ISIF

Perspectives

On Information Fusion

June 2023
Volume 6
Number 1

GOING FOR GOLD: FINDING LOST GOLD CITIES IN ECUADOR

FUSION PROCESSES AND SITUATION CONTROL

25 YEARS OF ISIF: KEY DEVELOPMENTS, MEMORIES AND TRIBUTES

Publication of the
INTERNATIONAL SOCIETY OF INFORMATION FUSION

ISIF 25th Anniversary

ISIF Perspectives

Perspectives seeks bridging articles, expository papers and tutorials, classroom notes, and announcements on topics of general interest to the ISIF FUSION community. Fresh points of view on established topics are especially welcome, as are articles on topics of interest to the ISIF annual FUSION conference. Papers containing significant original research should be directed to the ***Journal of Advances in Information Fusion*** (JAIF) or another research journal. The standing Call for Papers (CfP) for ***Perspectives*** is posted at https://isif.org/.

More detailed guidelines and submission instructions for authors may be found at http://perspectives.msubmit.net/cgi-bin/main.plex?form_type=display_auth_instructions. The average length for submissions is approximately six (6) pages (in JAIF two-column format). All submissions will be reviewed for content and style, as well as suitability for ***Perspectives***. All papers accepted for publication will be written in a relaxed, colloquial style that facilitates understanding by a wide audience.

Cover: ISIF Perspectives Editorial Board. Group pictures FUSION 2001, Paris; FUSION 2006, Florence. Title band artwork in the Key Development papers provided courtesy of © Fraunhofer FKIE: iStock® vadimmmus / 123RF® majcot

Perspectives (ISSN: 2831-4824, print; ISSN: 2831-4832, online) is published annually by the International Society of Information Fusion (ISIF). 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.

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

Table of Contents

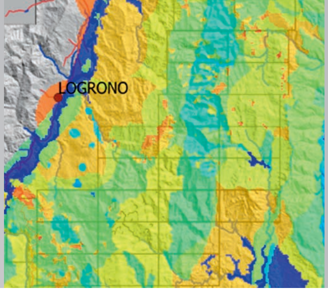

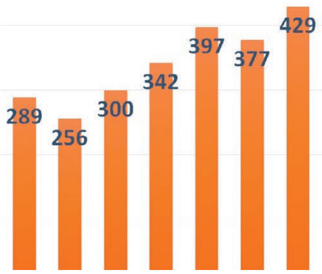

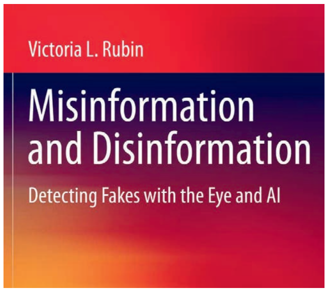

ISIF

PERSPECTIVES MAGAZINE

Welcome to the sixth issue of *Perspectives* magazine! This issue is particularly special as it celebrates the 25th anniversary of the International Society of Information Fusion (ISIF), which has grown, matured, and expanded its scope to embrace new fields and research over the course of 26 conferences and 25 years.

It is my great pleasure to be your guide on this journey as you read these lines, possibly while attending FUSION 2023 in Charleston, South Carolina. The conference is returning to the birthplace of the Society, the United States, where the first FUSION conference took place in Las Vegas, Nevada in 1998.

We start this journey with Larry Stone (Metron, USA) and Keith Barron (Aurania, Canada), who take us on the path of Gold Cities. Their feature paper "Going for Gold: Finding Lost Gold Cities in Ecuador" tells a thrilling, successful story illustrating how Bayesian search theory methods were used to fuse old Spanish records and maps with modern geophysics and geochemistry to find the location one of seven "Gold Cities" of the 16th Century.

The second feature paper by Jim Llinas from the University at Buffalo presents ideas preparing the future. In his paper "Fusion Processes and Situation Control", he highlights ISIF's research heritage in situation assessment and proposes an expanded framework for situation control, allowing for more systemic research in the future.

The core of this issue is a dedication to the 25th anniversary of ISIF. I invite you to read the paper on how ISIF evolved over the years from nothing to its current state, written by Chee-Yee Chong with the ISIF Board of Directors, as well as reflections from ISIF presidents on their presidency term(s). Notably, this section features a series of short papers by nearly 20 authors, summarizing several key developments on specific topics related to advances in information fusion since the inception of ISIF. Their purpose is to present the reader with the main challenges pertaining to each topic and highlight the advances made over the last 25 years. They are not intended to be comprehensive survey papers, so the reference lists may not be exhaustive, and the content reflects the authors' individual perspectives on the topic. As part of these key developments, a summary of the history of the Evaluation of Techniques for Uncertainty Representation is available in the section dedicated to Working Groups.

This 25th anniversary year is also the moment to pay tribute to departed Fusion Minds who have left a lasting impact, inspired young ISIF members in their scientific pursuits, and supported ISIF in many ways over the years. A special thank you is extended to their colleagues, friends and family who shared words, photos, and emotions to immortalize them in this issue.

This issue concludes nicely with a very interesting review by Claire Laudy (from Thales, France) of the book *Misinformation and Disinformation: Detecting Fakes with the Eye and AI* by Victoria L. Rubin. You will also have the privilege of listening to Victoria Rubin as a keynote speaker at the FUSION 2023 conference.

In keeping with tradition, this issue features a report on the 25th FUSION conference, held in Linköping, Sweden in 2022. The general chairs Fredrik Gustafsson, Gustaf Hendeby (Linköping University, Sweden), and Terence van Zyl (University of Johannesburg, South Africa) deserve special appreciation for successfully returning the conference to almost normal physical attendance. Additionally, the ISIF Awards section recognizes the 2023 ISIF Lifetime of Excellence, Distinguished Service, and Young Investigator awards recipients, as well as the FUSION 2022 Best Paper and Best Student Paper awards. To keep the excitement alive, you are encouraged to flip through the pages to discover the awardees.

As you can imagine, the successful completion of this *Perspectives* Volume 6 issue would not have been possible without the contributions of the authors, the diligent reviews, and personal inputs of Associate Editors Lyudmila Mihaylova, Jesus García, Paulo Costa, Wolfgang Koch, Murat Efe, Emre Özkan, and the invaluable "lifetime guidance" of AEiC Roy Streit, the unwavering efforts of Production Manager Kristy Virostek in ensuring everything stayed on track, and the consistent support of Administrative Editor David W. Krout. I would like to express my deepest gratitude to Chee-Yee Chong for being our Guest Editor for *Perspectives* as we celebrate the 25th anniversary of ISIF. Between California, Hong Kong, and Egypt, he took time to write articles and coordinate the various contributions.

A sincere thank you to everyone for your dedication, ideas, and energy.

Anne-Laure Jousselme
Editor-in-Chief
Perspectives on Information Fusion

December 2022 Volume 17 Number 2 ISSN 1557-6418

JAIF 22

Journal of Advances in Information Fusion

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

Regular Papers Page

www.isif.org

ISIF

Going for Gold: Finding Lost Gold Cities in Ecuador

Abstract—In the 16th century, Spanish Conquistadors established seven "Gold Cities" in the Audiencia de Quito in the Viceroyalty of Peru in what is now Ecuador. Records in the Archives of the Indies in Sevilla, Spain and the Apostolic Library of the Vatican indicate that these cities were extremely lucrative gold producers. However, by the beginning of the 17th century, these cities had disappeared. Five of these were rediscovered by the end of the 20th century, but two of them, Sevilla del Oro and Logroño de los Caballeros, remained lost into the 21st century. In 2016, Aurania, a Canadian mining exploration company, began its search for these lost cities. On June 6, 2022, Aurania announced that Logroño, reputed to be the richest of the seven gold cities, had been found along the Rio Santiago in Ecuador in the high likelihood area of Metron's probability map for the location of Logroño. How did Metron develop this map? Answer: By applying Bayesian search theory methods to old Spanish records and maps and combining that information with modern geophysics and geochemistry.

Lawrence D. Stone
Metron
Reston, Virginia, USA
stone@metsci.com

Keith Barron
Aurania Resources
Toronto, ON, Canada
keith@aurania.com

FINDING LOGROÑO DE LOS CABALLEROS

In 2016, Aurania acquired mineral-rights concessions in the remote Cutucu region of Southeastern Ecuador, which includes the Rio Santiago. They launched their exploration effort, which they called The Lost Cities Project (see Figure 1), on the presumption that the lost cities were located in these concession areas. At that time, there were no known gold deposits along the Rio Santiago, no concessions staked along the river, and no mining activity. However, all that changed as the price of gold surged to over $2,000/oz USD. In January of 2022, Aurania learned of a November 2021 television news report, which revealed that a sophisticated and well-funded group of "invaders" had claim-jumped legitimately held concessions along the Rio Santiago.

The invaders employed more than 50 excavators along the river and were recovering tens of thousands of dollars in gold each day. Eventually, the Ecuadorian military impounded their heavy equipment and ran off most of the illegal miners. Presuming that today's miners are extracting only the dregs of the gold that was present in the 16th century, it is likely that 450 years ago the Rio Santiago was as rich as Bonanza Creek in the Klondike when it was first discovered. The richness of the area conforms with Governor Juan de Alderete's 16th century account that in the first year of mining, almost 30,000 pesos of gold were produced at Logroño. One peso was equal to 4.6 g of "buen oro" (22.5 carat purity), making 30,000 pesos equal to approximately 4,100 troy ounces. A 1591 document contains one man's report that in one week he and six other miners extracted 350 pesos (more than 1.5 kg) of gold by hand. He said under oath, "this land was the richest in gold of all the Kingdoms of Peru".

Figure 1
Logo for the Lost Cities Project.

Aurania's subsequent sampling upstream from the excavations showed abundant fine gold in every panning sample taken. This confirmed to Dr. Barron that this section of the Rio Santiago was the location of the lost gold city of Logroño de los Caballeros. Moreover, it was situated in the high probability area of the map produced by Metron for the location of that city. The alluvial gold must have washed down from the nearby Cordillera de Cutucu, most likely from a location in one of Aurania's concession areas. Now Aurania's task is to find the source of the alluvial gold that made Logroño the richest of the Spanish gold cities.

HISTORY AND BACKGROUND

The term "gold cities" is misleading, though in conventional use at the time. These were mining camps similar to palisaded forts. The miners were laborers impressed from the indigenous people in the area. The cities typically lasted until the gold was depleted or the indigenous population died out from smallpox or other Western diseases. In the case of Logroño and Sevilla

Figure 2
Locations of Logroño and Sevilla del Oro on a map made by Diego Mendez in 1574.

del Oro, the native people repeatedly destroyed these cities until the Spanish abandoned them at the beginning of the 17th century when Spain lacked the will and resources to keep them open.

A map in the world's first published atlas (Abraham Ortelius, 1527–1598), labeled Peruviae Auriferae Regionis (Gold Regions of Peru), made by Diego Mendez in 1574 shows the location of these two cities (see Figure 2). Because of the remoteness of the area and the inability to accurately measure longitude[1] in that era, the locations shown on the map are rough estimates rather than precise locations. However, the general features of the map are accurate, which provided evidence of the existence of these cities as well as estimates of their locations.

More than four hundred years have passed since Spanish activity at Logroño ceased, and even though many of the records have been lost, what survives is a compelling narrative of gold mining in what was one of the most remote and isolated areas on Earth.

Dr. Barron first visited Ecuador in 1998 and decided to learn Spanish while staying with Professor of History, Dr. Octavio Latorre Tapia. To help Dr. Barron learn Spanish, they agreed that during the day they would speak Spanish while at night they would speak English. Dr. Latorre was lecturing at the Universidad Internacional in Quito at the time and specializing in the cartography of the age of the Spanish conquistadors in the New World.

Some years before the visit, the Ecuadorian government had engaged Dr. Latorre to perform archival research on lost gold settlements and mines. The premise being that mining in the 1600s was low-tech and inefficient so the Spanish could not have extracted all or even most of the gold from these mines. Interest in the lost Spanish gold cities was piqued after the accidental discovery of the Nambija mine complex in 1981 by two boys hunting in the forest. The mine had been abandoned after a smallpox epidemic killed the labor force around 1603. It was reactivated, and by 2000 it had officially produced some 2.7 million ounces of gold. Unfortunately, thousands of miners streamed into the city in a gold rush. This unregulated mining created an environmental disaster, as shown in Figure 3.

Later research revealed there were abundant references to Nambija in the archival literature, including a map from 1750 which gave the location. The Ecuadorian government believed that Nambija would have been rediscovered much earlier had there been a dedicated effort to search the archives for clues to its location. In addition, by regulating the rediscovery and development of lost gold cities, the Ecuadorian government hoped to avoid the environmental disaster that befell Nambija.

Even though the government stopped funding his research into lost gold cities, Dr. Latorre continued the research privately. He disclosed to Dr. Barron that this research had uncovered the existence of two lost gold settlements, "gold cities", that had still not been relocated: Logroño de los Caballeros and Sevilla del Oro. The first, named after the founder, Juan de Salinas Loyola's home in Rioja, Spain, and the second named after the Royal Seat of Seville. Both settlements were founded around 1560–1568; the last mention of them in the literature was in 1605. By 1630, they had vanished from almost all maps.

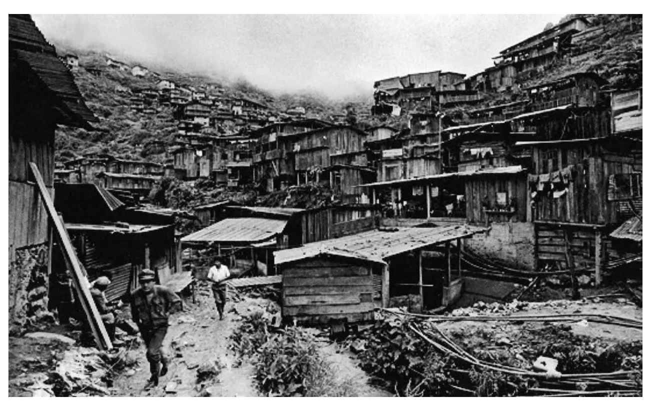

Figure 3
Nambija in 1993. Unregulated mining produced an environmental disaster.

[1] Accurate measurement of longitude requires measuring the time of the maximum height of the sun. This requires portable clocks (chronometers) that can keep accurate time over long periods. Chronometers were not generally available until the 19th century (see [1]). By contrast, the measurement of the latitude, which requires measuring only the Sun's peak altitude in degrees above the horizon, was substantially more accurate.

In 2000, Dr. Barron returned to Ecuador and discussed the possibility of forming a company with the purpose of re-locating the "Lost Cities" using both geological data and historical documentation. Dr Barron had read the book, *The Ship of Gold in the Deep Blue Sea* by Gary Kinder [2], which discussed the novel use of historical data of variable reliability in the application of Bayesian search theory. The book detailed the search for the wreck of the *SS Central America*, a side-wheel steamer that had gone down in a hurricane off the coast of South Carolina in September 1857 while carrying a shipment of gold from the United States mint in San Francisco. The Columbus-America Discovery Group engaged Dr. Stone to use their historical data search matrix to develop a probability map for the location of the wreck. This map was used to guide the search. Ultimately, the search was successful, locating the wreck in September 1988. It is considered the richest shipwreck of all time, and to date $1 billion USD in gold bars and coins have been recovered. Dr. Barron gave Dr. Latorre a copy of Kinder's book, which he quickly devoured; but the verdict was that there were insufficient historical and geographical clues to proceed in a similar way to find the two lost gold cities.

In January 2001, after receiving a tip from Ecuador's Regional Director of Mines, Dr. Barron visited the site of some gold mining in the Province of Zamora-Chinchipe, which borders on Peru. He realized that the artisanal gold miners, who were vacuum dredging in the rivers, were only 4 kilometers from the drainage divide that marked the international border with Peru. As the streams had their origins at the tops of nearby mountains in Ecuador, the source of the gold in these rivers was likely somewhere in those mountains. In April 2001, Dr. Barron began amassing exploration concessions in these mountains for Aurelian, an exploration company founded by Dr. Barron. Aurelian discovered the Fruta del Norte (FDN) gold mine in one of those concessions in March 2006 and was acquired by Kinross Gold Corp for $1.2 billion CAD in 2008.

Later in 2008, Dr. Barron again joined forces with Dr. Latorre to find the Lost Cities. By virtue of the sale of Aurelian, they now had the advantage of well-funded archival research. In addition to examining archives in Ecuador, they travelled further afield to the:

- Archivo Historico Arzobispal
- Riva Agüero Institute, Lima
- Biblioteca Nacional de España, Madrid
- Rare Book Division of the New York Public Library
- British Museum Library, London
- Archive of the Indies (Archivo General de Indias) in Seville, Spain
- Manuscript Section of the Apostolic Library of the Vatican, Rome

Over 100 historic documents relating to Logroño and Sevilla del Oro were discovered in Seville alone.

During a 2011 visit to the Vatican, Drs. Barron and Latorre found an anthology referencing the *Compendium and Description of the West Indies* written in 1628–29 by Carmelite priest Antonio Vazquez de Espinosa, which was eventually unearthed in the Manuscript Section of the Apostolic Library in the Vatican (Barb Lat. 3584). Dr. Barron was able to examine and obtain photographic reproductions of the original volume in the Vatican in May 2016. The book gives a description of how to get to Sevilla del Oro with place names that are still recognizable today and established that the Cordillera de Cutucu, a remote area in southeastern Ecuador some 100 km north of Fruta del Norte, was the site of the two lost settlements.

By coincidence, in March 2016, Dr. Barron had applied for 208,000 hectares of concessions in the Cordillera de Cutucu area. He had been interested in this property for some time, but a long-standing moratorium on granting concessions was in effect. During the Prospectors & Developers Association of Canada Conference in Toronto, the moratorium was lifted. At 12:01 am, while others were at the Ecuadorian Ministry of Mines cocktail reception, Barron and his Vice President for Exploration were typing in the coordinates to apply for the concessions. The concessions, awarded in December 2016, are shown in Figure 4.

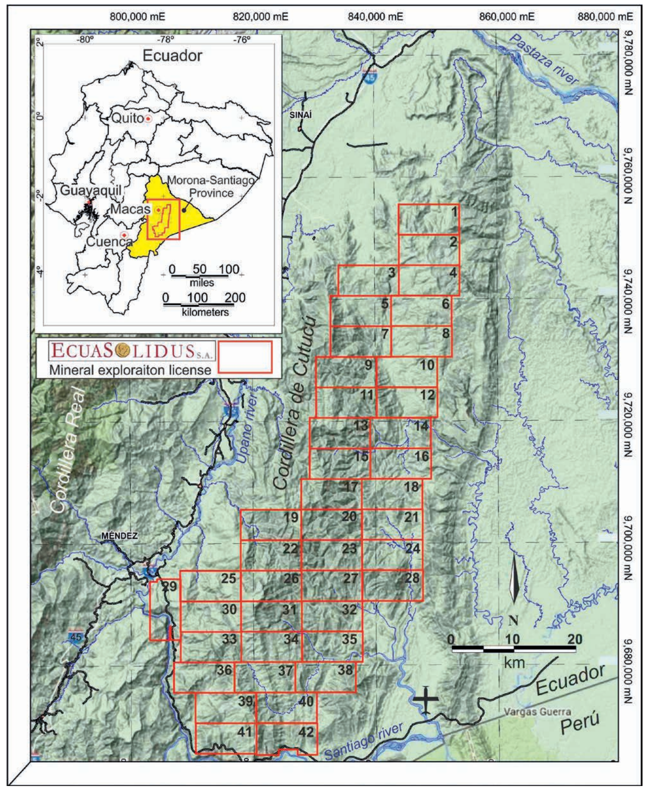

Figure 4
Cutucu project area. Red rectangles show Aurania's mineral concession areas.

PROBABILITY MAPS AND LIKELIHOOD RATIO SURFACES

While Aurania was performing extensive geophysical and geochemical surveys of the Cordillera de Cutucu region to determine the locations of commercially exploitable mineral deposits, Metron was performing the tasks below.[2]

- **Probability Maps.** Using historical information such as maps and descriptions from 16th and 17th centuries indicating the location of the cities, descriptions of journeys to these cities, and present-day topological features, Metron produced probability maps for the locations of Logroño de los Caballeros and Sevilla del Oro.
- **Likelihood Ratio Surfaces.** Using geophysical survey information such as magnetic anomaly measurements and geochemical information provided by stream bed samples, Metron produced likelihood ratio surfaces indicating likely locations of gold, silver, copper, lead, and zinc deposits.

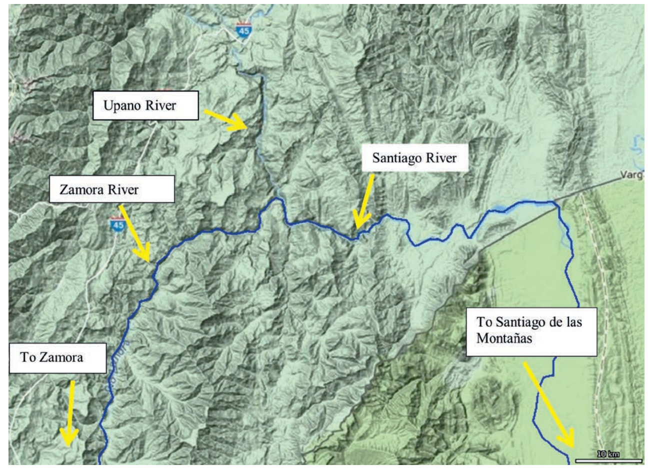

Figure 5
The Zamora River. Where the Upano joins the Zamora, the river becomes the Santiago River.

PROBABILITY MAPS FOR LOGROÑO AND SEVILLA

Metron developed the probability maps for Logroño and Sevilla using classic Bayesian search theory methods in which both objective and subjective information are incorporated into the probability distribution. Uncertainties in the information are represented by probabilities (possibly subjective). In addition, Metron considered multiple scenarios for estimating the location of each city. When information is gathered about the location of a city (or search object on the ocean bottom), it tends to coalesce into disjoint subsets of information. The information within a subset is consistent, but the information in one subset is inconsistent with that in another. As with the successful searches for the US nuclear submarine *Scorpion* lost in 1968 [3], the *SS Central America* that sank in 1857 [4], and the Air France AF 447 flight that crashed into the Atlantic in 2009 [5], Metron treated each subset as defining a scenario for the location of the city. It gave the information in each scenario a (subjective) credibility factor, produced a probability map based on each scenario, and computed a combined map that is the mixture of the scenario maps weighted by their credibility factors.

We describe the procedure for computing the Logroño map and compare the map (which was computed in 2020) to the location of Logroño discovered in 2022. The procedure for Sevilla was similar, but the location of Sevilla has not yet been discovered.

LOGROÑO PROBABILITY MAP

The Logroño probability map is based on two excellent pieces of historical information obtained by Dr Barron in his search of the archives mentioned above.[3]

- Based on the Mendez map (Figure 2), we estimated that Logroño is located at a point halfway between Zamora and Santiago de las Montañas as measured along the Zamora and Santiago rivers.
- "…the city of Logroño is within half a league of the Zamora River" [6]. As seen in Figure 5, the Zamora River becomes the Santiago River where the Upano joins it. The locations of the cities of Zamora and Santiago de las Montañas are off the map as indicated by arrows.

[2] The probability maps and likelihood ratio surfaces were computed by Joshua Hughes who is a Senior Research Scientist at Metron. Metron was aided in this effort by Drs. Camille and Richard Spencer, geochemists employed by Aurania.

[3] Dr. Barron collected a trove of information from the historical records, much of it in the original Spanish, which differs from modern Spanish much like 16th century English differs from modern English. Dr. Camille Spencer selected, curated, and translated crucial pieces of these records to provide the information used to form the Logroño probability map.

BASIC SCENARIO

Items 1 and 2 above form the basic scenario for the location of Logroño. We converted this information into a probability distribution as follows. Let d = distance between Zamora and Santiago de las Montañas as measured along the river Zamora. We estimated this distance as 423 km. Let X = the unknown distance from Zamora to Logroño as measured along the river. For the distribution of X we used a triangular distribution with density defined in (1)

$$\Pr\{X = x\} = p(x) = \begin{cases} 0 \text{ for } x < d/3 \\ \dfrac{36(x - d/3)}{d} \text{ for } d/3 \le x \le d/2 \\ 6 - \dfrac{36(x - d/2)}{d} \text{ for } d/2 < x \le 2d/3 \\ 0 \text{ for } x > 2d/3 \end{cases} \quad (1)$$

where Pr indicates probability density and d = 423 km.

We estimated that a league, as used in the historical documents, is between 4.18 km and 5.57 km (see [7]). So, half a league is roughly between 2 km and 3 km. Let H be the distance of Logroño from the Zamora River. We used the probability density in (2) to represent the uncertainty in this distance.

$$\Pr\{H=h\}=q(h)=\begin{cases}0 \text{ for } h<1.5\,\text{km}\\ \dfrac{1}{2\,\text{km}} \text{ for } 1.5\,\text{km}\le h\le 3.5\,\text{km}\\ 0 \text{ for } h>3.5\,\text{km}.\end{cases} \tag{2}$$

Aurania's geological experts estimated that Logroño was more likely to be located on the north side of the Zamora than the south side. We represented this by assuming:

Probability Logroño North of Zamora = p_N

Probability Logroño South of Zamora = $1-p_N$, where $p_N=0.75$.

SIMULATING THE DISTRIBUTION OF THE LOCATION OF LOGROÑO

We drew 100,000 points from the distribution defined above as follows:

- Make a draw from the distribution in (1) to obtain the distance of the sample point along the Zamora.
- Make a draw to determine if the point is North or South of the Zamora.
- Make a draw from the distribution in (2) to obtain the distance North or South for the point.

