine. A hidden Markov model-based algorithm for fault diagnosis with partial and imperfect tests. *IEEE Transactions on Systems, Man, and Cybernetics: Part C*, 30(4), 2000.
- [25] P. Zhang, S. X. Ding, P. M. Frank, and M. Sader. Fault detection of networked control systems with missing measurements. In *5th Asian Control Conference*, volume 2, pages 1258–1263, Jul. 2004.
- [26] P. Zhang, S. X. Ding, G. Z. Wang, and D. H. Zhou. Fault detection of linear discrete-time periodic systems. *IEEE Transactions on Automatic Control*, 50(2):239–244, 2005.



Sora Choi received her Ph.D. degree, in Mathematics, from Seoul National University, Seoul, Korea in 2003. She was a researcher at the Electronics and Telecommunications Research Institute (ETRI), Daejeon, Korea from 2003 to 2004, and was an instructor at Seoul National University between 2000 and 2002, Kyunghee and Konkuk Universities, in Korea, in 2007. She received a Graduate Predoctoral Fellowship in 2011 and is currently a graduate student/research assistant at the University of Connecticut, Storrs, where she is pursuing a Ph.D. degree in Electrical Engineering. Her research interests are in target tracking, detection theory, signal processing, and fault diagnosis.



Balakumar Balasingam earned his Ph.D. from McMaster University, Canada in 2008. Currently, he is an assistant research professor at the University of Connecticut in the Department of Electrical and Computer Engineering. Before joining UConn in 2010 as a postdoctoral fellow, he has worked as a postdoctoral fellow at the University of Ottawa, Canada from 2008. His research interests include signal processing, distributed information fusion and machine learning.



Peter Willett has been a faculty member in the Electrical and Computer Engineering Department at the University of Connecticut since 1986, following his Ph.D. from Princeton University. His areas include statistical signal processing, detection, sonar/radar, communications, data fusion, and tracking. He is an IEEE Fellow. He was EIC for *IEEE Transactions on AES*, and remains AE for the *IEEE AES Magazine* plus ISIF's *Journal of Advances in Information Fusion* and senior editor for *IEEE Signal Processing Letters*, and is a member of the editorial board of IEEE's Special Topics in Signal Processing. He is a member of the IEEE AESS Board of Governors (and is now VP Pubs) and of the IEEE Signal Processing Society's Sensor-Array and Multichannel (SAM) technical committee (and is now Vice Chair). He has served as Technical Chair for the 2003 SMCC and the 2012 SAM workshop; and as Executive/General Chair for the 2006, 2008 and 2011 FUSION conferences.

INTERNATIONAL SOCIETY OF INFORMATION FUSION

ISIF Website: <http://www.isif.org>

2013 BOARD OF DIRECTORS*

2011–2013	2012–2014	2013–2015
Sten F. Andler	Darin T. Dunham	Jean Dezert
Yvo Boers	Fredrik Gustafsson	Gee-Wah Ng
Lyudmila Mihaylova	Lance M. Kaplan	Anne-Laure Jousset

*Board of Directors are elected by the members of ISIF for a three year term.

PAST PRESIDENTS

Roy Streit, 2012	Erik Blasch, 2007	Yaakov Bar-Shalom, 2002
Joachim Biermann, 2011	Pierre Valin, 2006	Pramod Varshney, 2001
Stefano Coraluppi, 2010	W. Dale Blair, 2005	Yaakov Bar-Shalom, 2000
Elisa Shahbazian, 2009	Chee Chong, 2004	Jim Llinas, 1999
Darko Musicki, 2008	Xiao-Rong Li, 2003	Jim Llinas, 1998

SOCIETY VISION

The International Society of Information Fusion (ISIF) is the premier professional society and global information resource for multidisciplinary approaches for theoretical and applied information fusion technologies.

SOCIETY MISSION

Advocate

To advance the profession of fusion technologies, propose approaches for solving real-world problems, recognize emerging technologies, and foster the transfer of information.

Serve

To serve its members and engineering, business, and scientific communities by providing high-quality information, educational products, and services.

Communicate

To create international communication forums and hold international conferences in countries that provide for interaction of members of fusion communities with each other, with those in other disciplines, and with those in industry and academia.

Educate

To promote undergraduate and graduate education related to information fusion technologies at universities around the world. Sponsor educational courses and tutorials at conferences.

Integrate

Integrate ideas from various approaches for information fusion, and look for common threads and themes—look for overall principles, rather than a multitude of point solutions. Serve as the central focus for coordinating the activities of world-wide information fusion related societies or organizations. Serve as a professional liaison to industry, academia, and government.

Disseminate

To propagate the ideas for integrated approaches to information fusion so that others can build on them in both industry and academia.

Call for Papers

The Journal of Advances in Information Fusion (JAIF) seeks original contributions in the technical areas of research related to information fusion. Authors of papers in one of the technical areas listed on the inside cover of JAIF are encouraged to submit their papers for peer review at <http://jaif.msubmit.net>.

Call for Reviewers

The success of JAIF and its value to the research community is strongly dependent on the quality of its peer review process. Researchers in the technical areas related to information fusion are encouraged to register as a reviewer for JAIF at <http://jaif.msubmit.net>. Potential reviewers should notify via email the appropriate editors of their offer to serve as a reviewer.