This determines a possible location of Logroño. Repeat 100,000 times. The resulting distribution is shown in Figure 6. The distribution consists of 100,000 equal probability points. The graphical representation of this distribution is obtained by imposing a grid on the area near the Santiago River and summing the probabilities of the points contained in each grid cell. The cells are colored according to the probability in them with red indicating the highest probability cells and the colors through orange, yellow, and green to blue indicating lower probability ones.

We had an additional piece of information about the location of Logroño. Specifically, "A short distance from their meeting [of the Upano and Zamora rivers] was …the famous city of Santa Ana de Logroño de los Caballeros, by another name, the city of gold", [8]. We represented this information by the likelihood function described below, which we incorporated into the distribution in Figure 6 to compute the posterior distribution shown in Figure 7. Observe that this distribution is more concentrated and closer to the junction of the Upano and Zamora rivers than the distribution in Figure 6.

LIKELIHOOD FUNCTION

We represented the information that Logroño was located a short distance from the junction of the Zamora and Upano rivers using the gamma probability density function $g_{\alpha\beta}$ given in (3) with $\alpha = 2$ and $\beta = (1/12)$km.

$$g_{\alpha\beta}(x)=\frac{x^{\alpha-1}\beta^{\alpha}e^{-\beta x}}{\Gamma(\alpha)} \text{ for } x\ge 0 \text{ where } \alpha,\ \beta>0 \tag{3}$$

where Γ is the standard gamma function with $\Gamma(n) = (n-1)!$ for n a positive integer. This density function has its maximum at 12 km from the junction as shown in Figure 8. We used this gamma density as a likelihood function and combined it with the prior in Figure 6 in a Bayesian fashion to compute the posterior distribution in Figure 7. Specifically, for each point in the prior distribution in Figure 6, we calculated the distance x of the point from the junction of the Zamora and Upano rivers and multiplied the point's probability by

$$\Pr\{\text{"short distance estimate"} \mid \text{Logroño at } x\} = g_{\alpha\beta}(x)$$

with $\alpha = 2$, $\beta = 1/12$.

We then renormalized the probabilities on the points to sum to 1.

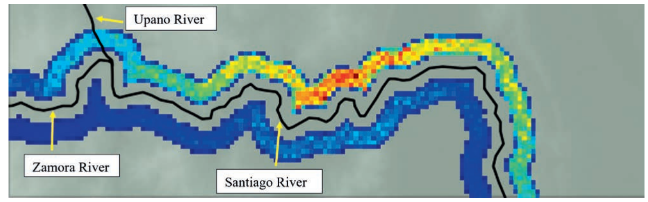

Figure 6
Probability distribution based on items (1) and (2). Red indicates high probability cells, blue low probability.

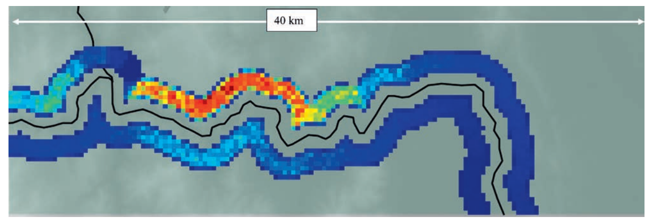

Figure 7
Logroño distribution updated with the gamma likelihood function representing the estimate of the distance of Logroño from the junction of the Zamora and Upano rivers.

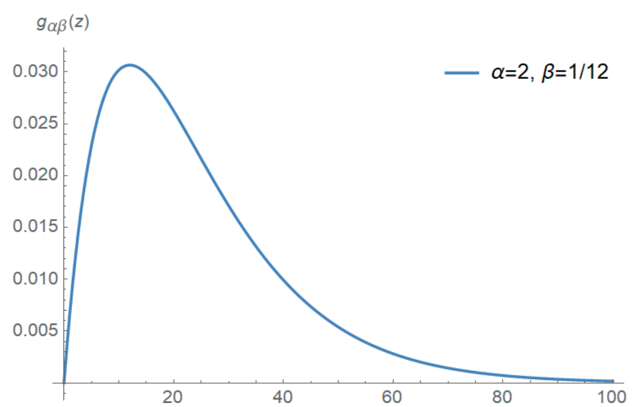

Figure 8
Likelihood function for a "small distance" from the junction of the Zamora and Upano rivers.

LOCATION OF LOGROÑO

Figure 9 shows the location of the recent illegal mining activity along the Santiago River. Undoubtedly, Logroño was located along the banks of the Santiago River where alluvial gold was panned in great quantities from the river by the Spanish.

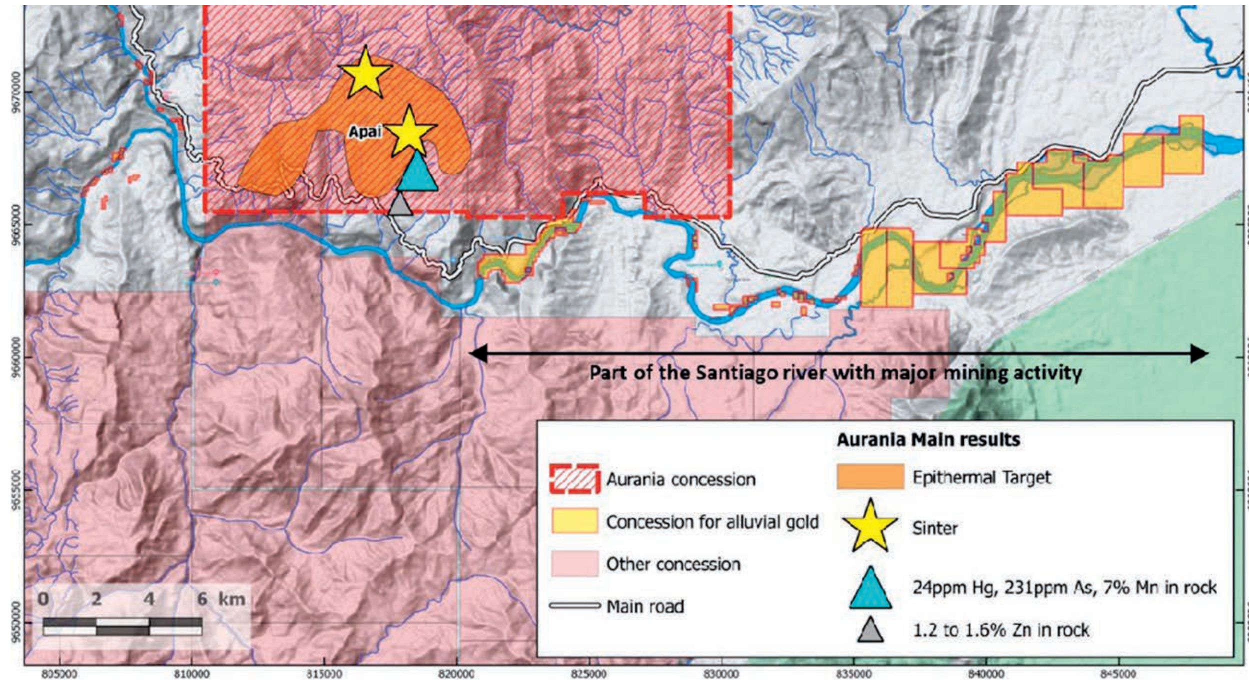

Figure 9
Areas of illegal mining activity are shown.

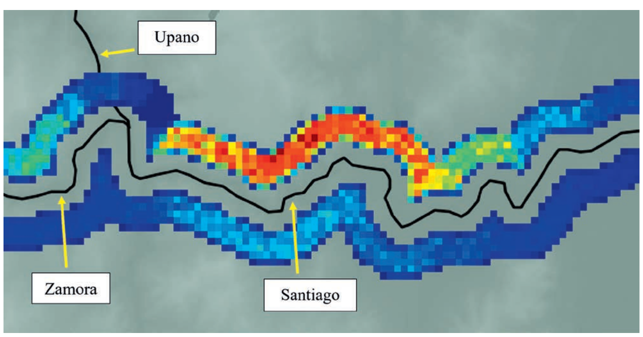

Figure 10
High probability regions for the location of Logroño.

Clearly, this gold has washed down from the mountains above the river. Figure 10 below shows an expanded version of the Logroño map in Figure *7*. The correspondence between the high probability areas on the map and the regions where illegal mining activity was taking place is striking.

Despite this preponderance of evidence, the location of ancient Logroño can never be known with absolute certainty. There will be no stone signposts, and any archeological site along the river would likely have been swept away long ago during high water or destroyed by modern mining activity.

Aurania's task now is to find the source of the gold in the alluvial deposits in the Rio Santiago. Already Aurania has found a site upstream where a landslide has revealed epithermal chalcedonic quartz veins and where gold can be panned. Epithermal deposits are likely to contain gold. There is also an epithermal deposit in the southwest corner of Aurania's concessions that is only a couple of kilometers north of the river. Aurania believes that dedicated and intensive geological mapping and prospecting will ultimately locate the source or sources of the alluvial gold in the Santiago River.

LIKELIHOOD RATIO MAPS

With help of Drs. Camille and Richard Spencer, expert geochemists employed by Aurania, Metron[4] prepared likelihood ratio maps indicating areas of high likelihood for containing gold, silver, copper, lead, and zinc. In this section, we discuss the method used to produce the likelihood ratio map indicating the presence of gold and compare it to the region near Logroño, especially the mountains above it.

For the presence of gold, we prepared likelihood ratio maps indicating the presence of epithermal deposits which are likely to contain gold and silver. We incorporated three types of evidence, lithology (types of rock found), magnetic anomalies found during magnetic surveys of the concession areas, and the minerals found in stream-bed samples in the concessions area.

We imposed a grid on a region that included the concession areas. For each piece of evidence and type of deposit, experts (Camille and Richard Spencer) estimated the following likelihood ratio at the locations x at which measurements were made:

$$LR(x) = \frac{\Pr\{\text{obtain measurements at location } x \mid \text{epithermal deposit present}\}}{\Pr\{\text{obtain measurements at location } x \mid \text{no epithermal deposit present}\}}. \tag{4}$$

The reader may be wondering why we have chosen to encode the subjective estimates of experts in terms of likelihood

[4] The likelihood ratio maps were prepared by Joshua Hughes, Senior Research Scientist at Metron.

Table 1

Likelihood Ratios (LR) for Indicator Scores for Presence of an Epithermal Layer

	Score							
	10	6	5	4	3	2	1	0
LR	9.0	6.0	4.5	3.0	2.25	1.8	1.5	0.56

ratios. The reason is very simple and very powerful. Likelihood ratios give the proper weight to each piece of evidence in the likelihood ratio surface. We can then combine likelihood ratios from diverse types of information by simply multiplying, for each cell x, the likelihood ratios for the evidence in that cell. If there is no measurement in a cell for a given type of evidence, the likelihood ratio is set to 1 in that cell. The resulting likelihood ratio surface represents a combination of the evidence wherein each piece of evidence is given its appropriate weight as determined by expert opinion. We illustrate this process for epithermal deposits using the results of stream sediment data.

Aurania experts defined the catchment area for each stream sampled. The catchment area of a stream is the area that drains into the stream. They identified eight key chemical elements as pathfinders (indicators) of epithermal deposits of gold and silver. They are: gold, silver, tellurium, arsenic, antimony, mercury, selenium, and thallium. For each element, they set a threshold. If the stream sediment analyses corresponding to a catchment area contained the element at a level above this threshold, it received a score of 1 point for all the elements except gold which received a score of 3. The scores for a catchment area were added together to obtain a total score for the area. The maximum score that a catchment area can receive is 10. Aurania experts assigned likelihood ratios to each catchment area according to its score as shown in Table 1.

Figure 11 shows log likelihood ratio maps for the presence of an epithermal layer based on geochemical evidence from lithology, stream sediment samples, and magnetic anomalies. Figure 12 shows the combined likelihood ratio surface obtained by multiplying pointwise the likelihood values in the maps in Figure 11. Figure 13 zooms in on the section of Figure 12 near the junction of the Upano and Zamora rivers which is just upstream from the area where the illegal gold mining took place on the Santiago River. The mountains just north of the Santiago River and east of the Upano are the likely source of the alluvial gold found in the Santiago River. Observe that these mountains are marked as high likelihood regions for epithermal deposits in the map in Figure 13. Notice also that one of Aurania's concession areas is located here.

CONCLUSIONS

The accuracy of the probability map generated for the location of Logroño is another striking example of the power of Bayesian search methods. These methods allow the analyst to combine all available information, both objective and subjec-

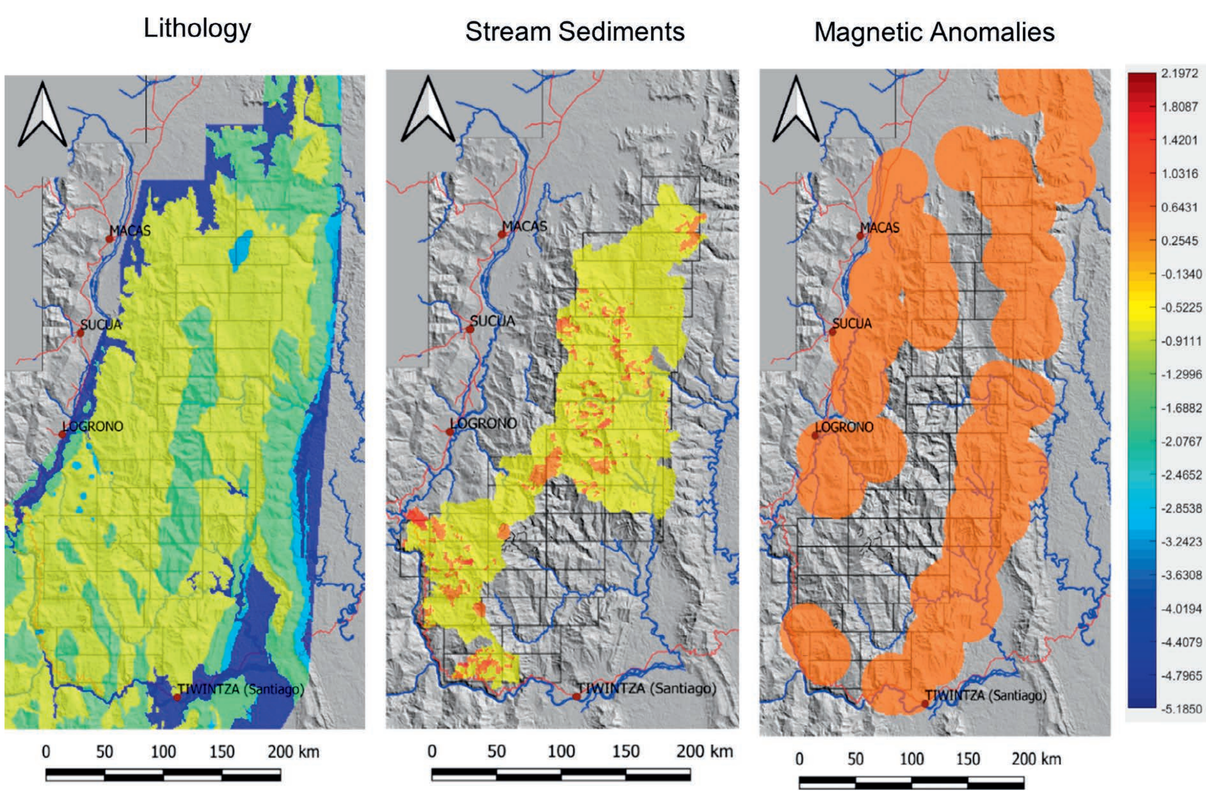

Figure 11
Log likelihood ratio maps for lithology, steam sediments, and magnetic anomalies.

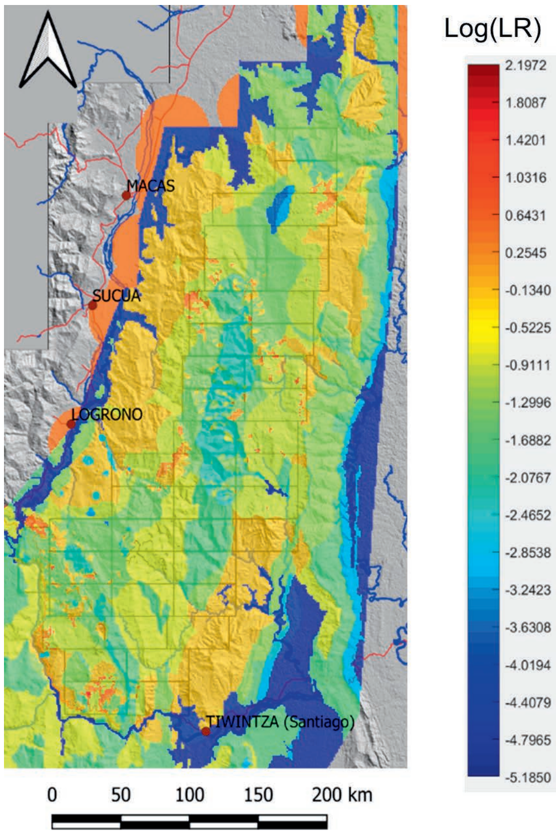

Figure 12
Combined log likelihood ratio map for epithermal deposits.

Figure 13
Log likelihood ratio map near the junction of the Upano and Zamora Rivers.

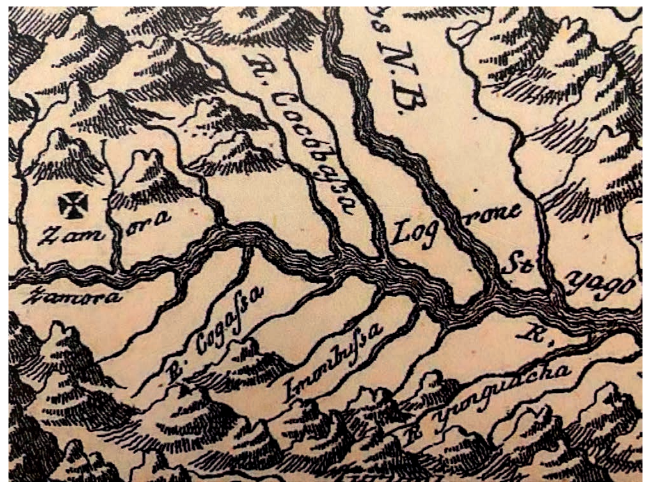

Figure 14
Maragnonii sive Amazonum Fluminis Terrarum in Orbe Maximi.

tive, in a principled Bayesian fashion, to produce a probability map for the location of the object of interest. In the Bayesian methodology, uncertainties in the information, both the objective and subjective, are represented in terms of probabilities. By constructing prior distributions using the scenario method and incorporating information using likelihood functions via Bayes rule, we can produce a probability distribution that represents the client's best understanding of the problem.

Creating a probability map for a lost or missing object is not a scientific endeavor. We do not have the luxury of repeating the experiment a thousand times to test our models and assumptions to determine which ones are correct. Instead, our goal is to produce a probability map that enables the client to search rationally and effectively. To do this, we use Bayesian methods which allow us to incorporate expert opinion and subjective judgements along with objective information. In many searches, this approach has proven to be effective and efficient. In fact, trying to take a "scientific" approach that uses only "hard" data and does not incorporate subjective information, such as expert opinion, can produce very ineffective searches.

The likelihood ratio maps produced for Aurania provide another example of the power of Bayesian methods to combine disparate types of information in a simple and effective manner. In the case of mining exploration, surveys performed by an exploration company can produce many types of evidence for the presence of economically significant mineral deposits. As in Aurania's case, that evidence can be geological, geochemical, and geophysical. Each type of information produces indications for likely location of mineral deposits. Using expert opinion to convert that information into likelihood ratios, puts the information on a common scale that reflects the strength of the evidence so that the evidence can be combined in a principled manner. We are not aware of another method that does this.

POSTSCRIPT

In August of 2022, Dr. Barron visited the Iglesia de la Compañia, in Quito, built by the Jesuit Order between 1605 and 1765. The altar, ceiling, and internal ornamentation are entirely covered with gold leaf. Over time, this church has also been called: "the Temple of Solomon of South America". Father Bernardo Recio, a traveling Jesuit, called it the "Golden Ember". On display during Dr. Barron's visit, and never examined before, was a very rare map "Maragnonii sive Amazonum Fluminis Terrarum in Orbe Maximi" (see Figure 14) published in Nuremberg in 1785 by the Jesuits, which documented the various explorations and settlements founded by them before their expulsion from the Spanish colonies in 1767. In this map, "Logrone" is placed near the beginning of the "St Yago R.", only some 20 kilometers from where Logroño was ultimately found. It is tantalizing to think that some of the gold used in the construction of this church came from the Lost Cities.

REFERENCES

1. Wikipedia. Longitude, https://en.wikipedia.org/wiki/Longitude, last access May 11, 2023.
2. Kinder, G. Ship of gold in the deep blue sea. New York, NY: Atlantic Monthly Press, 1998.
3. Richardson, H. R., and Stone, L. D. Operations analysis during the underwater search for *Scorpion. Naval Research Logistics*, Vol. 18, 2 (1971), 141–157.
4. Stone, L. D. Search for the *SS Central America*: Mathematical treasure hunting. *Interfaces*, Vol. 22, 1 (1992), 32–54.
5. Stone, L. D., Keller, C.M., and Strumpfer, J. P. Search for the wreckage of Air France AF 447. *Statistical Science*, Vol. 29, 1 (2014), 69–80.
6. Lopez, J. d. A. Expediente de Juan de Escobar, escribano de las minas de oro de Zaruma, sobre confirmación de su título. In *QUITO 24, N. 8*. Sevilla, Spain: Archive of the Indies, 1593.
7. Wikipedia. League (unit), https://en.wikipedia.org/wiki/League_(unit), last access May 11, 2023.
8. Maroni, P. 1738 Noticias auténticas del famosa río Marañón. In *Seguidas de las Relaciones de loa P.P. A. de Zárate y J. Magnin. (1735-1740). Monumenta Amazónica*. Lima: Institut Français d'Études Andines, 1988.

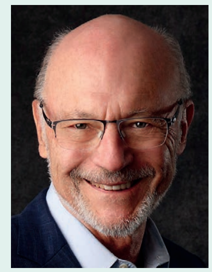

Lawrence Stone is Chief Scientist at Metron Inc. He is a member of the National Academy of Engineering and a fellow of the Institute for Operations Research and Management Science (INFORMS). He is a recipient the Jacinto Steinhardt Award from the Military Applications Section of INFORMS in recognition of his applications of Operations Research to military problems.

In 1975, the Operations Research Society of America awarded the Lanchester Prize to his text *Theory of Optimal Search*. In 1986, he produced the probability maps used to locate the *S.S. Central America* which sank in 1857, taking millions of dollars of gold coins and bars to the ocean bottom one and one-half miles below. In 2010, he led the team that produced the probability distribution that guided the French to the location of the underwater wreckage of Air France Flight AF447. Recently, he used search theory methods to help guide the Canadian exploration company, Aurania, to the location of one of the lost Spanish gold cities in Ecuador.

He coauthored the 2016 book *Optimal Search for Moving Targets*. He was one of the primary developers of the Search and Rescue Optimal Planning System (SAROPS) used by the U.S. Coast Guard since 2007 to plan searches for people missing at sea. He continues to work on a number of detection and tracking systems for the U.S. Navy. He is a coauthor of the 2014 book *Bayesian Multiple Target Tracking, Second Edition*.

Keith Barron is Chairman, President, and CEO of Ecuador gold and copper explorer, Aurania Resources Ltd. Dr. Barron has been involved as an Exploration Economic Geologist and entrepreneur for over 39 years, in more than nineteen countries. In 2001, he privately co-founded Ecuador gold explorer Aurelian Resources Inc., which was listed on the TSX-V in 2003 and made the colossal Fruta del Norte gold discovery in 2006. The company was bought by Kinross Gold in 2008 for $1.2 billion CAD. He is President and Director of Firestone Ventures (TSXV:FV) and President of Potentate Mining in Montana, USA, the world's second largest sapphire miner. He has served as a Director of several other listed companies.

At the PDAC convention in March 2008 he was awarded the Thayer Lindsley International Discovery Award for his role in the discovery of the Fruta del Norte gold deposit, and he was also jointly named the Northern Miner's Mining Man of the Year 2008. His major expertise is in gold and precious gems, although he has considerable experience in base metals, uranium and industrial minerals. His doctoral thesis documented the world's oldest epithermal gold-silver deposit at Springpole Lake in Canada. He holds a Ph.D. in Geology from the University of Western Ontario and a BSc. (Hons) in Geology from the University of Toronto.

6th School on Belief Functions and their Applications

https://bfasociety.org/BFTA2023

Japan Advanced Institute of Science and Technology, Ishikawa, Japan
Oct. 27- Nov. 1, 2023

Important Dates

1st May 2023: Registration opens

31st July 2023: Deadline for early-bird registration

15th September 2023: Deadline for registration

Students 190€ (early-bird) / 240 €

Regular 290€ (early-bird) / 340 €

Places are limited

The School

The BELIEF school is a biennial event organized by the Belief Functions and Applications Society (BFAS) that offers a unique opportunity for students and researchers to learn about fundamental and advanced aspects of the theory of belief functions (also referred to as Dempster-Shafer theory, or evidence theory), a formalism for reasoning with uncertainty.

The school will be organized around a set of lectures by prominent researchers. Lectures will gradually tackle basic to more advanced theoretical concepts. They will also highlight the links with other uncertainty theories such as random sets and possibility theory, and present applications of belief functions in various domains including machine learning, information fusion, statistical inference and materials science.

The Lectures

Introduction to the theory of belief functions
Thierry Denoeux, Université de Technologie de Compiègne, France

Evidential clustering
Thierry Denoeux, Université de Technologie de Compiègne, France

Differential privacy for belief functions
Chunlai Zhou, Renmin University of China

Pattern classification with belief functions
Zhunga Liu, Northwetsern Polytechnical Universityn Xi'an, China

Information fusion in the theory of evidence
Frédéric Pichon, Université d'Artois, France

Application of possibility theory to optimization and decision making
Masahiro Inuiguchi, Osaka University, Japan

Epistemic random fuzzy sets: theory and application to machine learning
Thierry Denoeux, Université de Technologie de Compiègne, France

Applications of Belief Functions for Exploring Novel Materials
Hieu Chi Dam, JAIST, Japan

Old and new developments in (consonant) belief functions for statistical inference
Ryan Martin, North Carolina State University, USA

Graphical Belief Function Models: Theory, Computation and Applications
Prakash Shenoy, School of Business, University of Kansas, USA

Measuring inconsistency in evidence theory
Anne-Laure Jousselme, CS Group Research Lab, France

Application of belief functions to ensemble classification and recommendation
Van-Nam Huynh, JAIST, Japan

Note: Financial support will be offered by BFAS to some students for attending the school. Details of the application procedure will be announced soon.

Organizing committee

Van-Nam Huynh (Chair), JAIST, Japan

Thierry Denoeux, Université de technologie de Compiègne, France

Anne-Laure Jousselme, CS Group Research Lab, France

Fréderic Pichon, Université d'Artois, France

Contact information:

Email: bfta2023@bfasociety.org

bfasociety.org/BFTA2023

BFAS

JAIST
JAPAN ADVANCED INSTITUTE OF SCIENCE AND TECHNOLOGY 1990

Fusion Processes and Situation Control

Abstract—The ISIF community has a long heritage in research directed to situation assessment. That research traces to the earliest days of the International Society of Information Fusion (ISIF) to include the first few conferences. Well thought out ideas and models of situation analysis/situation assessment/situation projection processes were developed across the international community. Companion efforts grew out of the Cognitive Situation Management (CogSIMA) community in the context of cognitive situation control that are complementary to the ISIF papers. However, continued maturation and integration of those ideas toward designing and developing prototype integrated fusion processes have not been realized. This paper offers some additional ideas that we call an expanded framework for situation control that will require such integrated and managed processes. At its heart, the paper is a call for the ISIF community to move away from functionally isolated research, and to develop a more systemic view of its research that will offer opportunities for more impactful roles in the research community.

James Llinas
University at Buffalo
Buffalo, NY, USA
llinas@buffalo.edu

PERSPECTIVES ON SITUATION ASSESSMENT

There is an extensive body of literature on the topic of estimating situational states in applications ranging from cyber-defense to military operations to traffic situations and autonomous cars. In the military/defense/intelligence literature, "situation assessment" seems to be the *sine qua non* for any research on surveillance and reconnaissance, command and control, and intelligence analysis. Virtually all of this work focuses on assessing the situation-at-the-moment; many if not most of the estimation techniques are based on data and information fusion (DIF) approaches, with some recent schemes employing artificial intelligence (AI) and machine learning (ML) methods. But estimating and recognizing situational conditions are processes often couched in a decision-making, action-taking context, implying that actions may be needed so that certain goal situations will be reached as a result of such actions, or at least that progress toward such goal states will be made; that is, situations are generally not being estimated just to be observed. This context thus frames the estimation of situational states in the larger context of a control-loop, with a need to understand the temporal evolution of situational states, not just a snapshot at a given time. Estimating situational dynamics requires the important functions of *situation detection*, *situation recognition*, *situation understanding*, *situation prediction, and situation comparison* that are also central to such an integrated estimation + action-taking control process architecture. The varied processes for all these combined capabilities lie in a closed-loop "situation control" framework, where the core operations of a stochastic control process move the situation to a desired goal state; see an earlier paper on this topic [1] and a longer version of this paper in [2].

SYSTEMIC VIEWPOINTS

The issues described above are DIF system-boundary issues in the systems-engineering sense. Much DIF research is couched in the sense of DIF as a value-adding but isolated process: e.g., how many papers on tracking, the most-studied function in the community, address the details of and synergies with multisensor operations or other system-level functions? Proportionally, very few, and virtually all rely on the sensing system somehow producing observations that satisfy the Mutual Exclusion criterion[1], among other assumptions about observational data. This paper suggests that more systemically expansive research is needed in the DIF community and is a paper that looks at some of the interdependencies among DIF situation-estimating processes and decision-making. Minimally, the DIF "Black Box" should be extended to synergies with Sensing and with Response Systems, and with Humans. In adversarial and many civilian situations, the core purpose of DIF will be to deliver information needed for optimal action-taking of some kind; fusion for disaster-response is a good example [3], involving situation assessment to enable coupled life-saving operations as a major purpose.

We propose several additional functionalities for this closed-loop control process as an expansion of some prior work on the situation-control topic and include remarks on the integration of some control-theoretic principles. Some remarks are also made on the state of the art of the schemas and computational technologies for situation detection, recognition, prediction, and understanding, as well as the roles for human intelligence in this larger framework. Our intent in this paper is to expand the framework of situation control in terms of our views of several other component processes briefly described herein, and in discussing these additional processes, to relate them to research and capabilities in the DIF domain.

INTRODUCTION

The concept of a "situation" can be thought of as describing a portion of a real-world that is of interest to a participant in

[1] One measurement per single target per sensor.

that portion of the world; see "Situation Estimation" for some details. An understanding of a situation is needed and useful toward guiding or assessing the need for possible assessment and action of the participant in that situation. Action of a participant may also be needed to possibly alter the situation if it is in an undesirable state (assuming resources capable of affecting the situation *in known ways* are available, and that a *goal state* can be specified), or for the participant to alter his position in the situation. For a human participant, the mental faculties of human cognition, such as consciousness (awareness), reasoning, formation of beliefs, memory, adaptation, and learning, frame the functional aspects of a process of *cognitive situational understanding*, related to the notion of "sensemaking" (see, e.g., [4], [5], [6])[2]. Acting on the situation, however, leads to the need for a process of *cognitive situation control (or management)*, as well described in various of Jakobson's papers [7]–[10] that, in part, motivated this work and provided its foundation. We build on and recognize Jakobson's work especially in [10]. We also recognize and draw on Roy's Fusion 2001 paper on Situation Analysis that also brought forth many of the ideas discussed herein [11]. Similar ideas were also described in Lambert's 2001 paper as well [12]. In our *Frontiers* publication [2], we offer an expanded view of the issues discussed here, including aspects of cognitive/neural situational understanding.

SITUATION ESTIMATION

Our abstraction of the notion of a situation is as "a set of entities in a set of relations". If this characterization is acceptable, then situation estimation (SE) involves inferencing about the *existence of relations* across entity sets. Philosophers have generally agreed that "relation-making characteristics" derive from *certain types* of "monadic" properties of entities, e.g., the heights of people form the *basis* of possible relations ("taller-than", etc.). In this view, such properties *enable* inferencing about the existence of relations. These shared properties that enable the existence of relations are called the "relata" (of relations), or "relative-making characteristics" [13]. This line of thought also suggests *that relations are the result of a process* of some type of comparison, i.e., [14], "an act of reasoning". Further, sensors and associated processing (feature/attribute extraction) provide "relata" or *entity properties* that would *support* reasoning from which inter-entity relations could be asserted, but sensors *do not* provide "observations" of relations; those need to be inferred from the relata, as just stated. Importantly, situation estimation is also complicated by the *combinatorics of relations* among entities and entity-sets in complex real-world cases; sets of relata and sets of entities impute these inherent combinatorics. We have not seen much continued research focused along the lines of these remarks, yet Roy [11] pointed out as far back as in Fusion 2001 the need to develop estimates of *sets of relations* among entities in a process he called "Situation element contextual analysis", as part of his situation analysis model. That contextual analysis "… thus develops a description of *all sorts of relationships* among situation elements: physical (is composed of), spatial, proximity, temporal, structural, organizational, perceptual, functional (involves/requires/provides), functional (e.g., supply, communications), process (performs the process of), causal, informational source/recipient, influence source/recipient, sequential dependency (occurs conditional upon), temporal dependency (occurs when), etc.", from [11]. Research directed to the complex machinery needed to estimate the component relations and relation-sets and their integration remains fairly absent in the fusion community at large.

SEMANTIC LABELING, MODELING, AND ONTOLOGIES

The entities at higher levels of DIF processing are not just the physical objects but can be actions, events, behaviors, and other things that may be of concern. Specifying the appropriate Level 1 entities requires looking ahead to Levels 2 and 3 because these inferencing/estimation processes are *interdependent*. In turn, these entities can have combinatoric sets of relations to each other, as just mentioned, but now across fusion Levels. Some type of semantic labeling of the entities and their relational constructs must be established to have a "language" with which to discuss and label DIF-produced estimates of situational conditions.[3] There are various ways to address this language requirement: examples are the use of a situation modeling language, e.g., [15], or the use of an ontology, e.g., [16], [17]. The situation modeling approach typically employs a graphical language for situation modeling (such as Frames), allowing the expression of primitive situations and complex situations involving the composition of situations (with temporal or other constraints when required). In an ontological approach, along with the entity ontology, a relation ontology is also needed so that the specifics of a labeled, specific situational state can be assembled from these components. That assembly requires a higher level of abstraction in inferencing. Thus, a situation detection or recognition process will need to be supported by an ontological foundation where entities, relations, and *labeled situational states* are coupled to the fusion and recognition processes that will have to assemble the recognized, labeled situational state by exploiting this framework and all of the relata. Steinberg [18] offers one example of the inferencing machinery for these operations, building on the situation logic processes of Barwise and Perry [19]. These processes also need to account for the various uncertainties in the integrated observational and inferential processes. Jousselme et al. [20] provide an overview of the principal typologies of uncertainty employed in situation analysis and inferencing and suggest that addressing reasoning and uncertainty in situation assessment will require frameworks having capabilities to integrate qualitative and quantitative pro-

[2] Sensemaking is not the same as understanding; sensemaking involves interplay between foraging for information and abstracting the information into a representation called a schema that will facilitate a decision or solution (http://www.peterpirolli.com/Professional/Blog__Making_Sense/Entries/2010/8/16_What_is_sensemaking.html).

[3] The Level 1 state estimates have a relatively simple set of semantic labels drawn from common language and not needing formalisms of ontology, or at least less so. Those labels provide enough semantic specificity from which to engineer solutions to required processes; e.g., what a "track" is and how to form an estimate of it–this is not the case for Level 2 and 3 processes.

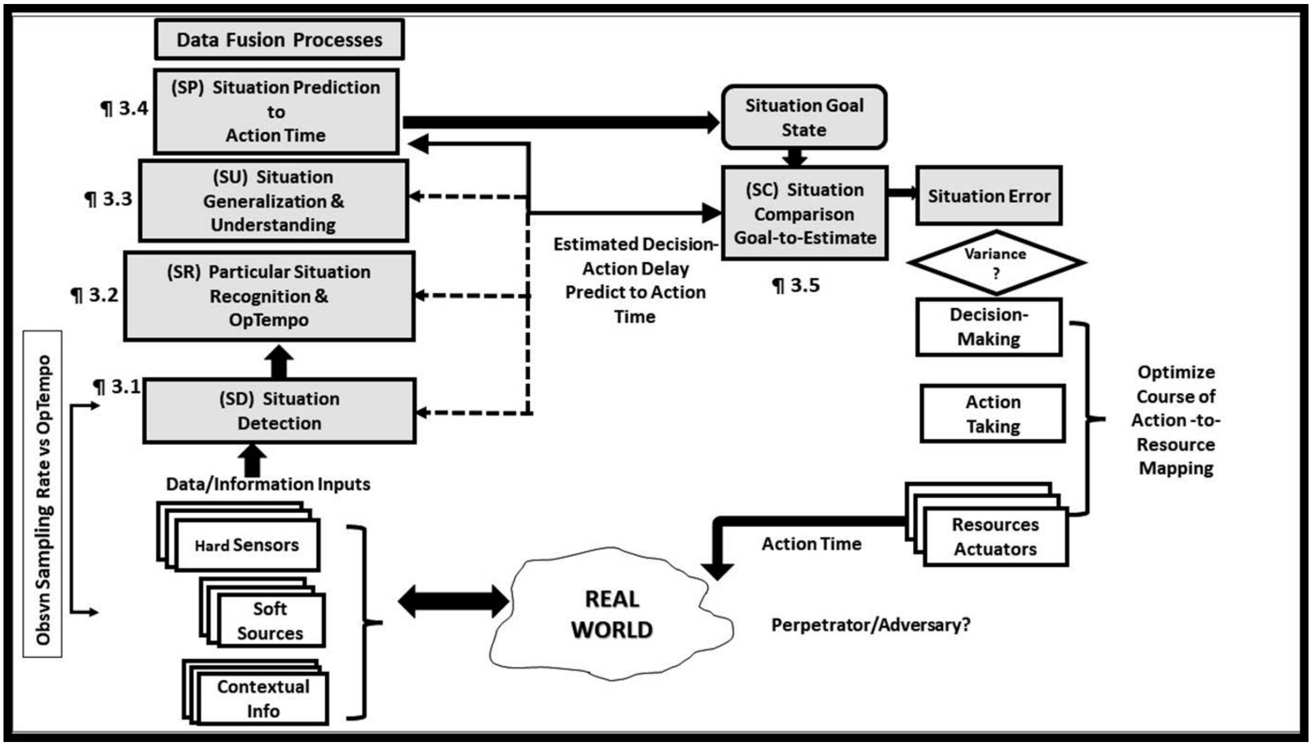

Figure 1
Fusion processes in situation control.

cesses. The modeling and relational specificity described here would also be needed if labeled training data were to be used in an AI/ML-based approach to situation estimation. There has been a fairly large number of publications that offer representational schemes for situations, some labeled as ontologically-based, but those models have not been broadly applied (see [16], [21]–[26], that are just a sampling).

PROPOSED FUNCTIONAL EXPANSION OF THE BASELINE FRAMEWORK: OVERVIEW

While Jakobson provides a sound initial foundation for a process description of situation control, we suggest various enhancements of this process description; our discussion is linked to Figure 1 below. The paragraph titles below are shown in the figure for clarity in following the discussion. Fusion-based situational processing begins with the flow of inputs from hard and soft sensors or sources and from incorporation of contextual information, in the lower-left of Figure 1.

SITUATION DETECTION (SD)

Our view of detection relates to prior knowledge and the aspect of observing the occurrence of something or some part of something that is known *a priori* (elsewise the process is discovery); integrated detection and discovery processes may very likely be needed in some applications, as bounded by prior knowledge but we do not address that issue. For a state such as a situation, an entity-relation complex, detection may relate to observing or measuring some parts of the relational complex. Thus, to detect a situation requires deciding if any entity or relation of a situational complex is of equal value for asserting the detection of the entirety of the situational state. On the other hand, if a situational complex is of large dimensions, detecting small components might lead to many false alarms. Construction of a detection methodology therefore requires setting thresholds and labels of a) numbers of things that need to be detected, and b) which of those things are most indicative of an evolving situational state. Notions of detection probability and probability of false alarm are present here in the same way as for a hard or soft sensor. There are various papers in the literature on this topic, e.g., using Bayes Nets and Fuzzy Logic for Situation Detection but it is rarely presented in the context of an integrated system approach or as a detection theory for situations.

SITUATION RECOGNITION (SR)

Given that a situation has been detected, situation estimation processes can begin toward estimation of the existence of some *Particular Situations*. In some works, these processes are asserted to be "the" Situation Estimation process. But we choose to call this the Particular Situation, arguing that it is just that, the collection of situational elements at any given moment, for a dynamic evolution process that has some time to completion. In the same way as for SD, methods for SR have to decide on the issue of "completeness" and set notions of thresholds toward asserting the existence of any particular situation. An issue that arises here is the degree to which particular situations in the situation ontology are similar to each other. If the "truth" situations specified in the ontology are not sufficiently disparate, labeling of particular situations will be yet more difficult because of overlapping similarities.

ESTIMATING SITUATIONAL RATE: "OPTEMPO"

We introduce a new requirement for DIF that we have not seen in the literature: the estimation of a factor that will be very important in determining the process context for Situation Management and Control: the assessed rate at which the situation is unfolding; that is, the *Operational Tempo* ("OpTempo") of the situation. This factor needs to be estimated early in the SR process and weighed in relation to both the scanning/sampling rate of the observational resources, the prediction interval, sensor resolution factors, and in fact the viability of the overall DIF process (again indicating the need for systems-level thinking). If the situation is unfolding at a rate faster than it can be feasibly observed (or perhaps acted upon), forming dependable situation estimates going forward will be very difficult, and situational predictions will be equally hard.[4] OpTempo can be roughly thought of as related to the hard sensor Nyquist-type sampling rate to capture sufficient information for estimation. This balance changes the dependencies of the Learning/Understanding process (see below) between *a priori* knowledge and real-time observational data; uncertainties in the consequent estimated situation will also be affected. Estimating situational OpTempo should therefore be a fundamental requirement of the SR function, as it is a critical process design and management parameter, setting the overall "clock" for this control process. The notion of OpTempo is also in the fashion of a meta-metric, since any situation will be comprised of multiple component multi-entity relational processes unfolding at varying rates (the combinatorics mentioned previously). Note too that there are optimization issues lurking here, as regards defining how optimal co-employment of bounded observational resources (OR) and situation processing (SP) will be managed across these process needs. That is, there will be competition between the use of OR and SP computational resources for situational state development and for co-estimating its evolution-rate/OpTempo that will for example require temporal comparison processes to be developed.

NATURAL AND ADVERSARIAL ENVIRONMENTS

In any setting involving situation state estimation, an early question has to do with whether the setting is a natural one where phenomena are driven by natural causes or whether the setting comprises a two-sided, adversarial context. The case involving adversaries can be related to the case of "Information Warfare", where the two sides are manipulating information, the bases for perception and inference, to their advantage. The larger purpose of these operations is to manage adversarial perceptions by structuring the information available to an adversary to be compliant with intended perceptual constructs. Another topic related to deception is denial of information by covertness, camouflage, jamming, and other means. Deception and denial strategies work because of exploitation of reasoning errors, cognitive limitations, and cognitive biases [27]. It can be argued then that another early function for SE is to assess and filter out any adversarially related data or states and make an early assessment of the quality and reliability of the data ("garbage-in/garbage-out"). For both natural and adversarial cases, situational models will need to be posed as bases for framing all situational estimates. Thus, as can be said for all DIF processes, process design will require making choices on issues of Data Quality; this is also a factor not seen very much in the SA literature; see [28].

SITUATIONAL UNDERSTANDING (SU)

While the particular situation estimate may be helpful to certain analytical or even decision-making purposes, in many applications, it is desirable or possibly necessary to know or estimate the *class or type of situation* the particular one is an instance of. One notion of understanding can be said to relate to an ability to "generalize from the particulars". Generalization allows the recognition of the similarities in knowledge acquired in one circumstance, allowing for transfer of knowledge onto new situations. A challenge now receiving considerable attention with the new thrusts into AI is to understand how humans are able to generalize from very limited sampling, as well as the issue of "transfer learning". In defense contexts, this type of generalization is often directed to gaining or asserting a "mission" context for the particular situation (the mission class that this situation is an instance of). For example, surveillance is a mission class, comprising phases such as ingress, tactics such as evading, and actions such as attacking; a type of taxonomy of mission-to-situations could help in the generalizations proposed here. Generalizing then allows estimation of a broader type and can also trigger layered estimates (e.g., particular-to-mission-to-tactics-to-strategic). Such broader, generalized views require application of prior knowledge, tacit knowledge, and contextual influences. Ideas along these lines are also seen in [11], where he asserts a need for a "Situation element interpretation" process that similarly focuses on forming a higher, generalized view of the fused results. Generalization can be done by exploratory excursions from the particular situation at the moment as a kind of extended induction, and also by methods drawn from argumentation. Similar techniques are employed in Sensemaking models where "Foraging" is a process that searches for related data and for plausible extensions to the current data set, related to "inductive generalization". Following [29], methods of elaboration and reframing are frequently employed by humans when people are confronted with, or discover, new information from developing situations. Other methods that may offer ways to generalize could come from Bayesian network-based probabilistic generative frameworks that, for example, employ Allen's interval relation network to represent local temporal dependencies in a generative way. These probabilistic generative methods may offer some possible approaches toward "generalizing from the particulars". Probabilistic generative methods have been successfully employed in data fusion-based classification and may offer methods extendable to Level 2 situational understanding. Generalization is also a rather pervasive topic in psychology. In [30], Austerwell et al. discuss the issue of learning how to generalize, which suggests that generalization requires postulating "overhypotheses" or constraints in effect on

[4] This same concern certainly applies at Level 1 (L1) fusion and again is often not an issue coupled to L1 tracking and classification operations because, in much of the community research, they are not couched in the systemic sense as influencing or controlling sensing operations. In the same way as for the Mutual Exclusion issue, "adequate" sensor sampling rates are typically assumed.

the hypothesis domain to be nominated. Some assert that such overhypotheses are innate but Austerwell et al. argue that they can be learned. In either case, the generalization framework is said to be Bayesian-based. Generalization has also been studied in [31] that suggests an exponential metric distance between the stimuli as a basis to assert similarity, and in [32] that discusses the overhypotheses issue. If *a priori* models of general/mission-level situations of interest can be formulated, then notions of "degree" to which the current, particular situation matches that model can be estimated. As situational elements are of natural graph form (nodes as entities, arcs as relation-models), graph-based methods can be applied toward assessing the degree to which the current situation is close to a generalized class of situation. Gross et al. [33] explore such an application and study various ways to make probabilistic-type comparisons between such graph structures.

SITUATION PREDICTION (SP)

The main requirement for a DIF-based situation prediction (SP) process is driven by another system boundary issue, in cases where the DIF SE processes are linked to action-taking and associated decision-making operations; see [1] for an early paper on this topic. This interdependency is driven by the need for synchronizing action-taking such that the action is being taken on the best-estimated situation *at the time of the action*. If that synchronization is not achieved, the acting resources will be acting on an incorrect situation and/or the incorrect entities. This issue also relates to the OpTempo issue; if the situation change rate is slow, some degree of mismatch in SP-action-taking synchronization may be tolerable, and also errors in SP are more tolerable. The opposite is true if the OpTempo is high. Further, as for most prediction, projection, or extrapolation processes, the difficulty and accuracy of such processes is linked to the temporal degree of projection (how far ahead) and the rate of observation and input of any data that the projections depend on; this is not just sensor/observational data but contextual and soft data as well. Some resources that act on situations may be more or less time-sensitive, and this also changes the SP requirement. Thus, synchronization across several interdependent processes may be of concern in this context; a mission goal-based analysis of these dynamics is needed to guide overall processing.

The degree to which an SP needs to predict ahead is related to the expected delay in the combined time it takes to a) decide to act and b) the action-time of any actionable resources. Presuming decision-making precedes action-taking, these projection requirements can also depend on the *type of decision-making style* being employed (see also [1]). That is, it is well-known that there are many variants of decision-making processes that humans and machines may make (see [34]), and so this projection-time estimate may also need to know the decision-making modality being employed.

SITUATION GOAL AND SITUATION COMPARISON (SC)

At some point in time or as part of an ongoing process, an assessment of whether the situation is satisfactory or not is typically carried out; this requires a specification of some desired or goal situational state that is the basis for comparison to real-time estimates. We note that the existence of a goal state is crucial to the overall process, and the placement of comparative operations. It is possible that Goal-to-Estimated SCs could be done quite early in these operations, such as at the moment of Situation Detection or Situation Recognition. Such comparisons, no matter where they occur, are the triggering process for decision-making if the situation is not somehow acceptable. But executing this step thus requires a process for SC. Goals may also change over situation development time, and thus multiple comparisons may be required, in a somewhat ongoing process. However executed, the SC process yields what could be called an "error signal" as would exist in any control process, as Jakobson [10] also points out. We assert that this error signal will have stochastic properties, since the estimated situational state, and perhaps the goal state as well, will have stochastic-type error factors embedded in the calculations. The error signal requires assessment as to whether any action is required, and so there is a question as to "degree" of error, and if the error is stochastic, issues of variance in this error variable will factor into the severity assessment. For example, if that error has "three-sigma" variance, no action may be decided, as the situation error estimate is poor. We see almost no research addressing these concerns.

SUMMARY

This article is intended to create discussion in the Information Fusion (IF) community about taking broader and systemic views of fusion process designs and addressing the consequently more-systemic impacts of such views on process designs. Here, we have probed into the Level 2 Situation Estimation space with some ideas on this type of thinking and about impacts to IF-based process designs. A main motivation here is toward realization of new opportunities and challenges for the IF community, and that addressing such challenges broadens the impact that this community can have across a very wide range of applications. We need to step away from functionally isolated data fusion R&D.

REFERENCES

1. Llinas, J. Reexamining information fusion-decision making inter-dependencies, in cognitive methods in situation awareness and decision support (CogSIMA). In *IEEE International Inter-Disciplinary Conference on Cognitive Situation Management*, San Antonio, TX, 2014, 1–6.
2. Llinas, J., and Malhotra, R. An expanded framework for situation control. *Frontiers in Systems Neuroscience*, Vol. 16, (2022), doi:10.3389/fnsys.2022.796100
3. Llinas, J. Information fusion for natural and man-made disasters. In *Proceedings of the Fifth International Conference on Information Fusion*, Annapolis, MD, July 2002, 570–576.
4. Endsley, M. R., Bolté, B., and Jones, D. G. *Designing for Situation Awareness: An Approach to User-Centered Design*. Boca Raton, FL: CRC Press, 2003.
5. Pirolli, P., and Card, S. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of International Conference on Intelligence Analysis*, McLean, VA, 2005, 6.
6. Klein, G., Phillips, J. K., Rall, E. L., and Peluso, D. A. A data-frame theory of sensemaking. In *Expertise Out of Context: Proceedings of the Sixth*

International Conference on Naturalistic Decision Making, R. R. Hoffman, ed. Mahwah, NJ: Lawrence Erlbaum Associates Publishers, 2007, 113–155.

7. Jakobson, G., Buford, J., and Lewis, L. A framework of cognitive situation modelling and recognition. In *Proceedings of the 25th IEEE Military Communications Conference (MILCOM)*, Washington DC, Oct. 2006, 1–7.
8. Jakobson, G., Buford, J., and Lewis. L. Situation management: Basic concepts and approaches. In *Proceedings of the 3rd International Workshop on Information Fusion and Geographic Information Systems*, St. Petersburg, Russia, 2007.
9. Jakobson, G. Introduction to cognitive situation management for tactical operations. *IEEE Communications Society, Distinguished Lecture*, (Jun. 2008).
10. Jakobson, G. A framework for cognitive situation control. In *Proceedings Cognitive and Computational Aspects of Situation Management (CogSIMA), 2017 IEEE Conference on Cognitive Situation Management*, Savannah, GA, Mar. 2017.
11. Roy, J. From data fusion to situation analysis. In *Proceedings of Fourth International Conference on Information Fusion (Vol. 2)*, Sun City, South Africa, Aug. 2001.
12. Lambert, D. A. Situations for situation awareness. In *Proceedings of the Fourth International Conference on Information Fusion (FUSION 2001)*, Montreal, Canada, Aug. 2001.
13. Brower, J. E. Abelard's theory of relations: Reductionism and the Aristotelian tradition. *The Review of Metaphysics*, Vol. 51, 3 (1998), 605–631.
14. Brower, J. E. Relations without polyadic properties: Albert the great on the nature and ontological status of relations. *De Gruyter*, Vol. 83, 3 (2001), 225–257, doi:10.1515/agph.2001.001
15. Costa, P. D., Mielke, I. T., Pereira, I., and Almeida, J. P. A. A model-driven approach to situations: Situation modeling and rule-based situation detection. In *Proceedings 2012 IEEE 16th International Enterprise Distributed Object Computing Conference*, Beijing, China, 2012, 154–163.
16. Baumgartner, N., and Retschitzegger, W. A survey of upper ontologies for situation awareness. In *Proceedings of the 4th IASTED International Conference on Knowledge Sharing and Collaborative Engineering*, St. Thomas, US VI, Nov. 2006, 1–9.
17. Kokar, M. M., Matheus, C. J., and Baclawski, K. Ontology-based situation awareness. *Information Fusion*, Vol. 10, (Jan. 2009), 83–98, doi:10.1016/j.inffus.2007.01.004
18. Steinberg, A. N. Principles and methods of situation assessment. *Advances and Challenges in Multisensor Data and Information Processing*, Vol. 8, (2007), 339.
19. Barwise, J., and Perry, J. *Situations and Attitudes*. Cambridge, MA: Bradford Books, MIT Press, 1983.
20. Jousselme, A. L., Maupin, P., and Bossé, É. Uncertainty in a situation analysis perspective. In *Proceedings of the Sixth International Conference of Information Fusion*, Cairns, QLD, Australia, 2003, 1207–1214.
21. Azarewicz, J., Fala, G., and Heithecker, C. Template-based multi-agent plan recognition for tactical situation assessment. In *Proceedings of the Fifth IEEE Conference on Artificial Intelligence Applications*, Miami, FL, 1989.
22. Cardell-Oliver, R., and Liu, W. Representation and recognition of situations in sensor networks. *IEEE Communications Magazine*, (Mar. 2010), 112–117.
23. Dousson, C., Gaborit, P., and Ghallab, M. Situation recognition: Representation and algorithms. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*, Chambery, France, 1993, 166–172.
24. Almeida, J. P. A., Costa, P. D., and Guizzardi, G. Towards an ontology of scenes and situations. In *Proceedings 2018 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA)*, Boston, MA, 2018, 29–35.
25. Little, E. G., and Rogova, G. L. Designing ontologies for higher level fusion. *Information Fusion*, Vol. 10, 1 (2009), 70–82.
26. Boury-Brisset, A.-C. Ontology-based approach for information fusion. In *Proceedings of the Sixth International Conference on Information Fusion*, Cairns, Queensland, Australia, 2003, 522–529.
27. Elsaesser, C., and Stech, F. Detecting deception. In *Adversarial Reasoning: Computational Approaches to Reading the Opponent's Mind*. Boca Raton, FL: Chapman and Hall, 2007, ch. 2.1.
28. Rogova, G. L., and Bosse, E. Information quality in information fusion. In *Proceedings 2010 13th International Conference on Information Fusion*, Edinburgh, UK, Jul. 2010, 1–8.
29. Klein, G., Moon, B., and Hoffman, R. R. Making sense of sensemaking 2: A macrocognitive model. *IEEE Intelligent Systems*, Vol. 21, 5 (Sep./Oct. 2006).
30. Austerweil, J. L, Sanborn S, and Griffiths, T.L. Learning how to generalize. *Cognitive Science*, Vol. 43, 8, (2019), e12777.
31. Shepard, R. N. Towards a universal law of generalization for psychological science. *Science*, Vol. 237, (1987), 1317–1323.
32. Kemp, C., Perfors, A., and Tenenbaum, J. B. Learning overhypotheses. In *Proceedings of the Twenty-Eighth Annual Conference of the Cognitive Science Society*, Vancouver, Canada, 2006, 417–422.
33. Gross, G., Nagi, R., and Sambhoos, K. Soft information, dirty graphs and uncertainty representation/processing for situation understanding. In *Proceedings 2010 13th International Conference on Information Fusion*, Edinburgh, UK, Jul. 2010, 1–8.
34. Kahneman, D. *Thinking, Fast and Slow*. New York: Macmillan Publishers, 2011.

James Llinas earned B.S. and M.S. degrees in Aerospace Engineering from New York University and carried out advanced rocket research at the Cornell Aeronautical Laboratory for several years. He transitioned to studies in applied mathematics and received a Ph.D. from the University at Buffalo, New York. This led to 10 years of applied research in intelligence-related technologies and programs where he led several programs in space-based intelligence. During this period, he joined the JDL Data Fusion Group as a consultant and was part of the core group that developed the JDL Data Fusion Model. He then joined the University at Buffalo where he founded the Center for Multisource Information Fusion (CMIF), the only academic center in the United States carrying out systems-based research in Data Fusion. Over more than 25 years, CMIF has carried out numerous funded programs studying both basic principles of Data Fusion as well as complex, systems-type research for major programs.

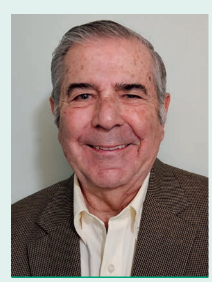

2023 ISIF Awards

W. Dale Blair
Georgia Tech Research Institute
Smyrna, GA, USA
dale.blair@gtri.gatech.edu

ISIF YAAKOV BAR-SHALOM AWARD FOR LIFETIME OF EXCELLENCE IN INFORMATION FUSION

The International Society of Information Fusion (ISIF) Yaakov Bar-Shalom Award for a Lifetime of Excellence in Information Fusion is the premier ISIF award. This award is given for a lifetime of contributions to information fusion. The ISIF Board of Directors and the ISIF Awards Committee are pleased to announce that Dr. Roy Streit is the recipient of the highly prestigious 2023 *ISIF Yaakov Bar-Shalom Award for a Lifetime of Excellence in Information Fusion*. Dr. Streit's contributions to multitarget tracking, his leadership in the ISIF community, and his founding of ISIF *Perspectives on Information Fusion* magazine as the Editor-in-Chief were the basis for his nomination by Stefano Coraluppi. The 2023 award is to be presented at the 2023 FUSION conference in Charleston, SC, USA in June 2023.

2023 recipient of the ISIF Yaakov Bar-Shalom Award for a Lifetime of Excellence in Information Fusion is Roy Streit.

Dr. Streit received his B.A. (Physics and Mathematics) from the East Texas State University (now Texas A&M-Commerce) in 1968, his M.A. (Mathematics) from the University of Missouri at Columbia in 1970, and the Ph.D. from the University of Rhode Island in 1978. He is a Distinguished Fellow at Metron in Reston, VA, a scientific consulting company specializing in search, tracking, and related problems. Prior to joining Metron in 2005, he was a Senior Research Scientist in the Senior Executive Service (SES) at the U.S. Naval Undersea Warfare Center (NUWC). The focus of his work at NUWC was the development, evaluation, and application of multi-sensor data fusion algorithms in support of submarine sonar and combat control automation.

Dr. Streit is an IEEE Life Fellow. He holds nine U.S. patents (two have D-10 status) and has published several books, over 40 journal articles, and numerous conference papers and book chapters. He served for many years on the ISIF Board of Directors including as ISIF President (2012). As mentioned, he was the founding Editor-in-Chief of the ISIF *Perspectives* magazine. He was General Co-Chair (with Murat Efe) for FUSION 2013 in Istanbul, Turkey, and General Co-Chair (with X. Rong Li) of FUSION 2017 in Xi'an, China.

The ISIF Lifetime of Excellence Award recognizes a researcher or engineer for outstanding contributions to the field of information fusion throughout his or her career. Contributions include technical advances, technical vision and leadership, education and mentoring, novel applications of information fusion and the associated engineering achievements, and service to ISIF. Following the award to Prof. Bar-Shalom, subsequent awardees include Dr. Chee-Yee Chong (2016), Prof. Pramod Varshney (2018), and Mr. Ed Waltz (2021).

2023 ISIF YOUNG INVESTIGATOR AWARD

The ISIF Young Investigator Award acknowledges the exceptional accomplishments of a young ISIF member within the field. With this award, ISIF aims to inspire personal dedication and promote involvement from emerging researchers and engineers. The recipient of the 2023 ISIF Young Investigator Award is Florian Pfaff. The basis for Dr. Pfaff's nomination is his contributions to the field of information fusion in estimation of nonlinear manifolds and information fusion for combined stochastic and set-membership uncertainties. Prof. Lyudmila Mihaylova nominated Dr. Pfaff for the 2023 award that is to be presented at the 2023 FUSION conference in Charleston, SC, USA in June 2023.

2023 recipient of the ISIF Young Investigator Award is Florian Pfaff.

Florian Pfaff is a highly accomplished researcher and an integral member of the FUSION community. His innovative filters for nonlinear manifolds have outperformed previous methods and offer effective solutions for various topologies. In addition, Pfaff has contributed to distributed estimation, multitarget tracking, and extended object tracking.

Dr. Pfaff's work has been widely recognized with numerous awards, honors, and distinctions, including the SICK Science Award 2018 for best Ph.D. thesis and the Otto von Guericke Award 2019 for the best project by the funding agency. Past recipients are David Crouse (2016), Marcus Baum (2017), Karl Granstrom (2018), and Florian Meyer (2021).

2023 ISIF BOB LYNCH DISTINGUISHED SERVICE AWARD

2023 recipient of the ISIF Bob Lynch Distinguished Service Award is Chee-Yee Chong.

ISIF sponsors the ISIF Robert Lynch Award for Distinguished Service to recognize an individual who has provided great service to the society. The award was established in memory of Bob Lynch, who contributed regularly and tirelessly over many years to the organization of the annual International Conference on Information Fusion (ICIF), the founding and production of the *Journal for Advances in Information Fusion* (JAIF), the founding of *ISIF Perspectives on Information Fusion*, and maintenance of the ISIF web site. Dr. Chee-Yee Chong is the recipient of the 2023 ISIF Bob Lynch Distinguished Service Award. Dr. Yaakov Bar-Shalom nominated Dr. Chong for his numerous contributions to ISIF. Dr. Chong was a co-founder of the International Society of Information Fusion (ISIF) and served as ISIF President in 2004 and ISIF Treasurer from 1998 to 2022, a quarter of a century. Dr. Chong was General Co-chair for the 12th International Conference on Information Fusion in 2009, Guest Editor for the special issue on Multiple Hypothesis Tracking in *ISIF Journal of Advances in Information* (JAIF), and a past Associate Editor for the JAIF. The 2023 ISIF Bob Lynch Distinguished Service Award is to be presented at the 2023 FUSION conference in Charleston, SC, USA in June 2023.

FUSION 2022 Conference Awards

Anne-Laure Jousselme
CS Group - France
La Garde, France
Anne-Laure.Jousselme@csgroup.eu

Linköping, Sweden welcomed the 25th International Conference on Information Fusion with an increased attendance and paper submission after the COVID-19 period. This year, 156 papers accepted to be published in the proceedings of the conference were in competition for the best paper awards in two categories: student and regular. Compared to previous years, the procedure for the award attribution was slightly updated, allowing any member of the organization committee to be considered, with the exception of the general chairs and the award chairs. The technical chairs first identified a long list of 12 papers in each category. To guarantee the fairness of the process, Joakim Jaldén, a technical chair with no conflict of interest, acted as a screener and helped the award chair to shorten the list. A short list of six papers in each category was then issued, using an objective ranking of papers based on the criteria *average score*, *average score weighted by confidence*, and *award points*, as provided for the reviewing process. The Award Committee was selected by the General Chairs: Chee-Yee Chong, Paulo Costa, Pieter de Villiers, Jean Dezert, and Anne-Laure Jousselme (Chair). The committee members independently reviewed all of the papers and ranked them within each of the two categories. Winners then emerged by a simple summation of the ranks, with a clear consensus. Because *two* student papers achieved the same top ranking, it was agreed that both should receive the Best Paper Award, without the need to further discriminate between them. Six papers were recognized at the gala dinner of the FUSION 2022 conference by General Cochair Gustaf Hendeby and Award Chair Anne-Laure Jousselme. On behalf of the International Society of Information Fusion, congratulations to all candidate papers and an obvious special mention to the winners!

Ive Weygers, Jean-Pierre Le Cadre Best Paper award recipient.

JEAN-PIERRE LE CADRE AWARD

- Best Paper: Daniel Laidig, Ive Weygers, Simon Bachhuber, and Thomas Seel, "VQF: A Milestone in Accuracy and Versatility of 6D and 9D Inertial Orientation Estimation"
- First Runner-up: Kailai Li, Florian Pfaff, and Uwe Hanebeck, "Circular Discrete Reapproximation"
- Second Runner-up: Simon Bachhuber, Daniel Weber, Ive Weygers, and Thomas Seel, "RNN-based Observability Analysis for Magnetometer-Free Sparse Inertial Motion Tracking"

BEST PAPER (JEAN PIERRE LE CADRE AWARD)

Daniel Laidig, Ive Weygers, Simon Bachhuber, and Thomas Seel, "VQF: A Milestone in Accuracy and Versatility of 6D and 9D Inertial Orientation Estimation"

Abstract—We present a novel quaternion-based inertial orientation estimation filter. Inclination drift from gyroscope strapdown integration is corrected from specific force measurements that are low-pass filtered in an almost-inertial frame to effectively compensate for instantaneous accelerations and decelerations. Heading drift is corrected via a scalar heading offset. The resulting decoupled state representation provides simultaneous 6D and 9D orientation estimation. We systematically evaluated the method on a rich orientation estimation benchmark dataset and show that the proposed method clearly outperforms three of the currently most commonly adopted and accurate inertial orientation estimation filters. The filter is available as open-source software, and its parameters are tuned to work well for a wide range of movements and application scenarios. Our fundamentally different filtering approach with a decoupled state representation and novel inclination correction resulted in a new level of accuracy, with a 41% improvement of the total orientation error and doubling the inclination accuracy. This facilitates new and exciting high-precision applications in the field of inertial motion tracking.

TAMMY L. BLAIR AWARD

- Best Student Paper (Tie): Alessandro D'Ortenzio, Costanzo Manes, and Umut Orguner, "A Model Selection Criterion for the Mixture Reduction Problem Based on the Kullback-Leibler Divergence"
- Best Student Paper (Tie): Runze Gan, Qing Li, and Simon Godsill, "A Variational Bayes Association-based Multi-object Tracker under the Non-homogeneous Poisson Measurement Process"
- Runner-Up: Mingchao Liang and Florian Meyer, "Neural Enhanced Belief Propagation for Data Association in Multiobject Tracking"

Qing Li and Runze Gan, Best Student Paper award recipients (tie).

TIE: BEST STUDENT PAPER (TAMMY L. BLAIR AWARD)

Runze Gan, Qing Li, and Simon Godsill, "A Variational Bayes Association-based Multi-object Tracker under the Non-homogeneous Poisson Measurement Process"

Abstract—The non-homogeneous Poisson process (NHPP) has been widely used to model extended object measurements where one target can generate zero or several measurements; it also provides an elegant solution to the computationally demanding data association problem in multiple target tracking. This paper presents an association-based NHPP system, coupled with which we propose a variational Bayes association-based NHPP (VBAbNHPP) tracker that can estimate online the target kinematics and the association variables in parallel. In particular, the VBAbNHPP tracker can be easily extended to include online static parameter learning (e.g., measurement rates) based on a general coordinate ascent variational filtering framework developed here. The results show that the proposed VB-AbNHPP tracker is superior to other competing methods in terms of implementation efficiency and in tracking accuracy.

Alessandro D'Ortenzio, Best Student Paper award recipient (tie).

TIE: BEST STUDENT PAPER (TAMMY L. BLAIR AWARD)

Alessandro D'Ortenzio, Costanzo Manes, and Umut Orguner, "A Model Selection Criterion for the Mixture Reduction Problem Based on the Kullback-Leibler Divergence"

Abstract—In order to be properly addressed, many practical problems require an accurate stochastic characterization of the involved uncertainties. In this regard, a common approach is the use of mixtures of parametric densities which allow, in general, to arbitrarily approximate complex distributions by a sum of simpler elements. Nonetheless, in contexts like target tracking in clutter, where mixture of densities are commonly used to approximate the posterior distribution, the optimal Bayesian recursion leads to a combinatorial explosion in the number of mixture components. For this reason, many mixture reduction algorithms have been proposed in the literature to keep limited the number of hypotheses, but very few of them have addressed the problem of finding a suitable model order for the resulting approximation. The commonly followed approach in those algorithms is to reduce the mixture to a fixed number of components, disregarding its features which may vary over time. In general, finding an optimal number of components is a very difficult task: once a meaningful optimality criterion is identified, potentially burdensome computational procedures must be devised to reach the optimum. In this work, by exploiting the optimal transport theory, an efficient and intuitive model selection criterion for the mixture reduction problem is proposed.

ISIF Working Group Report

Paulo C. G. Costa
George Mason University
Fairfax, VA, USA
pcosta@gmu.edu

Anne-Laure Jousselme
CS Group - France
La Garde, France
Anne-Laure.Jousselme@csgroup.eu

On the occasion of the 25th anniversary of ISIF, this paper presents the origin, challenges and key developments of the Evaluation of Techniques for Uncertainty Representation Working Group (ETUR WG), sponsored by ISIF since 2012.

UNCERTAINTY REPRESENTATION AND REASONING EVALUATION FRAMEWORK

Although in recent times chat generative pretrained transformer, commonly known as ChatGPT, has become the epitome of artificial intelligence ubiquity in human lives, the truth is that we have long been subjected to increasingly pervasive sensors, wide availability of large volumes of heterogeneous data, easily accessible machine learning frameworks, and other aspects that enable such systems to exist. Furthermore, the seamless fashion in which such systems pervade our daily lives usually disguises the complexity of the interactions that happen among the various information systems so that data can be properly accessed. Uncertainty management is a key aspect of these interactions and is a critical component to ensure sound results when using multiple data sources. This is especially true when the underlying sources of uncertainty are also heterogeneous, such as in systems that operate above level 2 of the Joint Directors of Laboratories (JDL) framework, a.k.a. high-level information fusion (HLIF) systems. Not surprisingly, the problem of uncertainty representation and reasoning in HLIF systems has attracted interest that extends beyond the information fusion (IF) community.

Even in modern times, fusing hard and soft information from diverse sensor or source types (human-as-a-sensor included) and the associated uncertainty is a task that still relies heavily on human intervention, creating a scalability conundrum that current technologies are incapable of solving. Despite the widespread acknowledgment that HLIF systems must support automated knowledge representation and reasoning in the presence of uncertainty, there is no consensus on the appropriate approach to adopt (which theory, uncertainty function or model, fusion rule, etc.), on the performance criteria that should guide the design of an HLIF system in terms of uncertainty handling, or on how to assess such criteria.

IF applications typically must deal with information that is incomplete, imprecise, inconsistent, and otherwise in need of a sound methodology for representing and managing uncertainty. Complex and dynamic use cases make such tasks even more difficult, because for the same input conditions, apparently minor differences in how uncertainty is handled may drastically affect the output of the IF process. Evaluation of information fusion systems (IFSs) presents intrinsic challenges due to their complexity and the sheer number of variables influencing their performance. Low-level IF tasks generally address random phenomena for which numerical data are collected. The impact of uncertainty representation is well understood and generally quantifiable. However, higher levels of IF tasks need to handle uncertainty not only due to the variability of data (aleatory uncertainty) but also due to lack of knowledge (epistemic uncertainty). The approach chosen for representing uncertainty has an overall impact on system performance that is hard to quantify or even to assess from a qualitative viewpoint.

The evaluation of how uncertainty is addressed within a given IFS is distinct from, although closely related to, the evaluation of the overall performance of the system. Metrics for evaluating the overall performance of an IFS are more encompassing in scope than those focused on the uncertainty handling within the system. The metrics for the overall system capture not only the effects of the uncertainty representation but also the effects of other aspects that can affect the performance of the system (e.g., the implementation).

In 2010, an Uncertainty Forum was organized by Simon Maskell and John Lavery from the U.S. Army Research Office as part of the International Conference on Information Fusion (FUSION) held in Edinburgh, United Kingdom, to discuss some different ways of representing and dealing with uncertainty using a common and single scenario as a reference point. As the organizers mentioned, "The goal of the Uncertainty Forum is not to come to specific conclusions about a linear or other ranking of approaches for representing uncertainty but rather to widen the spectrum of available options and link these options with situations in which they perform well".[1] Prior to this forum, a vehicle-borne improvised explosive device (V-IED) scenario was submitted to five scientists, with each expert "defending" one approach or framework: the Bayesian method (Simon Godsill), Dempster–Shafer theory (Arnaud Martin), transferable belief model (David Mercier), Dezert–Smarandache theory (Jean Dezert), and human intelligence/processing (Peter Gill). The analysis conducted in [1] revealed that beyond the mathematical framework selected, personal choices of modeling impact the solution provided and the results obtained.

Hence, to help fusion scientists better sail among the different approaches dealing with uncertainty, the International Society of Information Fusion (ISIF) charted the Evaluation of Technologies for Uncertainty Representation (ETUR) working group,[2] which has been discussing this topic since FUSION 2012 in Singapore. The goal of this group is to provide a forum to address the problem of the assessment and evaluation of the different uncertainty representation approaches developed so far. The ultimate objective would be to provide objective cri-

[1] http://isif.org/fusion/proceedings/fusion2010/plenary-speakers.htm
[2] https://eturwg.c4i.gmu.edu

teria to assess specifically how uncertainty is handled in fusion systems and define basic concepts to be eventually accepted and standardized.

The main outcome is the uncertainty representation and reasoning evaluation framework (URREF), which includes an ontology, evaluation procedures and associated datasets, applications, and other components that aim at providing the foundational theory, mechanisms, and standardization artifacts required to evaluate the impact of uncertainty in information fusion systems. The framework provides a means for relating evaluation criteria specifically focused on uncertainty handling, with other information quality aspects such as the nature of uncertainty (aleatory vs. epistemic), the derivation of uncertainty (objective vs. subjective), the type of uncertainty (imprecision vs. uncertainty), and the uncertainty theory (belief functions, probability, fuzzy sets, possibility, etc.) [2], [3].

Within the URREF, a major task is to formally identify the concepts that are pertinent to the evaluation of uncertainty of an IFS, which is a seen as a means to ensure that all evaluations follow the same semantic constraints and abide by the same principles of mathematical soundness. This is enabled by the ontology reference model developed for the framework, known as the URREF ontology. The first stone to formal representation of the uncertain reasoning domain was put by the Uncertainty Reasoning for the World Wide Web Incubator Group of the World Wide Web Consortium, which published in March 2008 an uncertainty ontology "to demonstrate some basic functionality of exchanging uncertain information".[3] This effort was then pursued through the ISIF ETUR working group.

The URREF and its ontology component were developed through an iterative process, an essential part of which was to apply the framework to a set of use cases. The use cases not only serve as benchmarks but also reflect a range of considerations relevant to evaluation of uncertainty representation within the context of an overall fusion application. The focus is on high-level fusion tasks, which require a closer human interaction (human as a source or as a decision-maker). Applying the framework to use cases and their associated datasets grounds the ideas in concrete application areas and helps to uncover requirements that emerge as the framework is applied to a concrete problem.

As such, the work on developing these use cases has been generating new insights and requirements for the URREF (e.g., [4], [5], [6]). Among the different use cases proposed through the years, the three that have been mostly consistent throughout the discussions are *maritime surveillance,* where a harbor area is monitored by a set of sources mixing sensors and humans [6]; *rhino poaching,* which involves a decision support system that directs the patrol effort of the rangers to the areas with elevated risk of poaching [7]; and *cyber threat,* which comprises an expert model for cyber threat analysis [8].

The URREF is not a system or software application that can be "directly applied" to a use case. Yet the use cases are essential for the group to achieve an understanding of all the nuances and idiosyncratic aspects of the process of evaluating techniques that are fundamentally different in their assumptions and views of the world. They provide the grounding for establishing the URREF concepts and mechanisms needed to mitigate the effects that the underlying assumptions of each theory have in biasing the design of evaluations—each usually geared toward the strengths of one technique at the expense of the others. The URREF does not completely remove the subjectivity and biases involved in evaluating uncertainty representation techniques, but it is a strong step in that direction.

We offer a final thought about the ETUR working group, the URREF, and its unconventional nature and contributions to the IF society. On the one hand, the problem being addressed is a fundamental issue that requires a deep understanding of the many aspects of IF in general, of HLIF in particular, and of uncertainty theories. It is no coincidence that the group has such a wide background among its members, mixing expertise across the whole spectrum of the JDL model. On the other hand, its results can be assessed in terms of knowledge shared and formally captured about a difficult problem. This is when the group's work and contributions shine. Since 2012, it met in all FUSION conferences, with roughly 150 biweekly meetings at the time of this writing and 12 ETUR special sessions. In addition, a *Journal of Advances in Information Fusion* special issue, more than 70 articles in FUSION conferences, and tutorials, panels, and other events have brought a more throughout understanding of its topic and its importance to our community.

REFERENCES

1. Jousselme, A.-L., and Maupin, P. Uncertainty representations for a vehicle-borne IED surveillance problem. In *Proc. Int. Conf. Information Fusion,* Singapore, July 2012.
2. Costa, P. C., Laskey, K. B., Blasch, E., and Jousselme, A.-L. Towards unbiased evaluation of uncertainty reasoning: The URREF ontology. In *Proc. Int. Conf. Information Fusion,* Singapore, July 2012.
3. Costa, P. C. G., Jousselme, A.-L., Laskey, K., Blasch, E., Dragos, V., Ziegler, J., et al. URREF: Uncertainty representation and reasoning evaluation framework for information fusion. *Journal of Advances in Information Fusion,* Vol. 13, 2 (Dec. 2018), 137–157.
4. Dragos, V., Ziegler, J., and de Villiers, J. P. Application of URREF criteria to assess knowledge representation in cyber threat models. In *Proc. 21st Int. Conf. Information Fusion,* Cambridge, United Kingdom, 2018.
5. Pavlin, G., Jousselme, A.-L., De Villiers, J. P., Costa, P. C. G., and de Oude, P. Towards the rational development and evaluation of complex fusion systems: a URREF-driven approach. In *Proc. Int. Conf. Information Fusion,* Cambridge, United Kingdom, July 2018.
6. Jousselme, A.-L. Semantic criteria for the assessment of uncertainty handling fusion models. In *Proc. 19th Int. Conf. Information Fusion,* Heidelberg, Germany, July 2016.
7. Koen, H., De Villiers, J. P., Pavlin, G., de Waal, A., de Oude, P., and Mignet, F. A framework for inferring predictive distributions of rhino poaching events through causal modelling. In *Proc. Int. Conf. Information Fusion,* Salamanca, Spain, 2014.
8. Kiesling, T., Krempel, M., Niederl, J., and Ziegler, J. A model-based approach for aviation cyber security risk assessment. In *Proc. 11th IEEE Int. Conf. Availability, Reliability and Security,* Salzburg, Austria, 2016, 517–525.

[3] https://www.w3.org/2005/Incubator/urw3/group/draftReport.html

FUSION 2022 Report

REPORT ON THE 25TH INTERNATIONAL CONFERENCE ON INFORMATION FUSION

Fredrik Gustafsson
Gustaf Hendeby
Linköping University
Linköping, Sweden
fredrik.gustafsson@liu.se
gustaf.hendeby@liu.se

Terence van Zyl
University of Johannesburg
Johannesburg, South Africa
tvanzyl@gmail.com

The 25th International Conference of Information Fusion, FUSION 2022, was held in Linköping, Sweden, July 4–7, 2022. This was the second time FUSION was held in Sweden; the first time was in Stockholm in 2004. We considered Linköping to be the most appropriate place for this second time for many reasons. For the past 100 years, Linköping has been the aerospace capital of northern Europe, with two airfields, aircraft manufacturing industries, air force training, and an entire ecosystem of innovation and exploration built around this. Furthermore, the city hosts the Swedish Defence Research Agency, the Swedish Air Force Museum, an upcoming national cyber security centre, and many local industries with strong competencies and large activities in the information fusion area.

Gustaf Hendeby welcoming the group at the Swedish Air Force Museum.

Linköping University was founded in the 1960s as an engineering university to support Swedish aerospace and other industries. Since it was established, it has created more than 150% growth across the entire city. The science park has attracted 350 companies that employ more than 7,000 people, focusing on areas such as automotive safety, telecommunications, medical technology, and the Internet of Things. Linköping University, with its electrical engineering, computer science, and mechanical engineering departments, has contributed hundreds of papers to FUSION over the years. Its strong position in Sweden is manifested by it being part of the $600 million (6.1 SEK) Wallenberg Artificial Intelligence, Autonomous Systems and Software Program, focusing on autonomous systems and machine learning.

Beyond science and engineering, Linköping has a compact and beautiful city centre within walking distance of the conference venue, which made it easily accessible to conference participants. Linköping has a rich history, dating back more than 700 years. It hosts both the bishop of one of Sweden's first and most notable dioceses (1107) and a beautiful 800-year-old cathedral. The group tour of this historical heritage prior to the icebreaker was well received, as was the welcome reception at the Swedish Airforce Museum.

A long time has passed since we presented our winning bid to host FUSION in Linköping while in Cambridge, United Kingdom, in 2018. We had expressed our intention to bid for 2021 to the International Society of Information Fusion board while at FUSION in Xi'an, China, in 2017. That being said, the whole planning process lasted for more than five years. Indeed, our planning was interrupted by the COVID-19 pandemic. First, the conference was postponed by a year to give South Africa a second opportunity. Furthermore, we had the challenging task of planning for the conference in the face of significant uncertainty surrounding the unfolding pandemic. Fortunately, we had close collaborations with the local conference organiser, venue, and hotels, which were more than understanding and supportive. In the end, we are grateful that the pandemic

Fredrik Gustafsson downtown on the Linköping main square.

FUSION 2022 gala dinner at the Linköping Concert and Congress Center.

did not severely affect our conference again. However, some countries still had complex travel restrictions, and we had to be flexible and accommodating to alternatives for these authors. As a result, slightly more than 10% of the presentations had prerecorded videos instead of regular presentations. We could also see that the pandemic had residual consequences which did affect both the number of submitted papers and the number of attendees travelling to Sweden.

Looking at FUSION 2022 in numbers, we received in total 228 paper submissions. These were evaluated in 866 reviews, on average 3.8 reviews per paper, and no paper got fewer than three reviews. In total, 168 papers were accepted, leading to an acceptance rate of 74%.

There were 311 participants in all, consisting of 222 early-bird and 51 regular registrations. Five students from across the globe were awarded travel grants and free registration. In fact, of all the participants, more than a third registered as students. Furthermore, besides the volunteers, we had 21 registered sponsors. We are proud to have attracted many sponsors and are thankful for their contributions. Out of the 14 sponsors, four opted to be gold sponsors.

Despite the aftermath of COVID-19, participants from 28 countries attended, from every continent short of Antarctica. The conference had representatives from 170 different universities, companies, and organisations. It was obvious that many had a strong desire to meet in person to discuss, collaborate, eat, and drink amongst friends and colleagues after the prior few years of disruptions.

This excitement manifested itself in a record for the average number of participants per paper, with more than 1.8 participants per paper, which was an all-time high in the history of FUSION (excluding the fully online conference in 2020). The evening social events were fully booked, with many of our colleagues bringing partners, spouses, and family along. The need for personal interactions and a meeting of colleagues and friends to discuss science and engineering was obvious. Everyone was exited that it was once again possible to meet our FUSION family! Who would have thought that being normal would be what we want the most? We are, starting with FUSION 2022, hopefully back to normal again. As the general chairs, we thank you on behalf of the organizing committee.

Wolfgang Koch presents the Maria Window during the cathedral tour.

SDF 2022 Report

IMPRESSIONS OF THE 14TH IEEE AESS SYMPOSIUM "SENSOR DATA FUSION – TRENDS, SOLUTIONS, APPLICATIONS" (SDF 2022)

Felix Govaers
Wolfgang Koch
Fraunhofer FKIE
Wachtberg, Germany
felix.govaers@fkie.fraunhofer.de
w.koch@ieee.org

After a break of two years due to COVID-19, the 14th IEEE AESS Symposium *Sensor Data Fusion: Trends, Solutions, and Applications* took place from October 12–14, 2022 in Bonn, Germany. It was organized by the Department of Sensor Data and Information Fusion at Fraunhofer FKIE, with active participation of the international expert community, technical sponsorship from IEEE AESS and financial sponsorship from ISIF and industry partners IBM Defence Germany, Diehl Defence BGT, Spherea, Hensoldt, and Schönhofer SSE. The ISIF sponsorship made it possible to have the conference happen in the wonderful Uniclub Bonn.

The SDF Symposium continued the tradition of a single-track conference style at the location Uniclub Bonn, which is located next to the Rhine River in the center of the city. The number of participants (60) was slightly above the average of the past years. It was great to meet friends, colleagues and related researchers again; the participants enjoyed the in-person discussions after the involuntary break during the pandemic.

Participants from industry, universities, and research institutes from overseas and Europe presented and discussed current developments in modeling and applications of data fusion for intelligent systems as well as theoretical findings. The 24 presentations were grouped into seven sessions, because in 2022 the trend of the past years with increasing applications of machine learning for data fusion could not be observed. The session on deep learning-based methods consisted of four presentations mostly on classification tasks. For instance, Warre Geeroms et al. (University of Ghent) identified speakers based on fusion concept to combine audio streams with facial features detected in a synchronized video, whereas Jingxuan Su et al. (University of Sheffield) presented an approach to improve imbalanced data for semantic image segmentation. The audience could listen to and discuss model-based approaches from eight contributions on estimation theory and target tracking as well as from eight presentations on navigation and localization. Selected talks were, for instance, the presentation of Audun Hem et al. (Norwegian University) on a smart compensation on radar rotations within the framework of the Joint Integrated Probabilistic Data Association (JIPDA) and the presentation of Sutthiphong "Spot" Srigrarom (National University of Singapore) on a track-to-track association approach for person re-identification in multi-camera applications based on geometrical considerations of the convex hull topology.

Wolfgang Koch opening the SDF 2022 Symposium.

Simon Maskell during his keynote speech on large scale density sampling.

A highlight of the symposium was the keynote on "Towards Using Large-Scale Sequential Monte Carlo to Get Big Information out of Small Data" given by Simon Maskell, where he presented innovative approaches for using recent computer technology on sampling from target distributions in large data sets by parallelizing Monte Carlo methods. In 2022, the SDF Symposium introduced the novelty of an Industry Talk, which in this case was given by Martin Kugelmann from Spherea. In his talk entitled "What is Common to a Drilling Machine, a Medium-sized Company, a Research Institute and Music?" he moved along a proverbial path from innovation for technology, which enhances our daily life, to the sounds of music, which connects humanity by a common ground of sensing, feeling, and enjoying. His talk was the perfect transition to the piano recital which was given afterwards as an opening event for the conference gala dinner. The pianist Julia Rinderle[1] introduced the composers of the evening which were Beethoven (born in Bonn), Schumann (died in Bonn), and Chopin (might have come through Bonn on his way from Poland to Paris). Despite the technical background of her audience, the feedback we received was enthusiastic. Data fusion and classical music definitely are a match.

All publications from SDF 2022 can be found in IEEE Xplore. The upcoming SDF 2023, to be held again in Bonn, November 27–29, will be a joint event with the IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI).[1]

[1] https://juliarinderle.de/

SDF mfi23 Bonn

Combined SDF and MFI Conference

Call for Papers

Combined SDF and MFI

This year the combined IEEE 2023 Symposium Sensor Data Fusion and International Conference on Multisensor Fusion and Integration (SDF-MFI) will take place in Bonn, Germany. We are happy to announce the collaboration of two great conferences on robotics, data fusion, automation and intelligent systems in combined one-track conference. The Uniclub Bonn next to the Rhine river at the center of the former capital provides a great venue.

Topics

Theory: Probability theory, Bayesian inference, nonlinear estimation, Dempster-Shafer, fuzzy sets, logic, machine learning, neural networks, distributed architectures.

Sensors: RGB cameras, depth cameras, radar and sonar devices, laser scanner, infrared sensors, IMU, gyroscopes.

Algorithms for: tracking and localization, recognition, perception, AI in robotics, cognitive systems, sensor registration, big data, sensor management, distributed sensor systems, SLAM, visual servoing, learning by demonstration.

Applications: Sensor networks, multi-robot systems, distributed and cloud robotics, bio-inspired systems, service robots, automation, biomedical applications, autonomous vehicles (land, sea, air), manipulation planning and control, multifinger hands, micro/nano systems, surveillance, multimodal interface and human robot interaction, navigation, Internet-of-Things, smart cities, cyber-physical systems, Industry 4.0, search/rescue/audition, field robotics, swarm robotics, force and tactile sensing, surgical robotics, humanoids, soft-bodied robots.

Fees

€ 299.-	Students
€ 499.-	Regular

- For the student registration a proof of the student status is required.
- One registration covers one paper only.

Contributions

Prospective authors are encouraged to submit high-quality full draft papers (6-8 pages, IEEE format). All submissions are subject to a peer-review process by the technical program committee. Accepted and presented papers will be submitted to IEEE for publication. At least one of the authors of each accepted contribution is expected to register for the conference and to present the paper. For details contact `www.fkie.fraunhofer.de/sdf2023`.

Important Dates

31.08.2023	Submission of full draft papers
13.10.2023	Notification of acceptance
03.11.2023	Submission of the final version
27.11.2023	Start of the conference

Organisation

Executive Chairs:

- **Wolfgang Koch**, Fraunhofer FKIE and University of Bonn, w.koch@ieee.org
- **Uwe D. Hanebeck**, Karlsruhe Institute of Technology KIT, uwe.hanebeck@kit.edu

Technical Program Chairs:

- **Florian Pfaff**, Karlsruhe Institute of Technology KIT
- **Felix Govaers**, Fraunhofer FKIE, Germany

Technical Program Committee

Marcus BAUM, University of Göttingen, GER; Jürgen BESTLE, HENSOLDT, GER; Christian BRANDLHUBER, 21strategies, GER; Chee CHONG, Consultant, CA, USA; Stefano CORALUPPI, STO, MA, USA; Armin B. CREMERS, University of Bonn, GER; Daniel CREMERS, Technical University Munich, GER; Klaus DIETMAYER, University of Ulm, GER; Darin DUNHAM, Lockheed Martin, USA; Bharanidhar DURAISAMY, Daimler, GER; Murat EFE, Ankara University, TK; Frank EHLERS, FWG, GER; Dietrich FRÄNKEN, Hensoldt, GER; Jesus GARCIA, University Carlos III, Madrid, ES; Fredrik GUSTAFSSON, Linköping University, SW; Uwe D. HANEBECK, Karlsruhe Institute of Technology KIT, GER; Bernhard KRACH, Airbus, GER; Joerg KUSHAUER, Diehl BGT Defence, GER; Henry LEUNG, University of Calgary, CA; Lyudmila MIHAYLOVA, University of Sheffield, UK; Gee Wah NG, DSO, SGP; Umut ORGUNER, University of Ankara, TR; Johannes REUTER, University of Applied Sciences Konstanz, GER; Stefan REUTER, Robert Bosch GmbH, GER; Lauro SNIDARO, University of Udine, IT; Klaus-Dieter SOMMER, University of Ilmenau, GER; Roy L. STREIT, Metron Inc., USA; Jörn THIELECKE, Universität Erlangen, GER; Reiner THOMÄ, Technical University Ilmenau, GER; Martin ULMKE, Fraunhofer FKIE, GER;

VDE ITG KIT Karlsruhe Institute of Technology IEEE Robotics & Automation Society ies UNIVERSITÄT BONN Fraunhofer FKIE

25 Years of ISIF

INTRODUCTION

The International Society of Information Fusion (ISIF) was formed in 1998 to be the sponsor of FUSION 1999. Over the next 25 years, ISIF has sponsored 25 annual FUSION Conferences on five continents, published the *Journal of Advances in Information Fusion* (JAIF), and the ISIF *Perspectives on Information Fusion* magazine. This paper reflects on the progress of ISIF at its 25th anniversary and summarizes its achievements in serving the information fusion community.

ORGANIZATION

ISIF was incorporated as a non-profit organization in September 1998, mainly to be the sponsor of FUSION 1999 [1]. Since there were no initial members, an organizing committee was formed to elect the first ISIF Board in December 1998. Jim Llinas was the first president. The Board decided that attendees of FUSION conferences would automatically become members. With members from the attendees of FUSION 1999, the Board for 2000 was elected and Yaakov Bar-Shalom became the president. Such a procedure was made official later by the ISIF Constitution and Bylaws.

The organization of ISIF has evolved over the years. A constitutional amendment in 2004 created the position of President-Elect, who will be the President in the subsequent year. Another amendment in 2008 enlarged the Executive Committee to include the most recent past presidents. These changes provide more continuity in the management of ISIF. Membership in the Executive Committee continues to evolve to reflect the changing needs of ISIF. For example, the Vice President (VP) Social Media was added in 2022. The current ISIF Board of Directors consists of the President, President-Elect, Treasurer, Secretary, VP Communications, VP Conferences, VP Membership, VP Publications, VP Social Media, VP Working Groups, JAIF Editor-in-Chief (EiC), *Perspectives* EiC, the last two Past Presidents, and nine elected members, each serving three-year terms.

ISIF Board of Directors

FUSION CONFERENCES

ISIF was formed to sponsor FUSION 1999. Over the past 25 years, the International Conference on Information Fusion, simply known as the FUSION Conference, has emerged as the premier venue for the interchange of the latest research in information fusion and discussion of its impacts on our society. The conference is known for its inclusiveness for accepting innovative and valuable ideas that may not be completely polished. It is ultimately the people that are the lifeblood of any research community, and the FUSION Conference offers ample opportunity to meet with the leading experts in information fusion in both technical and social settings. It is the impromptu discussions that has led to new ideas and research projects. The combination of technical and social programs provides a unique experience for our research community.

The FUSION Conference that we know and love did take some time to evolve. A personal perspective of this evolution by the various conference organizers was reported last year in [2]. We will highlight a few critical milestones. The initial conference was a gamble by two Ph.D. entrepreneurs to hold a new

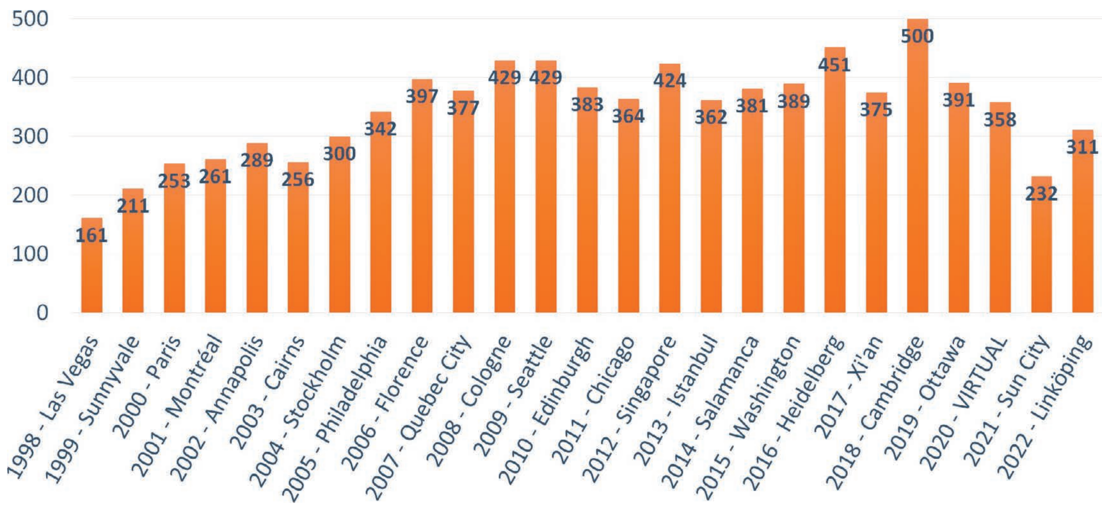

Figure 1
Attendance numbers for all 25 FUSION conferences.

data fusion conference in Las Vegas that was open to researchers from all over the world. Such an event did not exist before, and the two entrepreneurs had to start from scratch. They searched the internet for an influential, energetic, and willful data fusion researcher to help and attract papers. They identified and recruited X. Rong Li as that researcher and he agreed. Rong was able to rally a core group of experts to serve on the steering committee and convinced Yaakov Bar-Shalom to deliver the first keynote speech. The conference attracted 161 attendees and was a technical (but not financial) success that led to the formation of ISIF by attendees of FUSION 1998. The team then organized a second conference in Sunnyvale, CA that continued the technical success (attendance grew to 211) and provided surplus funds to help build up ISIF. The rest, as they say, is history.

Over the next decade, the technical program of the conference converged into the existing format that we all now expect with a tutorial day followed by three days each consisting of a plenary talk followed by three sessions of six–eight parallel tracks of oral presentations, along with a student paper program. The tutorial program and the student paper program were started at FUSION 2002. While the first two conferences did not have social programs, such programs started with FUSION 2000 in Paris—the first FUSION conference outside the United States—with a gala dinner on the Bateaux Mouches on the Seine River. Over time, each conference began to organize two–three social programs to provide experiences for our attendees that are unique to the location. The 2011 conference included a "5K fun run" that has been held ever since and offers attendees to the conference to participate in an early morning race around the hosting city.

Of the first few FUSION conferences, FUSION 2002 was particularly successful financially—it received record financial support and made a good profit of USD 46K, which contributed greatly to ISIF's "primitive accumulation of capital". It was also the first time for the IEEE to support the FUSION conference. By 2009, attendance grew to 432 participants with a record financial surplus of over USD 100K despite the Great Recession of 2008 and the swine flu pandemic of 2009. FUSION 2009 also marked the last year of significant sponsorship by US Government agencies such as the Army Research Office (ARO), Air Force Office of Scientific Research (AFSOR), and Office of Naval Research (ONR).

In the first decade, the conference locations expanded beyond the US into Europe, Canada, and one event in Australia. This led to a core cluster of ISIF researchers in North America and Europe. FUSION 2012 in Singapore marked the very first time the event was held in Asia. The intent was to grow ISIF beyond its research cores in North America and Europe and the conference attracted a record of 450 attendees and did help to create a new core of researchers within Asia and other parts of the world. The conference returned to Asia in 2017 with Xi'an, China serving as the host. Attendance peaked at 491 attendees at FUSION 2018 in Cambridge UK. At the 2017 conference in China, South Africa was picked to host FUSION 2020 with the idea of expanding the international reach of ISIF by building upon a strong cluster of researchers from that region.

Then everything changed in March 2020 with COVID-19. The South African team did an excellent job adjusting FUSION 2020 for the virtual format serving 278 attendees via the Zoom and Whova platforms. While the technical content remained strong, the lack of impromptu interactions within the hallways and at social events was clearly lacking. The ISIF Board, in conjunction with the South African team, decided to try again in 2021, hoping life would return to normal. The conference was even pushed back to November 2021 to improve the odds that people could travel. Then Omicron came and travel (while possible) was still very restricted. FUSION 2021 was a hybrid event with 170 virtual attendees and 50 in-person attendees, of whom 25 were international. The in-person attendees enjoyed a safari game drive. Nevertheless, most ISIF participants never experienced South Africa, and we hope the organizing team will eventually be able to rectify this unfortunate situation in the future after they recover from their two-year ordeal of decision making under unprecedented uncertainty.

FUSION 2022 returned to a fully in-person format in Sweden, but the effects of COVID-19 still lingered. Attendance increased to 311 but was still below pre-pandemic levels. As stated earlier, the combination of inclusivity, technical rigor, social interactions, and access to top fusion researchers is what makes the FUSION conferences so great. These features were difficult to offer during the height of the pandemic. We are hopeful that the conference will get back on its pre-pandemic trajectory (see Figure 1) and continue to grow its global influence, with FUSION 2024 in Venice, Italy and FUSION 2025 in Rio de Janeiro, Brazil.

JOURNAL OF ADVANCES ON INFORMATION FUSION

ISIF is surely best known for its highly successful FUSION conference series. However, the Society is much more than a series of conferences! Other key elements include its flagship publication, the *Journal of Advances in Information Fusion* (JAIF), the past and current working groups, the sponsorship of smaller workshops and symposia, and the website platform and social media presence that promote a constant exchange of ideas in the ISIF community and beyond.

JAIF was founded in 2005 and published its first issue in June 2006. An important characteristic of JAIF is that all published papers are freely available to the research community, without access restrictions. In fact, JAIF was the one of the first open-access journals. JAIF has a semi-annual cadence, and to date has yielded 34 issues over 17 years. The intent was, and remains, to provide a forum for high-caliber archival publications on information fusion, on par with the top IEEE journals and Transactions. Though JAIF does overlap in scope somewhat with *IEEE Transactions on Signal Processing* (T-SP) and *IEEE Transactions on Aerospace and Electronic Systems* (T-AES), as well as some other prestigious publications, it is unique in its focus on research and cutting-edge applications in high-level and low-level information fusion. As such, JAIF maintains a privileged position to publish top-tier work in these areas, while drawing quite naturally on expanded versions of the best FUSION papers presented each year. There are a mix of regular and dedicated, special-topic issues. To date, there have been five special issues, with more in early stages of planning for the coming years.

Dale Blair, who had considerable experience as EiC of T-AES, was the founding JAIF Editor-in-Chief (EiC). The founding of ISIF received strong encouragement and support from Yaakov Bar-Shalom, the ISIF VP Publications at that time, as well as from Bob Lynch, the first JAIF Administrative Editor. After eight years as EiC, Dale was followed by Uwe Hanebeck as EiC in 2014. Uwe remained in this position for six years and was then replaced by Stefano Coraluppi, the current EiC, in 2020. To this day, Dale remains active in his support to JAIF, serving in an oversight role as ISIF VP Publications. Another key contributor to the success of JAIF is David Krout, who replaced Bob as Administrative Editor in late 2015. The JAIF Editorial Board includes world-class experts in their respective fields. All are encouraged to explore available JAIF papers at https://isif.org/journals/all, and to consider submitting their best work to the journal.

PERSPECTIVES ON INFORMATION FUSION

The ISIF *Perspectives on Information Fusion* magazine was started to be a sister publication of JAIF to publish articles of general interest to the information fusion community. These include expository papers on new research areas, tutorials, classroom notes, book reviews, and announcements. Archival papers containing new research are still published in JAIF.

Publication of *Perspectives* was approved by the ISIF Board at FUSION 2014 with Roy Streit as founding Editor-In-Chief. While Roy, his area editors, and Dale Blair, VP-Publications, had experience with editing/publishing technical journals, none of them had started a magazine before. Roy very quickly discovered that this was not an easy task because he could not use the same processes and tools as JAIF [3]. The first issue of *Perspectives* was published in 2016 and distributed to FUSION 2016 attendees in Heidelberg. Including this issue, six issues of *Perspectives* have been published, with hard copies distributed at (physical) FUSION conferences. The current goal is to continue to publish one issue per year, to appear just before the annual FUSION conference. Anne-Laure Jousselme is the current Editor-in-Chief.

AWARDS

ISIF has established three society awards to recognize individuals for their contributions to information fusion. Nominations are solicited from ISIF membership, and the selection is made by the Awards Committee, currently chaired by Dale Blair.

The premier award is the ISIF Lifetime of Excellence in Information Fusion award, given to a researcher or engineer for outstanding contributions to the field of information fusion throughout his/her career. It was first given in 2015 and subsequently renamed in 2016 for the first recipient, Yaakov Bar-Shalom. Subsequent recipients are Chee-Yee Chong (2016), Pramod Varshney (2018), Ed Waltz (2021), and Roy Streit (2023).

The ISIF Young Investigator Award recognizes a young ISIF member for outstanding contributions to information fusion. The goal is to encourage individual efforts and foster increased participation by younger researchers and engineers. This award was established in 2016. The recipients to date are David Crouse (2016), Marcus Baum (2017), Karl Granstrom (2018), Florian Meyer (2021), and Florian Pfaff (2023).

The ISIF Robert Lynch Award for Distinguished Service recognizes an individual who has provided significant service to the Society. It was established in 2016 in memory of Robert (Bob) Lynch, who was involved in the organization of the annual FUSION conferences and co-chaired FUSION 2009 in Seattle. He was a key contributor in founding and production of JAIF and founding of the *Perspectives* magazine, and single-handedly maintained the ISIF web-site for many years. Chee-Yee Chong is the recipient of the 2023 ISIF Bob Lynch Distinguished Service Award.

In addition to the three Society awards, there are two awards for best papers at the FUSION conferences. The Jean-Pierre Le Cadre Award was established in 2010 in memory of Jean-Pierre Le Cadre for the best paper of the FUSION conference. The Tammy Blair Best Student Paper Award was established in memory of Tammy Blair, who passed away from the swine flu the week after FUSION 2009 in Seattle, for which she was Administrative Chair. The best paper awards are managed by the Awards Committee of each conference. The list of best paper and best student paper awards from FUSION 2004 to FUSION 2017 can be found in [4]. Best paper awards for subsequent years are published in *Perspectives*.

WORKING GROUPS AND OTHER MEETINGS

A lesser known ISIF activity to serve its members and advance the state-of-the-art in fusion is the sponsorship of working groups (WGs). ISIF sponsors working groups to bring together researchers who share a common interest in a technical area of information fusion, generally an emerging area that needs collaboration by a group of dedicated researchers to develop a theoretical framework or software tools. ISIF support includes providing a free meeting room during a FUSION conference and related website links and support for virtual meetings.

The Multistatic Tracking WG (MSTWG) was the first ISIF WG. Its objective was to promote collaboration among its members in multisensor fusion and tracking, with a focus on multistatic sonar and radar. From 2007 to 2016, when it was decommissioned, it held 17 regular meetings, five teleconferences, organized 11 special sessions at conferences, and analyzed seven common data sets.

There are currently two active working groups sponsored by ISIF. The first is the Evaluation of Techniques for Uncertainty Representation Working Group (ETURWG). The ETURWG has been meeting for over 10 years to refine, update, clarify, and implement the Uncertainty Representation and Reasoning Evaluation Framework (URREF) ontology. The working group activities include developing a URREF tutorial, incorporating artificial intelligence and machine learning (AI/ML), and defining metrics.

The other working group is the Open-Source Tracking and Estimation Workshop (OSTEWG), which supports Stone Soup in conjunction with a NATO Team Activity. Stone Soup has developed a software repository for state-of-the-art filtering/tracking algorithms and other algorithms, as well as user interfaces.

In addition to working groups, ISIF sponsors small meetings. Past meetings sponsored by ISIF include the:

- BELIEF international conference that addresses theoretical advances of belief functions, and promotes and expands the application fields of belief functions
- Canadian Tracking and Fusion Group (CTFG) workshop, and
- Sensor and Data Fusion (SDF) Symposium organized by Fraunhofer FKIE, Germany.

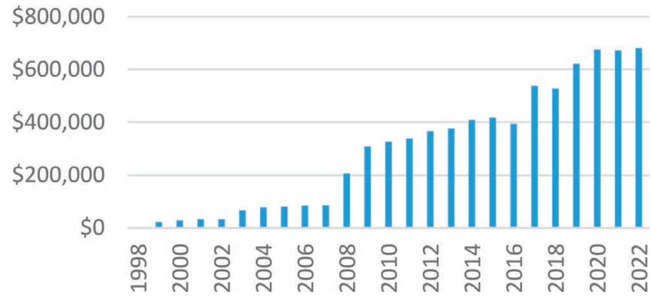

Figure 2
ISIF net worth (in USD) over time.

These sponsored events are smaller gatherings that connect many of the ISIF members outside of the annual conference.

MEMBERSHIP

FUSION conference attendees automatically become ISIF members with a part of their registration fee going to ISIF as membership dues. Thus, ISIF membership statistics closely track those of FUSION attendance, with a small deviation due to renewals by members who do not attend a conference, e.g., due to COVID-19. The geographical distribution of membership varies with the location of each conference. Membership in the host country generally surges the year FUSION is held and drops the following year since many conference attendees do not return to the next conference or renew their membership online. This phenomenon is especially prevalent with conferences held outside of North America or Europe. However, there is a core of members who attend FUSION conferences year after year. These are generally senior researchers in the field and include the Board of Directors. It is interesting to note that only two members, Yaakov Bar-Shalom and Chee-Yee Chong, have attended every FUSION conference since 1998.

Statistics about memberships are possible thanks first and foremost to Pierre Valin who built the first ISIF membership database. This database was maintained later by Elisa Shahbazian and subsequently by Anne-Laure Jousselme. The database allows ISIF to keep records of its membership since the very first conference.

To grow ISIF membership and reduce its dependence on FUSION attendance, we need to give people a reason a join ISIF besides conference attendance. One membership benefit we are exploring is an email address, member@isif.org, similar to member@ieee.org that IEEE provides for its members. Such email addresses are useful for members in industry or government where use of their employer email address for professional communication may be inappropriate.

FINANCES

ISIF was founded in 1998 with a private loan from Daniel Zhu, the entrepreneur who organized the first two FUSION conferences. The loan was repaid with the surplus from FUSION 1999. With the income from subsequent FUSION conferences that were mostly financially successful, and careful management of expenses, ISIF has accumulated a healthy reserve (Figure 2). This reserve enables ISIF to support its publications, working groups, and other meetings, and take some financial risk in sponsoring FUSION conferences outside the usual locations of Europe and North America.

Since ISIF has no other sources of income, its net worth fluctuates with the surplus from the FUSION conferences, with a lag that depends on when ISIF receives the surplus. The jumps in the net worth are due to the conferences with the big surpluses. In descending order, the top three are: 2016 (Heidelberg), 2018 (Cambridge), and 2009 (Seattle), all with surplus over USD 100K.

The ISIF Board is trying to figure out how much reserve is needed to prepare for the proverbial "rainy day". Without the surplus from the top three conferences, the ISIF reserve would be reduced by more than 50%. This can be very risky if we encounter another situation that affects conference attendance such as another pandemic. We are indeed fortunate that skillful financial management produced a surplus for each of the last three conferences despite attendance affected by COVID-19.

LOOKING FORWARD

The information fusion landscape has changed drastically since the founding of ISIF 25 years ago. Information fusion is now a crucial component in many applications, not just traditional defense and aerospace systems. As an example, sensor fusion is a prerequisite for driver assistance and autonomous driving. At the same time, the traditional model-based information fusion approach is being challenged by data-driven machine learning, which has taken the world by storm.

These changes present challenges to ISIF because each application or technology area has its own conferences and journals that compete with FUSION and JAIF. However, there are also opportunities because ISIF is not tied to a particular application or technology. As a community that focuses on the common core issues of information fusion across applications, and agnostic to a particular technical approach, ISIF is in a unique position to advance the state-of-the-art in information fusion. The future of ISIF is bright.

REFERENCES

1. Chong, C.-Y. Early history of International Society of Information Fusion. *Perspectives on Information Fusion*, Vol. 1, 1 (Jun. 2016), 37–42.
2. Chong, C.-Y., et al. 25 years of FUSION Conferences: collection of memories. *Perspectives on Information Fusion*, Vol. 5, 1 (Jul. 2022), 36–56.
3. Streit, R. Comments on the founding of *Perspectives*. *Perspectives on Information Fusion*, Vol. 6, 1 (Jun. 2023), 35.
4. Li, X. R. Best paper and best student paper awards, FUSION 2004–2017. *Perspectives on Information Fusion*, Vol. 2, 1 (Mar. 2019), 42–49.

COMMENTS ON THE FOUNDING OF *PERSPECTIVES*

In Salamanca, Dale Blair and I were having a casual conversation on topics of mutual interest while queuing for an event at FUSION 2014. At some point he volunteered that the Board had decided to start a magazine that would publish information of general interest to the ISIF community, but not archival research. I agreed that was a good idea. I innocently asked, "What kind of information?" Almost whatever you want, was Dale's reply. I walked straight into that one, eyes wide open. Dale's implied offer was serious, and so was his reply. It took a few minutes, but I agreed to serve as the Founding Editor in Chief. That role was confirmed at the Board meeting a few days later, together with several Associate Area Editors who agreed to help.

What I did not realize at the time was how much work is entailed in starting a new publication. The magazine had no history, no name, no publication rate, no departments, no reviewing standards, no layout format, etc. The only resource it had was a manuscript submission system in which every submission was a "refereed journal article," which meant that the editor's correspondence letters were unsuitable. (That last bit sounds silly, but I was appalled when the system sent the first submission, the obituary of a friend and colleague, for technical review.) It is necessary to design/configure/fix a great host of seemingly small things that we all take for granted. But I digress. The first thing I did was choose a name. I chose *Perspectives*, because of the diversity of interests in the ISIF community. I gathered a half dozen similarly purposed magazines and decided what I wanted a notional Table of Contents to look like. That is when I asked the AEs to join the conversation.

There is a missing piece, a missing role. Can you see it? That role was pivotal, but completely unrecognized by me, and hence it went unfilled for the first few issues. I didn't even know the name for this role, but now I do. It is the Production Manager (PM). Belatedly discovering the role and asking the Board to fill it changed everything. The first and current PM is Kristy Virostek, and without her *Perspectives* could not be a regularly published magazine.

There is much more that could be said. Maybe I will write more for the next issue, but I will wrap these notes by saying the role of Founding Editor in Chief of any publication is no small task. More than once I regretted saying yes, but – in truth – it was incredibly satisfying to see the very first issue of *Perspectives* in print. Would I do it all over again? Yes, absolutely, but I'd do things a little differently – I would start by finding the right PM.

—Roy Streit

EVOLUTION OF THE JDL MODEL

Alan N. Steinberg
Independent Consultant
Alexandria, VA, USA
alaneilsteinberg@gmail.com

Christopher L. Bowman
Data Fusion & Neural Networks, LLC
Arvada, CO, USA

Franklin E. White
FEWsionworks
San Diego, CA, USA
fewsionworks@icloud.com

Erik P. Blasch
MOVEJ Analytics
Fairborn, OH, USA
erik.blasch@gmail.com

James Llinas
University at Buffalo
Buffalo, NY, USA
llinas@buffalo.edu

Dedication—This paper is dedicated to the memory of Chris Bowman: brilliant colleague and close friend; a pioneer of the data fusion community in the United States and worldwide. Chris played a leading role in refining the JDL data fusion model that is the topic of this paper. He died unexpectedly before we could complete this article.

MODELS

The well-known Joint Directors of Laboratories (JDL) Data Fusion model has served as a paradigm for much of the subsequent discussion and development of data and information fusion.

The model was conceived in the late 1980's by the JDL Data Fusion Subgroup, consisting of prominent fusion experts and representatives from various US Government agencies [1], [2], [3]. The model was formulated as a scheme for clearly defining and differentiating concepts concerning the then-new field of data fusion. The model gained considerable influence by its articulation in Waltz and Llinas's landmark book, *Multisensor Data Fusion* [4].

Developments in the succeeding decades in applications and in applicable methods–in problem spaces and solution spaces–have strained the taxonomy, boundary assumptions, and partitioning scheme assumed in the early model. This has prompted numerous revisions and alternatives to the model.

Concepts and terms have been broadened to apply data fusion methods beyond the JDL's initial tactical military domain. Data fusion itself, initially defined as:

> *a process dealing with the association, correlation, and combination of data and information from single and multiple sources to achieve refined position and identity estimates, and complete and timely assessments of situations and threats, and their significance. The process is characterized by continuous refinements of its estimates and assessments, and the evaluation of the need for additional sources, or modification of the process itself, to achieve improved results [1]*

was defined more simply and comprehensively as:

> *the process of combing data to estimate or predict the state of some aspect of a world state [5].*

With the wisdom of age, we now prefer to define *Data Fusion* in even simpler and broader terms as:

> *the process of combining data to estimate entity states;*

where an entity can be any aspect of a universe of discourse at any degree of abstraction. To maximize breadth of applicability, we forgo distinctions of sensor fusion, data fusion, information fusion, knowledge fusion, etc.; considering "data fusion" as the encompassing term.

A data fusion process has the role of estimating entity states of interest within a problem domain on the basis of multiple data. As such, data fusion is a particular topic of epistemology: learning on the basis of multiple pieces of data. The specific data fusion problem is that of determining what data are relevant to a state estimation problem and using such data in deriving estimates; accounting for uncertainty in data relevance, data accuracy and in the performance of the inference method.

The JDL model introduced the notion of fusion "levels" as in Figure 1, distinguishing classes of fusion processing methods as applicable to major distinguishable classes of problems: processes that relate to the refinement of estimates or understanding of "objects" (Level 1), "situations" (Level 2), "threats"(Level 3), and "processes" (Level 4) [2], [3], [4].

The JDL model and its progeny have had to confront issues of the semantics of such terms. When the initial JDL model was used in for integrating across US Navy C4I and Combat systems, the issue arose as to the use of the terms "entity" and "object". Although commonly used interchangeably, the software community takes an entity to be a real world "thing" and an object to be a machine representation thereof (we'll see that clarification of usage in later revisions to the model). There's no space here either to describe common usage or to prescribe preferred usage. However, a fusion model will need an ontology and taxonomy to clarify such terms as:

- attribute//property//feature//signal//observable
- entity//object//individual//target

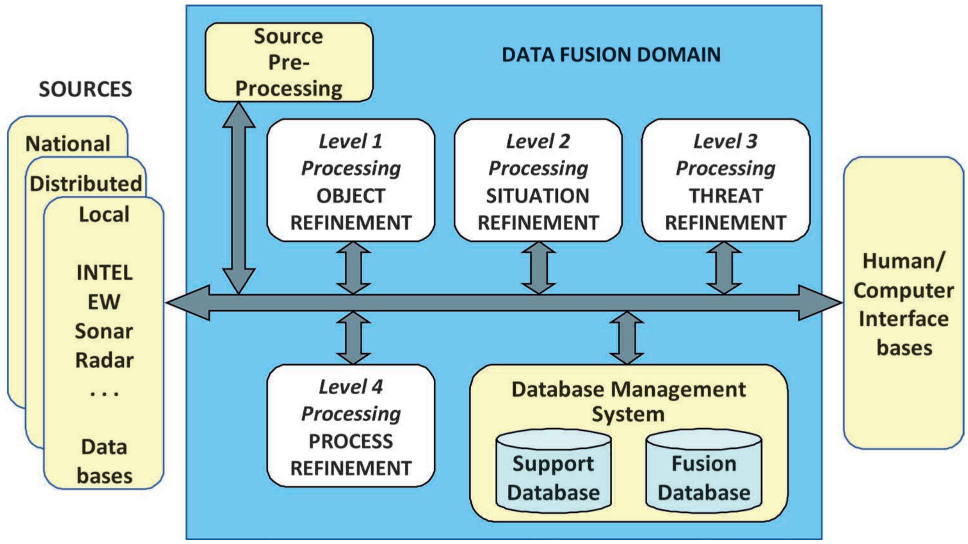

Figure 1
Early JDL data fusion model, 1990 [4].

- relation//relationship
- structure//complex//situation//scenario
- detection//contact//perceived entity//track (*vide* [7])

The original JDL model depicts levels as interacting via a bus architecture, such that processing sequences and access to data are free design variables. The prominence of 'threat" and the illustrative list of "sources" reveal the initial military focus. Later refinements generally had the purpose of broadening this perspective. After all, the very purpose of data fusion is inclusivity: to exploit *all* available information pertinent to a given problem. This motivates the broadest possible generalization and abstraction of problems and of solutions. It should be appreciated that significant research and development has occurred across the widest range of application, far beyond those of a military nature.

Challenges and refinements to the model are to be expected and welcomed to meet changing needs and perceptions. The model was not revealed to the JDL Data Fusion Subgroup on tablets from Mount Sinai. We made it up.

Ongoing developments in various technologies have obliged consideration of the relationship and role of data fusion in respect to new forms of knowledge representation and of uncertainty management, of data mining, cloud-based information retrieval, multimedia information exploitation, artificial intelligence and machine learning, joint human/machine problem-solving, etc.

A reexamination of the model was undertaken in the late '90s to clarify terms, broaden the model as much as possible from its initial focus on tactical military applications, refine the partitioning scheme, and explore relationships of data/information fusion with resource management, data mining, human situation awareness, and decision-making [5]. Source pre-processing was ennobled as Level 0 fusion to encompass data association and estimation at the feature/signal level (e.g., calibration, filtering, pulse train deinterleaving, modulation characterization). Level 4 fusion was divorced from resource management to clarify and exploit the distinction and duality of fusion/estimation vs management/control functions:

- *L0, Feature/Signal Assessment*: estimation of patterns: paradigmatically signal or feature modulations in 1, 2, or more dimensions; but can extend to most any abstract pattern: numeric or geometric patterns; musical or literary themes; rhyme schemes, etc.
- *L1, Individual Entity Assessment*: estimation of states of entities considered as individuals
- *L2, Situation Assessment*: estimation of relational states and of complexes of relationships
- *L3, Scenario/Impact Assessment*: predictive or forensic estimation of courses of action, scenarios, and outcomes
- *L4, System Assessment*: estimating states of the system itself: e.g., sensor and data alignment, estimation or control performance, fidelity of predictive models

Blasch has led the examination of several alternative approaches over the years [8–12]. As a recognition of the characteristic role that human cognition plays in understanding information, he and his colleagues introduced a DF Level 5, "User Refinement", similarly proposed by Hall and Mullen as "Cognitive Refinement" [13].

These ideas were incorporated in 2004–05 in a significant variant developed by the ISIF Data and Information Fusion Group (DIFG) [9]. As depicted in Figure 2, the DIFG model distinguishes fusion levels as transforming information between entities of various types. It effectively partitions fusion processes on the basis of *agency*, in terms of classes of entities providing and receiving the data. This is an *information*

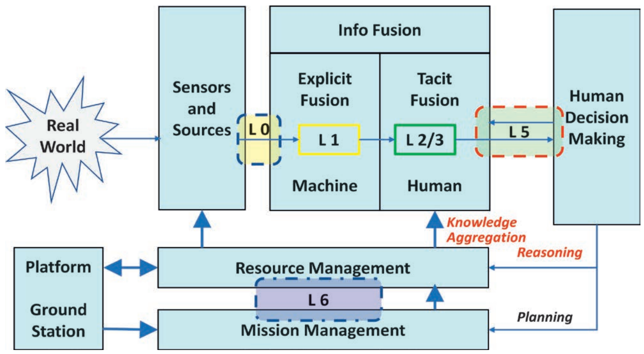

Figure 2
DIFG model, 2005, redrawn from [9].

exploitation model in that it includes planning and control at its Levels 4 and 5.[1]

Other model variants [11], [12], [14] similarly distinguish high-level from low-level information fusion (HLIF vs. LLIF), both in terms of types of processes and in types of products:

> *The low-level functional processes support target classification, identification, and tracking, while high-level functional processes support situation, impact, and fusion process assessment. LLIF concerns numerical data (e.g., target locations, kinematics, and attribute types). HLIF concerns abstract symbolic information (e.g., threat, intent, and goals) [12].*

This seems an unnecessarily constraining and perhaps forced marriage. Symbolic methods are certainly applicable to "low level" target classification, numerical methods (e.g., belief networks) to relational and situational assessment and to process assessment. Recognition and prediction of relational, situational, and system states are clearly akin to low-level individual state recognition and prediction. Similar classification, characterization, and tracking methods may apply. HLIF can provide context for predicting and understanding LLIF state and HLIF states can provide context for one another.

BUT WHICH MODEL IS RIGHT?

How then, to select among the multitude of JDL model variants and alternatives? These models tend to differ either:

- in scope: do they include control as well as estimation processes? Do they encompass human as well as machine techniques?
- in partitioning scheme: are elements differentiated by type of input, processes, output, or agencies (i.e., who or what does the fusing)? or
- in purpose: is it an ontological, epistemic, management, or engineering model?

Many fusion models, including various versions of the JDL model, are based on one or another of these distinctions, and sometimes straddle the distinctions.

We need to be clear as to the reason for having a data fusion model. To the extent that it is meant to support system design and evaluation, a data fusion model is a *management model* and, specifically, an *engineering model*. As such, we would like it to partition the problem space in a way that tends to support different types of solutions. For example, the stated objectives of [5] were (a) to provide a useful categorization representing logically different types of problems, which are commonly solved by different techniques; and (b) to maintain a degree of consistency with the mainstream of technical usage.

Let us propose three desired qualities for engineering models, to include data fusion models:

- ***Avoid Confusion*** with a clear distinction of problems that tend to require different solution methods
- ***Constrain Profusion*** of models and methods by generalizing concepts and constructs so as to apply across a wide range of problem domains, facilitating integration, technology re-use and deeper understanding
- ***Mitigate Diffusion*** of communities of practice by clearly defining the relationships of the modeled domain to other domains, promoting coordination and synergy. As in international politics, a data fusion model shouldn't erect borders that impede the useful flow of goods and services, either internally between fusion levels or with neighboring domains: planning, data mining, machine learning, etc.

In short, practitioners desire clear and comprehensive modeling of fusion problems, solutions, and problem domains. Internal and external synergy is facilitated by a common representational framework across fusion functions and with neighboring disciplines, as discussed in [19], and a comprehensive functional architecture [18], [20].

Although there have been many revisions and rivals to the JDL model, nearly all of them partition the fusion domain in terms of fusion "levels". The partitioning criteria in the early versions of the JDL model were easily blurred: do we differentiate "levels" based on types of input, types of processes, or types of outputs? None of these criteria is necessarily right or wrong but they may serve different needs.

In [5], [6], [15–18] we successively proposed refinements to the early definition of levels (e.g., Figure 3). The explicit goals were to clarify the partitioning and to broaden applicability beyond the original tactical military domain. We suggested partitioning levels according to fusion products: specifically,

[1] The DIFG model's distinction between "Explicit Fusion", performed by machines, from "Tacit Fusion", performed by humans, is a bit anthropocentric. There's no fundamental reason why machines can't perform higher-level (L2/3) fusion or people–or animals for that matter–can't perform lower level (L1) fusion. Perhaps we should view the reference to "human" agency as an exemplar for internal or external processes by sentient beings. Developments in AI, not to mention SF, blur that distinction. Also, the boxes labeled "Platform" and "Ground Station" in the figure can be viewed merely as examples of model instantiation.

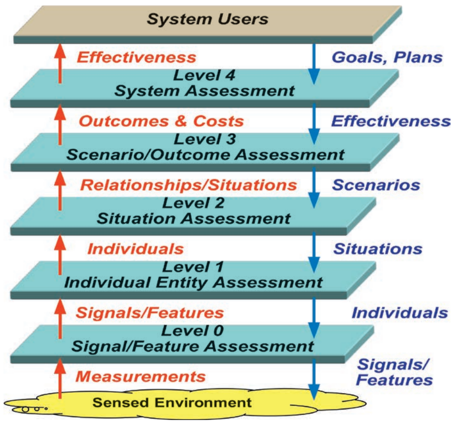

Figure 3
1999 revision [5].

the types of state variables to be estimated. In this way states of interest can be distinguished in terms roughly corresponding to the levels described in earlier versions of the JDL model.[2]

We also extended the formal and functional duality between data fusion and resource management functions by defining a set of corresponding management levels [17–20]. This extension helps clarify the role of data fusion within the broader field of information exploitation. Bowman applied this broader purview to formulate a dual-node network architecture, comprised of paired data fusion/management nodes, each pair acting as a quasi-autonomous agent that acquires and processes data to meet its evolving objectives in the system context [20].

These DF and RM levels map into a categorization of entity state variables which a DF system is tasked to estimate or which an RM system is tasked to control.

But how do the traditional fusion "levels" fare given this insight? Levels 0 through 2 clearly can be distinguished by types of variables: signal/feature parameters vs. individual metric and kinematic variables vs. relational variables.

Level 3 fusion estimates or predicts courses of action, events, and impacts. As these generally concern projected entity states and relationships, many versions of fusion models refer to a blended "Level 2/3" (as in Figure 2). We can broaden the earlier label "Impact" to "Outcome", which may include impacts on various entities, including on "our" system and mission.

As for Level 4, we have indicated the importance of differentiating estimation from control and, therefore, fusion from management [20]. *System Assessment* is therefore preferable as a fusion level to the original model's *Process Refinement*. However, L4 fusion is still an awkward fit. The distinction of L4 from other fusion levels is more a matter of ownership than of type of process or product. In L4, a system assesses its own signal/feature parameters, individual metrics and kinematics, and relationships; i.e., its own L0–3 variables. However, because system boundaries and ownership can be partial, mutable, and uncertain, so can the distinction between system assessment and assessment of external variables. Therefore, L4 challenges our preference for clear boundaries and partitioning criteria.

The original fusion Level 4, Process Refinement, as well as the proposed Levels 5 and 6, relates to the resource management side of data exploitation. Indeed, the original JDL documentation addressed these concerns as Level 4 machine control, Level 5 user control, and Level 6 control of the data collection and processing. These "levels" reflect the multi-dimensionality of data exploitation.

Even within the single dimension that distinguishes data fusion Levels 0 through 3, the term "levels" can be misleading. The sequential numbering of levels (or depictions as in Figure 3) should not be construed as a constraint. Fusion/management processes must be free to employ data types and sources within or across levels as needed [20].[3]

Figure 4 presents the original, non-hierarchical JDL model, refined and extended to improve clarity and breadth in modeling fusion problems, solutions, and problem domains:

a. clearer partitioning scheme, based on classes of variables to be evaluated or managed (*to avoid Confusion*)

b. generalization of concepts to extend to all applications, both in leveling the levels in a bus configuration and expanding their scope (*to constrain Profusion*)

c. expanding the model to include resource management to encompass all aspects of information exploitation (*to mitigate Diffusion*)

The notional agent bus architecture–as in the original model of Figure 1–allows data to flow unconstrained by the model within and among the data fusion and resource management levels, enabling flexible, opportunistic data exploitation and response. System users external to the fusion processes are shown, with the proviso that people can perform any of the functions internal to data fusion.

SUMMARY

The JDL model was developed to define the concepts and structure of the data fusion problem: that of estimating entity states of interest within a problem domain. The model has been refined over the years (a) to extend as broadly as possible across diverse problem domains to facilitate common solutions to common problems and (b) to recognize synergies with other disciplines related to information understanding and information exploitation. A testament to the contribution of the model has been in the wide use of its structure and taxonomy not only by researchers and practitioners, but also in data fusion product specifications for acquisition and deployment.

[2] As argued in [17], [18], [19], generality is improved by partitioning inference problems on the basis of types of entity state variables rather than by type of entity. A given entity–say, an aircraft–can be addressed at more than one level: as an individual (at Level 1) or as a complex (Level 2) such that the relationships among its components or subassemblies are being estimated. The aircraft may also be addressed at Level 3 as a dynamic process; or at Level 4, if it happens to be the system performing the estimation.

[3] Machine learning methods may operate across non-adjacent levels, inferring situations directly from measurements. Conversely, states of individuals may be inferred from situations or courses of action.

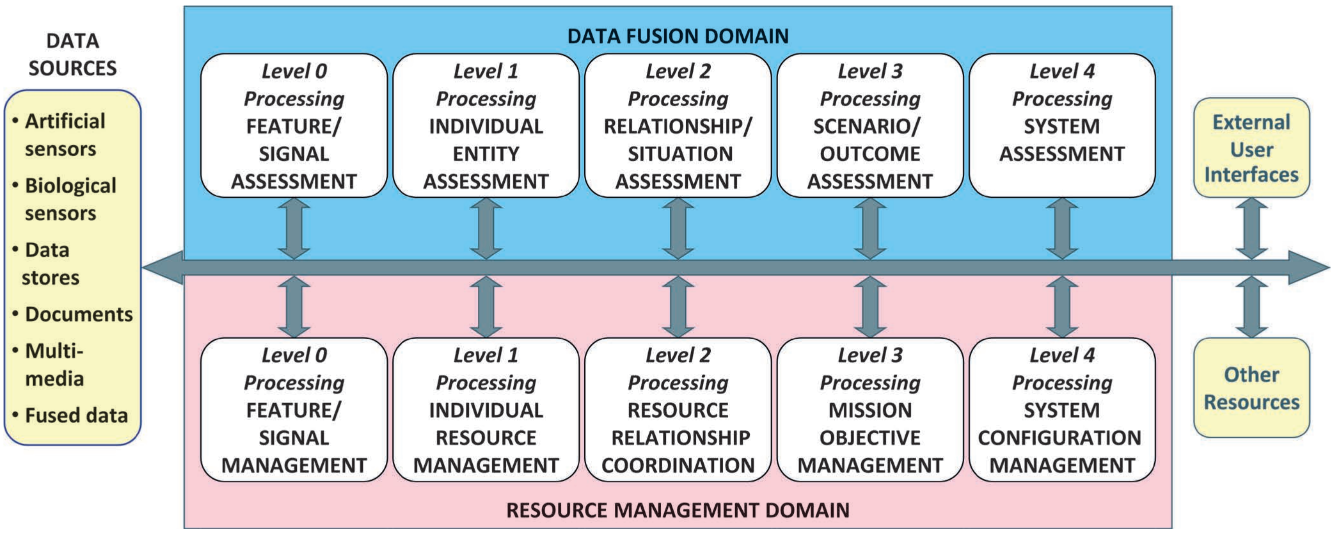

Figure 4
Candidate data fusion process model [18].

REFERENCES

1. White, F. E. Data fusion lexicon. Joint Directors of Laboratories, San Diego, CA, Technical Panel for C3, Data Fusion Sub-Panel, Naval Ocean Systems Center, 1987.
2. White, F. E. A model for data fusion. In *Proceedings First National Symposium on Sensor Fusion*, Chicago, IL, 1988.
3. Kessler, O. Functional description of the data fusion process. Technical report for the Office of Naval Technology Data Fusion Development Strategy, Warminster, PA, NADC, Nov. 1991.
4. Waltz, E., and Llinas, J. *Multisensor Data Fusion*. Boston, MA: Artech House, 1990.
5. Steinberg, A. N., Bowman, C. L., and White, F. E. Revisions to the JDL model. In *Proceedings Joint NATO/IRIS Conference*, Quebec, Canada, Oct. 1998; reprinted in *Sensor Fusion: Architectures, Algorithms and Applications, Proceedings SPIE*, Vol. 3719, (1999), 430–441.
6. Steinberg, A.N., and Bowman, C.L. Rethinking the JDL data fusion model. In *Proceedings MSS National Symposium on Sensor and Data Fusion*, Laurel, MD, June 2004.
7. Steinberg, A. N. An ontology for multilevel data fusion. In *Proceedings International Conference on Information Fusion*, Linköping, Sweden, 2022.
8. Blasch, E., and Plano, S. JDL Level 5 Fusion model "user refinement" issues and applications in group tracking. *Proceedings SPIE,* Vol. 4729, (2002), 270–279.
9. Blasch, E., and Plano, S. DFIG Level 5 (User Refinement) issues supporting situational assessment reasoning. In *Proceedings International Conference on Information Fusion*, Philadelphia, PA, 2005.
10. Blasch, E., Kadar, I., Hintz, K., Biermann, J., Chong, C., and Das, S. Resource management coordination with level 2/3 fusion issues and challenges. *IEEE AES Magazine*, Vol. 23, 3 (2008) 32–46.
11. Blasch, E., Lambert, D. A., Valin, P., Kokar, M. M., Llinas, J., Das, S., Chong, C.Y., and Shahbazian, E. High level information fusion (HLIF): survey of models, issues, and grand challenges. *IEEE Aerospace and Electronic Systems Magazine*, Vol. 27, 9 (Sept. 2012), 4–20.
12. Blasch, E., Steinberg, A. N., Das, S., Llinas, J., Chong, C-Y., Kessler, O., Waltz. E., and White, F. Revisiting the JDL model for information exploitation. In *Proceedings International Conference on Information Fusion*, Istanbul, Turkey, 2013.
13. Hall, D. L., and Mullen, S.,A.,H. *Mathematical Techniques in Multisensor Data Fusion* (2nd ed.). Boston, MA: Artech House, 2004.
14. Lambert, D. A. A unification of sensor and higher-level fusion. In *Proceedings International Conference on Information Fusion*, Florence, Italy, 2006.
15. Steinberg, A. N., and Bowman, C. L. Revisions to the JDL data fusion model. In *Handbook of Multisensor Data Fusion*, M. Liggins, D. Hall, J. Llinas, Eds. London, England: CRC Press, 2009, ch. 3.
16. Llinas, J., Bowman, C., Rogova, G., Steinberg, A. N., Waltz, E., and White, F. E. Revisiting the JDL data fusion model II. In *Proceedings International Conference on Information Fusion*, Stockholm, Sweden, 2004.
17. Steinberg, A. N. Foundations of situation and threat assessment. In *Handbook of Multisensor Data Fusion*, M. Liggins, D. Hall, J. Llinas, Eds. London, England: CRC Press, 2009, ch. 18.
18. Steinberg, A., and Snidaro, L. Levels? In *Proceedings International Conference on Information Fusion*, Washington, DC, 2015.
19. Steinberg, A. N. Situations and contexts. *Perspectives Magazine*, Vol. 1, 1 (Jun. 2016), 16–24.
20. Bowman, C. L. The dual node network (DNN) data fusion & resource management architecture.In *Proceedings AIAA Intelligent Systems Conference*, Chicago, IL, Sept. 20–22, 2004.

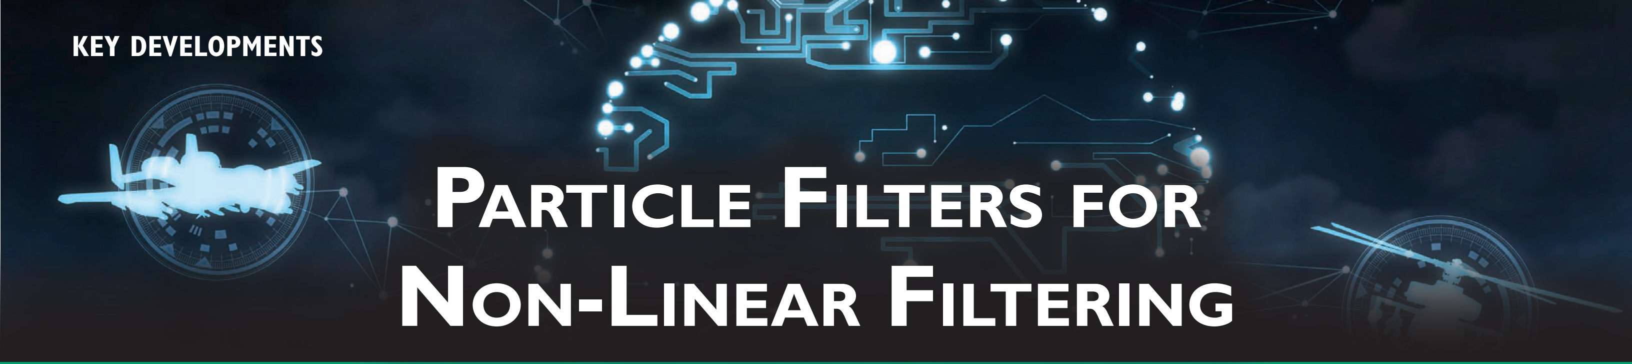

Particle Filters for Non-Linear Filtering

Simon Maskell
University of Liverpool
Liverpool, UK
smaskell@liverpool.ac.uk

PROBLEMS

We often encounter problems where we need to process an incoming stream of data to make inferences about a time-evolving quantity of interest. Examples of such problems emerge in a diverse range of applications spanning tracking, GPS-free navigation, robotics, epidemiology, and finance.

To make statistical inferences in such contexts, we capitalise on models that capture our understanding of both the time-evolution of the quantities of interest and the relationship between these quantities and the observed data. Were these models to be linear and Gaussian, the uncertainty can be exactly characterised using a sequence of analytic calculations. Unfortunately, the models are rarely linear or Gaussian. One widespread strategy is to approximate the models as being (locally) linear and Gaussian and/or to decompose the problem into a set of sub-problems each of which involves linear and Gaussian models: the result are approaches typified by the extended and unscented Kalman filters, multi-hypothesis tracker, and the interacting multiple model. Given that it is the models that capture our understanding, approximating the models necessarily compromises our ability to use that understanding to inform the inferences that we make. Particle filters adopt a different approach whereby we explicitly approximate the result of the inference but fully capitalise on the model fidelity when we do so. Particle filters achieve this via Sequential Monte Carlo, i.e., at each time step they characterise the uncertainty using a set of (weighted) sampled values for the inferred quantities of interest.

When the models concerned are well approximated as linear and Gaussian, particle filters may offer some benefit relative to alternative techniques, but this benefit is often (sensibly) argued to not be warranted by the computational expense required to propagate the (potentially large number of) samples over time. Perhaps as a result, while particle filters were initially hoped to offer improved performance in contexts where existing filters were struggling (typified by bearing-only tracking), they arguably failed to make significant gains in these applications: it transpires that the problems were limiting performance in these contexts, not the filters.

However, there are many important problems where the models are not well approximated as linear or Gaussian. It is these contexts where particle filters have shone as a result of their ability to solve problems that other approaches simply cannot tackle. Examples are diverse and range from GPS-free navigation [1] to localising earthquakes using data extracted from social media [2].

PARALLELISM

A fundamental strategy when developing faster processors is to make the processors smaller. However, when processors switch state, they generate heat and it becomes increasingly challenging to dissipate this heat as the devices shrink in size. The result is that single processors have failed to deliver on Moore's law since approximately 2013. Since then, increases in processing power (typified by the GPUs being used for deep learning) have been achieved by maximising the number of processors on a chip. If algorithms are to exploit such hardware, parallelism is a necessity.

When a particle filter processes each datum, it propagates each particle and calculates each particle's weight. These operations are independent from one particle to another such that particle filters are often claimed to be readily parallelisable. However, there is an issue with this claim. As a particle filter iterates through time, a weight is recursively updated for each particle: this weight is the extent to which the particle will contribute to any inference. It is inevitable that the weights for different particles will come to differ significantly and inferences will become dominated by a small subset of the particles. To address this wastefulness, particle filters employ a 'resampling' step. This involves removing the particles with low weights and replacing them with replications of the particles with high weights. It is the introduction of resampling that gave rise to the first working particle filter [3]. However, this same resampling step is non-trivial to parallelise. This has motivated research into approaches to both modifying the resampling step to make it amenable to parallel implementation [4] and approaches to defining a parallel implementation without such modifications [5].

PROPOSAL DISTRIBUTIONS

However much we exploit parallelism, we will be limited by how efficiently each particle is processed and, more specifically, the "proposal distribution", how the state associated with each particle is proposed. Research has focused on how to design efficient proposals. One exemplar such approach is the use of Kalman filter techniques inside the proposal [6]. Another closely related approach, particle flow [7], involves defining a process that takes the place of the proposal, and is highly reminiscent of numerical approaches (e.g., Hamiltonian Monte Carlo) that have proven both effective and popular in the context of other numerical Bayesian algorithms.

It transpires that one can achieve further improvements by proposing refinements to historic samples retrospectively in the light of recent data. Note that such an a fixed-lag approach [8] is also known as "blocking" in the context of particle filters.

PARAMETER ESTIMATION

As already explained, particle filters capitalise on models to make their inferences. Developing models and fine-tuning their parameters is time consuming. Techniques have been developed to learn such parameters from data: see, for example, [9].

As well as offering the potential to use high fidelity models, this capacity to learn models' parameters from data also makes it possible to apply particle filtering with the models used by deep learning algorithms such as long short-term memory (LSTM) networks and transformers. In this context, particle filters can be seen as a machine learning approach that enables users to understand the uncertainty associated with such deep learning approaches.

PROSPECTS

It is currently challenging for an applied researcher to capitalise on the advances to particle filtering that have happened in the last 25 years and are exemplified above. This contrasts with neighbouring domains where probabilistic programming languages (PPLS) such as Carptenter et al. [10] have made it straightforward to define and then use probabilistic models to make inferences from (fixed) data. One exciting avenue for future research into particle filters is to extend such PPLS to make it straightforward to perform parameter estimation and use parallel implementations of particle filters with high-performance proposal distributions.

REFERENCES

1. Gustafsson, F., Gunnarsson, F., Bergman, N., Forssell, U., Jansson, J., Karlsson, R., et al. Particle filters for positioning, navigation, and tracking. *IEEE Transactions on Signal Processing*, Vol. 50, 2 (2002), 425–437.
2. Sakaki, T., Okazaki, M., and Matsuo, Y. Earthquake shakes Twitter users: real-time event detection by social sensors. In *Proceedings of the 19th International Conference on World Wide Web*, Raleigh, NC, 2010, 851–860.
3. Gordon, N. J., Salmond, D. J., and Smith, A. F. M. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEEE Proceedings F (Radar and Signal Processing)*, Vol. 140, 2 (1993), 107–113.
4. Bolic, M., Djuric, P. M., and Hong, S. Resampling algorithms and architectures for distributed particle filters. *IEEE Transactions on Signal Processing*, Vol. 53, 7 (2005), 2442–2450.
5. Varsi, A., Maskell, S., and Spirakis, P. G. An O(log2n) fully-balanced resampling algorithm for particle filters on distributed memory architectures. *Algorithms*, Vol. 14, 12 (2021), 342.
6. Van Der Merwe, R., Doucet, A., De Freitas, N., and Wan, E. The unscented particle filter. *Advances in Neural Information Processing Systems*, Vol. 13, 2000.
7. Daum, F., and Huang, J. Particle flow for non-linear filters with log-homotopy. *SPIE Signal and Data Processing of Small Targets 2008*, Vol. 6969, (2008), 414–425.
8. Doucet, A., and Senecal, S. Fixed-lag sequential monte carlo. In *Proceedings 2004 12th IEEE European Signal Processing Conference*, Vienna, Austria, 2004, 861–864.
9. Chopin, N., Jacob, P. E., and Papaspiliopoulos, O. Smc2: an efficient algorithm for sequential analysis of state space models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, Vol. 75, 3 (2013), 397–426.
10. Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., et al. Stan: A probabilistic programming language. *Journal of Statistical Software*, Vol. 76, 1 (2017), 1–32.

IMM Estimator in Multisensor Multitarget Tracking for Air Traffic Control and Autonomous Driving

Henk A. P. Blom
Delft University of Technology
Delft, Netherlands
h.a.p.blom@tudelft.nl

Kaipei Yang
Yaakov Bar-Shalom
University of Connecticut
Storrs, Connecticut, USA
kaipei.y@gmail.com
yaakov.barshalom@uconn.edu

Many dynamical systems undergo switches in their dynamical configuration, shortly referred to as mode switching. For example, an observed aircraft or car switches from uniform motion to a maneuver mode, or switches back from a maneuver mode to uniform motion. In nonlinear filtering, the simplest model of this type is a Markov jump linear Gaussian process x_t satisfying:

$$x_t = A(\theta_t)x_{t-1} + B(\theta_t)w_t \tag{1}$$

where θ_t is a hidden Markov chain, that switches per time step with probability Π_{ij} from i to j in the set of models $\{1, \ldots, N\}$.

The problem is to estimate x_t from the noisy observation:

$$y_t = F(\theta_t)x_t + G(\theta_t)v_t \tag{2}$$

where y_t is the R^m-valued observation of the R^n-valued system state x_t. Matrices A, B, F, and G depend on θ_t, and w_t and v_t are independent white Gaussian noises.

The optimal non-linear estimator involves a number of Kalman filters that increases exponentially with time t. The interacting multiple model (IMM) estimator [1], [2] involves N Kalman filters only, one for each possible mode. To compensate for the reduction in number of filters, at the start of each estimation cycle, there is a controlled interaction/mixing between the estimates from the N Kalman filters. [1] has formally proven that these interaction/mixing equations are exact, not an approximation. At the end of each estimation cycle, the IMM estimator calculates the filter weights (mode probabilities), as well as the overall mean and covariance. Bar-Shalom et al. [3] give an in-depth explanation of the IMM estimator and its application in tracking and navigation. Kalman filters for kinematic models [3] are low-pass filters. With small noise gain B in (1) they have a low bandwidth, suitable for nearly constant velocity motion. With large B, they have a higher bandwidth and are suitable for maneuvering targets. The IMM with such models is an adaptive bandwidth estimator.

In case of no mode switching, Π is the identity matrix and IMM reduces to the well-known MM estimator. As explained by [4], the success of the IMM estimator can be attributed to its simplicity in extending the MM estimator with the exact interaction equations at the beginning of each estimation cycle, which makes IMM the natural approximation of the optimal estimator for mode-switching systems.

IMM IN AIR TRAFFIC CONTROL

In air traffic control (ATC), multisensor multitarget tracking (MTT) is a basic functionality in fusing observation data reports from various sensors into a reliable and accurate real-time air traffic situation. One of the problems to be handled by MTT is to track a sudden maneuver start and stop for aircraft. Additional problems are that sensor reports may include outlier and false measurements, both of which can be mistaken for a maneuver. Another problem is that a data report typically does not include the identity of the aircraft source or may include an erroneous identity.

In the eighties, the first author had the opportunity to investigate this problem at Netherlands Aerospace Laboratory NLR. The novel approach was to study the problem within the theory of nonlinear filtering of a jump-diffusion that evolves in a hybrid state space. This resulted in a characterization of IMM's interaction in a continuous time setting [5]. Subsequently, this interaction was developed for a discrete-time version of the IMM estimator [1]. Initially, at NLR, this research was judged to be esoteric, rather than of practical use. This view completely changed when for an IMM based track maintenance algorithm remarkably good performance was demonstrated on simulated and live data from primary radar observations of air traffic [6]. The modes of flight modelled in this research are uniform motion, speed change, right turn and left turn, while outliers, false measurements and missing identities were covered by probabilistic data association (PDA). In the follow-on phase, the research was widened to the development of a Bayesian MMT design for ATC [7].

These remarkable tracking results motivated EUROCONTROL to start the development of its multisensor multitarget tracking system ARTAS. The first ARTAS version fused data reports from multiple primary and secondary radars, and its tracking architecture was largely based on the IMM inspired design of [7]. Halfway through the nineties, ARTAS started its ATC operational use in a steadily increasing number of EUROCONTROL member states. The use of ARTAS by a steadily increasing number of ATC centres has also stimulated further development. One important development is the replacement of PDA by a joint PDA approach that avoids coalescence of neighbouring tracks [8]. Other important developments concern the fusion of data reports from new sensor types, such as Mode

S, automatic dependent surveillance (ADS), wide area multilateration (WAM), and surface movement radar (SMR) [9].

Today, ARTAS is operational for ATCs in 43 member states of EUROCONTROL, as well as in several other states, including the USA. In parallel, its further development is ongoing, such as fusing new sensor types for the tracking of an increasing number of drones.

IMM IN AUTONOMOUS DRIVING

An advanced driver assistance system (ADAS) or autonomous driving (AD) system must be capable of estimating 1) the ego vehicle's motion, orientation, behavior, and trajectory, as well as 2) the perception of surrounding objects such as other vehicles, bicycles, and pedestrians, to ensure the safety and efficiency of autonomous vehicles.

In an autonomous driving system, different sensors and sources of information provide different types of data, such as LiDAR [10], radar, cameras, GPS, inertial measurement unit, and so on. Each sensor has its strengths and weaknesses, and none of them alone can provide a complete picture of the vehicle's environment. For example, LiDAR can provide high-resolution three-dimensional point cloud data, but it can be affected by weather conditions such as rain and snow. On the other hand, radar can penetrate some weather conditions but provides lower-resolution data. By using IMM, the autonomous driving system can combine data from different sources and sensors, and rely on multiple models of the vehicle's environment to make more accurate and reliable decisions in real time. Each model is designed to capture a specific aspect of the environment, such as object detection, motion estimation (which is subject to different behavior modes), or localization. These multiple models are then used to generate a more accurate and reliable estimate of the vehicle's surroundings [11].

In a variety of autonomous driving scenarios, IMM estimation has demonstrated significant efficiency, robustness, and reliability in integrating onboard vehicle sensors in multiobject tracking (MOT) and vehicle localization. These are used for applications such as estimating road conditions and predicting drivers' turning intentions at urban intersections, i.e., can handle different behavior modes [12]. Compared to single-model-based tracking, IMM has been shown in practice to improve the accuracy of motion estimation and overall, MOT performance with less track segmentation, less object ID switching, and higher recall.

Being a model-based approach that incorporates prior knowledge, IMM fills a gap between autonomous driving and data-driven algorithms because the latter solely relies on patterns in data and may not be able to capture the full range of driving scenarios. Optimal performance of autonomous driving can be achieved by using a combination of model-based algorithms and data-driven approaches, with IMM delivering robust and reliable tracking results and machine learning and neural networks capturing more subtle patterns in the sensor data and providing additional insights. Overall, the IMM estimator will continue to be critical in the advancement of autonomous driving technology.

LOOKING AHEAD

In this short paper, IMM applications in multisensor multitarget tracking have been highlighted for air traffic and for road traffic. From these highlights it has become clear that by their objectives, these applications involve very large sets of live data streams. From this perspective, MTT has been decades ahead of the current era of large data research. This also means that the results obtained from MTT research can provide novel insight in large data research. To speak in IMM terms, this defines great opportunities for Interaction between research in Bayesian estimation and in large data.

REFERENCES

1. Blom, H. A. P. An efficient filter for abruptly changing systems. In *Proceedings 23rd IEEE Conference on Decision and Control*, Las Vegas, NV, 1984, 656–658.
2. Blom, H. A. P., and Bar-Shalom, Y. The interacting multiple model algorithm for systems with markovian switching coefficients. *IEEE Transactions on Automatic Control*, Vol. 33, 8 (1988), 780–783.
3. Bar-Shalom, Y., Li, X. R., and Kirubarajan, T. *Estimation with Applications to Tracking and Navigation*. New York, NT: Wiley, 2001.
4. Bar-Shalom, Y., Challa, S., and Blom, H. A. P. The IMM estimator as the natural approximation of the optimal estimator for hybrid systems. *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 41, 3 (2005), 986–991.
5. Blom, H. A. P. Markov jump-diffusion models and decision-making-free filtering. In *Proceedings 6th International Conference Analysis and Optimization of Systems*, Berlin, Germany, 1984, 568–580.
6. Blom, H. A. P. A sophisticated tracking algorithm for Air Traffic Control surveillance radar data. In *Proceedings 1984 International Conference on Radar*, Paris, France, 1984, 393–398.
7. Blom, H. A. P., Hogendoorn, R. A., and van Doorn, B. A. Design of a multisensor tracking system for advanced Air Traffic Control. In *Multitarget-Multisensor Tracking: Applications and Advances*, Vol. II, Y. Bar-Shalom, Ed. Boston, MA: Artech House, 1992, pp. 31–63.
8. Blom, H. A. P., and Bloem, E. A. Probabilistic data association avoiding track coalescence. *IEEE Transactions on Automatic Control*, Vol. 45, 2 (2000), 247–259.
9. EUROCONTROL. ARTAS service movement surveillance, https://www.eurocontrol.int/sites/default/files/publication/files/artas-article-paper-update-08032017.pdf, last access 10 May 2023.
10. Li, Y., and Ibanez-Guzman, J. Lidar for autonomous driving: the principles, challenges, and trends for automotive lidar and perception systems. *IEEE Signal Processing Magazine*, Vol. 37, (2020), 50–61.
11. Yuan, T., Krishnan, K., Duraisamy, B., Maile, M., and Schwarz, T. Extended object tracking using IMM approach for a real-world vehicle sensor fusion system. In *Proceedings 2017 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, Daegu, South Korea, 2017, 638–643.
12. Lefkopoulos, V., Menner, M., Domahidi, A., and Zeilinger, M. N. Interaction-aware motion prediction for autonomous driving: a multiple model Kalman filtering scheme. *IEEE Robotics and Automation Letters*, Vol. 6, 1 (2021), 80–87.

Distributed Estimation

Chee-Yee Chong
Independent Researcher
Los Altos, CA, USA
cychong@ieee.org

Felix Govaers
Wolfgang Koch
Fraunhofer FKIE
Wachtberg, Germany
felix.govaers@fkie.fraunhofer.de
w.koch@ieee.org

Distributed estimation addresses the problem of combining local estimates that are based on measurements of individual sensors. This setup is particularly useful in spatially distributed applications, since data transmitted over potentially low-bandwidth communication links is reduced and the computational burden is shared among multiple nodes. Since local estimates have errors that are correlated over time for the same sensor, and across sensors at the same time, most distributed estimation research has focused on how to address this correlation [1], either by removing double counting to reconstruct the centralized estimate, or using the correlation to find the best linear estimate given the local estimates.

FUSION BY DECORRELATION

Early research on distributed estimation aims at reconstructing the centralized Kalman filter estimate $\hat{x}$ and error covariance $\hat{P}$ from local Kalman filter estimates $\hat{x}_i$ with covariances $\hat{P}_i$, *for i = 1,..,S.* The fusion equation [2] is:

$$\hat{P}^{-1}\hat{x}^{-1} = \bar{P}^{-1}\bar{x}^{-1} + \sum_{i}^{S}\left(\hat{P}_i^{-1}\hat{x}_i^{-1} - \bar{P}_i^{-1}\bar{x}_i^{-1}\right) \qquad (1)$$

where $\bar{x}$, $\bar{P}$, $\bar{x}_i$, $\bar{P}_i$ are the global and local predictions. There is a similar equation for the error covariances.

Eq. (1) is sometimes called information matrix fusion because it is based on the information matrix form of the Kalman filter. It is also known as tracklet or equivalent measurement fusion [3] because the summand in (1) represents the new information in the measurements in the local estimates between fusion times. Because of its intuitive form and simple implementation, it is widely used in track fusion, and shown to have good performance even when the underlying assumptions of zero process noise or full-rate communication are violated.

Since a set of measurements is conditionally independent given the states at multiple observation times, optimal state estimate fusion can be achieved by using an augmented state that consists of the states at multiple times. The fusion equations have the same form as (1) except that x is the augmented state. Good estimation performance can be obtained using a small number of augmented states but the main benefit of using augmented state estimates is in track association [4].

FUSION OF PSEUDO ESTIMATES

Fusion of local estimates to compute the exact Kalman filter estimate in the presence of process noise and non-full-rate communication had been a challenging research problem for many years. This solution to this problem was published as the distributed Kalman filter (DKF) [5]:

$$\hat{P}^{-1}\hat{x}^{-1} = \sum_{i}^{S}\check{P}_i^{-1}\check{x}_i \qquad (2)$$

where $\check{x}_i$ and error covariance $\check{P}_i$ are pseudo estimates different from the local Kalman estimates because either the prediction or update equations use global information. If the local knowledge deviates from the actual model, then the fusion equation will not produce the global estimate. The performance of DKF with pseudo estimates is compared with tracklet fusion in [6]. The distributed accumulated state density (DASD) filter [7] has a similar fusion equation but uses the ASD, which is the density of the augmented state.

The DKF and DASD filter can compute the optimal global estimate with no assumptions on the process noise and communication rate. However, the local estimates are pseudo-estimates and not Kalman estimates. Furthermore, the local pseudo-estimates are computed with global models; thus, these algorithms are more suitable for distributed processing and not for distributed estimation or fusion of local tracks.

ESTIMATE FUSION USING CROSS-COVARIANCE

Another popular fusion approach does exactly the opposite of decorrelation by exploiting the covariances and cross-covariances of the local estimates. This has advantages such as ignoring the dependence of the estimates due to prior communication and process noise, and the need to identify additional local estimates for decorrelation. However, the result is a constrained estimate which may be different from the centralized estimate given all the measurements.

The earliest work using this approach is the Bar-Shalom Campo rule [8]. Since the late 1990s, estimate fusion given the cross covariance has become a very active area of research because of its general applicability. Two popular ones are the maximum a posteriori (MAP) estimation [9], and the best linear unbiased estimation (BLUE) or weighted least-squares (WLS), [10], both first presented at FUSION 1999.

For two sensors, the MAP estimate has the same form as (1):

$$\hat{x} = \overline{x} + L_1\left(\hat{x}_1 - \overline{x}\right) + L_2\left(\hat{x}_2 - \overline{x}\right) \tag{3}$$

However, the gain matrix $L = \left[L_1, L_2\right]$ is calculated from the covariance matrix between the state x and the local estimates $\left[\hat{x}_1, \hat{x}_2\right]$. The BLUE or WLS [10] is a generalization of the MAP approach and can handle arbitrary correlations in the local estimates. For estimates $\left(\hat{x}_i\right)_{i=1}^{S}$ to be fused, a BLUE fusion rule is:

$$\hat{x} = \sum_{i=1}^{S} W_i \hat{x}_i \tag{4}$$

where the matrix weights $\left(W_i\right)_{i=1}^{S}$ are computed using the error covariances and satisfy the condition $\sum_{i=1}^{S} W_i = I$.

BLUE is very flexible because it can handle all types of local inputs and arbitrary correlations as long as the covariance between the inputs are known. The performance of BLUE fusion rules depends on the choice of the local inputs, and implementation requires knowing all the covariances. If the local inputs are chosen properly, the BLUE fusion rule can generate the centralized estimate given the measurements in the local estimates, such as augmented state estimate fusion. However, BLUE does not provide guidelines for selecting the local state estimates.

COVARIANCE INTERSECTION

The covariance intersection (CI) algorithm [11] was motivated by map building, where the cross-covariance between the thousands of variables is hard to model. CI assumes no knowledge of cross covariance. For two local estimates $\hat{x}_1$ and $\hat{x}_2$ with error covariances P_{11} and P_{22}, the CI algorithm is:

$$P_{CI}^{-1}\hat{x}_{CI} = \omega P_{11}^{-1}\hat{x}_1 + \left(1-\omega\right)P_{22}^{-1}\hat{x}_2 \tag{5}$$

where $\omega \in \left[0,1\right]$ is a parameter to be chosen such that the fused covariance P_{CI}^{-1} is minimal. CI produces a consistent estimate with a conservative error covariance. It is very popular when very little information on correlation in the estimates is available. The recently developed inverse covariance intersection [12] yields a good compromise between the conservative CI and other optimistic fusion rules such as naïve fusion. This is achieved by an application of the CI rule on the joint information of local estimates.

CONSENSUS FILTERS FOR SENSOR NETWORKS

Distributed estimation over a sensor network is difficult because the local estimates have correlations that depend on the information path. Fusion by decorrelation or using cross-invariances requires communication to share model and network information. Since communication is expensive, distributed estimation requires robust algorithms that assume only local network information, with performance measured by other metrics besides estimation accuracy. When distributed estimation is used to support distributed control, consensus in the estimates is more important than estimation accuracy. Thus, consensus filtering has become a very active area of research since the early 2000s [14]. It is based on the principle that a consensus estimate can be obtained by exchanging local information between observation times.

CONCLUSIONS

Much progress has been made in advancing the state of the art in distributed estimation over the past 25 years. However, not much has been done to provide guidance on selecting the appropriate algorithm for a particular problem. Further research to characterize the estimation performance, communication and computation requirements, and robustness of the algorithms is needed. Standard data sets and performance metrics will facilitate algorithm development and testing.

REFERENCES

1. Chong, C.-Y., Chang, K.-C., and Mori, S. Forty years of distributed estimation. In *Proc. 22nd Int. Conf. Information Fusion,* Ottawa, Canada, 2019.
2. Chong, C. Y. Hierarchical estimation. In *Proc. 2nd MIT/ONR Workshop Distributed Information and Decision Systems Motivated by Naval Command Control Communication Systems*, Monterey, CA, 1979.
3. Miller, M. D., Drummond, O. E., and Perrella, A. J. Tracklets and covariance truncation options for theater missile tracking. In *Proc. 1st Int. Conf. Multisource-Multisensor Data Fusion*, Las Vegas, NV, 1998.
4. Chong, C. Y., and Mori, S. Track association using augmented state estimates. In *Proc. 18th Int. Conf. Information Fusion*, Washington, DC, 2015.
5. Govaers, F., and Koch, W. Distributed Kalman filter fusion at arbitrary instants of time. In *Proc. 13th Int. Conf. Information Fusion*, Edinburgh, UK, 2010.
6. Chong, C. Y., Mori, S., Govaers, F., and Koch, W. Comparison of tracklet fusion and distributed Kalman filter for track fusion. In *Proc. 17th Int. Conf. Information Fusion*, Salamanca, Spain, 2014.
7. Koch, W., and Govaers, F. On decorrelated track-to-track fusion based on accumulated state densities. In *Proc. 17th Int. Conf. Information Fusion*, Salamanca, Spain, 2014.
8. Bar-Shalom, Y., and Campo, L. The effect of the common process noise on the two-sensor fused track covariance. *IEEE Trans. Aerospace and Electronic Systems*, Vol. 22, 6 (Nov. 1986), 803–805.
9. Mori, S., Barker, W. H., Chong, C. Y., and Chang, K. C. Track association and track fusion with non-deterministic target dynamics. In *Proc. 2nd Int. Conf. Information Fusion*, Sunnyvale, CA, 1999.
10. Zhu, Y., and Li, X. R. Best linear unbiased estimation fusion. In *Proc. 2nd Int. Conf. Information Fusion*, Sunnyvale, CA, 1999.
11. Uhlmann, J. K. General data fusion for estimates with unknown cross covariances. In *Proc. SPIE, AeroSense Confe*rence, Vol. 2755, 1996.
12. Noack, B., Sijs, J., Reinhardt, M., Hanebeck, U. D. Decentralized data fusion with inverse covariance intersection. *Automatica*, Vol. 79, (May 2017), 35–41.
13. Olfati-Saber, R., and Shamma, J. S. Consensus filters for sensor networks and distributed sensor fusion. In *Proc. 44th IEEE Conf. Decision Control*, Seville, Spain, 2005.

Target Tracking in the Framework of Possibility Theory

Branko Ristic
RMIT University
Melbourne, Australia
branko.ristic@rmit.edu.au

Uncertainty is the essential attribute of target tracking. It needs to be included in mathematical modelling of all aspects of a target tracking system, such as target dynamics, the target birth/disappearance process, measurement characteristics (probability of detection, characterisation of false alarms, and measurement noise), and prior contextual information (domain knowledge, such as maps and corridors, and historical data). The vast majority of target tracking algorithms are formulated as Bayesian inference problems, with uncertainty characterised by probabilistic models.

In practical real-world applications of target tracking, however, the specification of these (precise) probabilistic models for all uncertain aspects of this complex problem is often difficult. For example, the probability of detection as a function of range would depend on many unknown factors (e.g., environmental conditions, object size, and its reflective characteristics); hence, a precise trustworthy model would be almost impossible to put forward. This inherent misspecification often implies that several heuristics have to be introduced to compensate for the discrepancies between the model and the real data. It also relates to the common aphorism in statistics: all models are wrong, but some are useful.

In technical terms, probability theory deals with only one aspect of uncertainty involved in modelling complex systems (such as target tracking): uncertainty due to randomness [1]. This relates to the *known unknowns* feature of a model, where the probability function is known but the actual realisation is random (and hence unknown). Arguably, there is another layer of uncertainty involved in modelling, the *unknown unknowns* factor, often referred to as epistemic uncertainty. The existence of epistemic uncertainty has been the motivation behind several recent theories for quantitative modelling of uncertainty, such as possibility theory, Dempster-Shafer theory, and imprecise probability theory [2]. The focus in this article is on possibility theory, because (1) the standard probabilistic concepts can be (relatively easily) extended to this context and (2) at present, the last two aforementioned theories are primarily developed for and limited to discrete state spaces.

THEORETICAL FOUNDATIONS

UNCERTAIN VARIABLE

The concept of *uncertain variable,* in the adopted framework of possibility theory, plays the same role as a random variable in probability theory. The main difference is that the quantities of interest are not random but simply unknown, and our aim is to infer their true values out of a set of possible values. The theoretical basis of this approach can be found in [3], [4]. Briefly, the uncertain variable is a function $\mathbf{X}:\Omega\to\mathbb{X}$, where Ω is the sample space and $\mathbb{X}$ is the state space (the space where the quantity of interest lives). Our current knowledge about $\mathbf{X}$ can be encoded in a function $\pi_X:\mathbb{X}\to[0,1]$, such that $\pi_X(\mathbf{x})$ is the possibility (credibility) for the event $\mathbf{X}=\mathbf{x}$. Function π_X is not a density function; it is referred to as a possibility function, being the primitive object of possibility theory [5]. It can be seen as a membership function that determines the fuzzy restriction of minimal specificity (in the sense that any hypothesis not known to be impossible cannot be ruled out) about $\mathbf{x}$ [6]. Normalisation of π_X is $\sup_{\mathbf{x}\in\mathbb{X}}\pi_X(\mathbf{x})=1$ if $\mathbb{X}$ is uncountable and $\max_{\mathbf{x}\in\mathbb{X}}\pi_X(\mathbf{x})=1$ if $\mathbb{X}$ is countable.

The objective is to carry out inference on dynamical systems in a manner analogous to the Bayesian formulation. Then, it is natural to consider sequences $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_k\equiv\mathbf{x}_{1:k}$ of uncertain variables, with k being a discrete-time index and $\mathbf{x}_k$ representing the state of the target of interest in $\mathbb{X}$ at time k. Such a sequence of uncertain variables is an uncertain process (chain) [4]. An uncertain process is Markovian if $\pi_{k|1:k-1}(\mathbf{x}_k\mid\mathbf{x}_{1:k-1})=\pi_{k|k-1}(\mathbf{x}_k\mid\mathbf{x}_{k-1})$, for any $\mathbf{x}_1,\ldots,\mathbf{x}_k\in\mathbb{X}$.

NONLINEAR FILTERING

Nonlinear filtering in the framework of possibility theory is formulated next. Let the target dynamics be specified by the transition possibility function $\rho_{k|k-1}(\mathbf{x}\mid\mathbf{x}')$, which specifies the uncertain evolution of the state from time $k-1$ to time k. Let the uncertain relationship between the target-originated measurement $\mathbf{z}\in\mathbb{Z}$ and the (hidden) target state $\mathbf{x}$ at time k be specified by the likelihood function $g_k(\mathbf{z}|\mathbf{x})$, expressed as a possibility function. Here, $\mathbb{Z}$ is the measurement space. Given the dynamics model $\rho_{k|k-1}(\mathbf{x}|\mathbf{x}')$ and the measurement model $g_k(\mathbf{z}|\mathbf{x})$, the goal of the possibilistic nonlinear filter is to estimate recursively the posterior possibility function of the state, denoted $\pi_{k|k}(\mathbf{x}|\mathbf{z}_{1:k})$, where $\mathbf{z}_{1:k}$ is the sequence of target-originated measurements up to time k. Assuming the initial $\pi_0(\mathbf{x})$ at $k=0$ is known, the solution can be presented in two stages: prediction and update [7], [8]. The prediction equation is given by

$$\pi_{k|k-1}(\mathbf{x}\mid\mathbf{z}_{1:k-1})=\sup_{\mathbf{x}'\in\mathbb{X}}\rho_{k|k-1}(\mathbf{x}\mid\mathbf{x}')\pi_{k-1|k-1}(\mathbf{x}'\mid\mathbf{z}_{1:k-1}),\tag{1}$$

and it represents the possibilistic analogue of the Chapman-Kolmogorov equation. The update equation is given by

$$\pi_{k|k}(\mathbf{x}\,|\,\mathbf{z}_{1:k}) = \frac{g(\mathbf{z}_k\,|\,\mathbf{x})\pi_{k|k-1}(\mathbf{x}\,|\,\mathbf{z}_{1:k-1})}{\sup_{\mathbf{x}\in\mathbb{X}} g(\mathbf{z}_k\,|\,\mathbf{x})\pi_{k|k-1}(\mathbf{x}\,|\,\mathbf{z}_{1:k-1})}, \qquad (2)$$

and it represents the possibilistic analogue of the Bayes update. The only difference between the (standard) probabilistic formulation of nonlinear filtering [9, Ch.1] and the possibilistic formulation expressed by (1) and (2) is that integrals in the former are replaced with supremums in the latter. The possibilistic nonlinear filter is a special instance of an outer measure class of nonlinear filters defined in [7].

Application of the possibilistic nonlinear filter to spatiotemporal tracking using natural language statements was studied in [3]. The filter was implemented using a grid-based method. Application to target motion analysis (TMA) using bearings-only measurements was presented in [8]. The filter was implemented using a particle filter. The conclusions of [8] are noteworthy: in the absence of a model mismatch, the probabilistic TMA and the possibilistic TMA filters perform identically. However, if there is a model mismatch, either in the dynamic model or in the measurement model, the possibilistic TMA filter is more robust, resulting in a significantly lower rate of filter divergences. Application to space object tracking was presented in [10].

AN OVERVIEW OF RECENT DEVELOPMENTS

The Bernoulli filter for single-target joint detection and tracking in the presence of false detections and misdetections was developed in the possibilistic framework for two cases: for a point target in [11] and for an extended target in [12]. In both cases, it was demonstrated that the possibilistic approach is more robust if the probability of detection is known only as an interval value.

The analogue of the probability hypothesis density (PHD) filter, for joint estimation of the number of targets and their states, was derived in the framework of possibility theory in [4]. This filter provides modelling flexibility in terms of facilitating the introduction of measurement-driven birth schemes and modelling the absence of information on the initial number of targets. However, it loses the ability of the standard PHD filter to estimate the number of targets by integration of the intensity function.

The first multitarget tracking algorithm in the probabilistic framework was reported in [13]. It was developed as a possibilistic analogue of the δ-generalised labelled multi-Bernoulli (δ-GLMB) filter. As such, it inherits all the capabilities of the standard probabilistic δ-GLMB filter, with the additional ability to deal with partial knowledge of dynamic model parameters, measurement model parameters, and the initial number and states of newborn targets. The possibilistic δ-GLMB filter is implemented using the concept of a Gaussian max-mixture (a weighted combination of Gaussian possibility functions).

A reward function for sensor control using the possibilistic nonlinear filter was studied in the context of bearings-only tracking in [14]. The reward was defined as the uncertainty reduction, where a measure of uncertainty contained in a posterior $\pi_{k|k}(\mathbf{x}|\mathbf{z}_{1:k})$ is defined as the volume under $\pi_{k|k}(\mathbf{x}|\mathbf{z}_{1:k})$.

SUMMARY

The formulation of target tracking algorithms in the framework of possibility theory is an exciting recent development. However, one should see it not as a 'silver bullet' for all situations but rather as an alternative to the standard Bayesian framework, with the potential to provide an additional layer of robustness due to epistemic uncertainty. Although early studies suggest promising results, further work is necessary to establish in a more universal context the benefits and pitfalls of the proposed framework.

REFERENCES

1. O'Hagan, T. Dicing with the unknown. *Significance,* Vol. 1, 3 (2004), 132–133.
2. Ristic, B., Gilliam, C., Byrne, M., and Benavoli, A. A tutorial on uncertainty modeling for machine reasoning. *Information Fusion,* Vol. 55 (2020), 30–44.
3. Bishop, A. N., Houssineau, J., Angley, D., and Ristic, B. Spatiotemporal tracking from natural language statements using outer probability theory. *Information Sciences,* Vol. 463 (2018), 56–74.
4. Houssineau, J. A linear algorithm for multi-target tracking in the context of possibility theory. *IEEE Transactions on Signal Processing,* Vol. 69 (2021), 2740–2751.
5. Dubois, D., and Prade, H. Possibility theory and its applications: where do we stand? In *Springer Handbook of Computational Intelligence*. Berlin, Germany: Springer, 2015, 31–60.
6. Zadeh, L. A. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems,* Vol. 1 (1978), 3–28.
7. Houssineau, J., and Bishop, A. N. Smoothing and filtering with a class of outer measures. *SIAM/ASA Journal on Uncertainty Quantification,* Vol. 6, 2 (2018), 845–866.
8. Ristic, B., Houssineau, J., and Arulampalam, S. Robust target motion analysis using the possibility particle filter. *IET Radar, Sonar & Navigation,* Vol. 13, 1 (2019), 18–22.
9. Ristic, B., Arulampalam, S., and Gordon, N. *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Boston, MA: Artech House, 2004.
10. Cai, H., Xue, C., Houssineau, J., Jah, M., and Zhang, J. Possibilistic space object tracking under epistemic uncertainty. *Advances in Space Research,* 2023, in press.
11. Ristic, B., Houssineau, J., and Arulampalam, S. Target tracking in the framework of possibility theory: the possibilistic Bernoulli filter. *Information Fusion,* Vol. 62 (2020), 81–88.
12. Chen, Z., Ristic, B., and Kim, D. Y. Possibilistic Bernoulli filter for extended target tracking. In *Proceedings of the IEEE International Conference on Acoustic, Speech and Signal Processing (ICASSP),* Rhodes, Greece, June 2023.
13. Cai, H., Houssineau, J., Jones, B. A., Jah, M., and Zhang, J. Possibility generalized labeled multi-Bernoulli filter for multitarget tracking under epistemic uncertainty. *IEEE Transactions on Aerospace and Electronic Systems,* Vol. 59, 2 (April 2023), 1312–1326.
14. Chen, Z., Ristic, B., Houssineau, J., and Kim, D. Y. Observer control for bearings-only tracking using possibility functions. *Automatica,* Vol. 133 (2021), 109888.

Artificial Intelligence, Machine Learning and Sensor Data Fusion

Lyudmila Mihaylova
The University of Sheffield
Sheffield, United Kingdom
l.s.mihaylova@sheffield.ac.uk

Wolfgang Koch
Fraunhofer FKIE
Wachtberg, Germany
w.koch@ieee.org

Recent years witnessed tremendous developments in artificial intelligence (AI), machine learning (ML), computer vision, and autonomous systems. While AI focusses on incorporating human intelligence to machines, ML can be seen as a range of tools aimed to empower computer systems with the ability to "learn". AI is seen as a broader concept compared with ML [1]. Figure 1 shows the relationship between these three related areas.

Considered in the light of sensor data fusion and the International Society of Information Fusion (ISIF), the area of ML has been present with different developments and in various ways–from biologically inspired neural networks to sequential Monte Carlo probabilistic methods for non-linear systems with non-Gaussian distributions. However, it is mainly in recent years, when ML methods became popular and expanded towards trustworthy ML and explainable AI. These are especially linked with the necessity to introduce different levels of autonomy [2], [3] and find the reasons or causality of events which brings the level of explainability. These two are especially linked with sensor data and nowadays data come both from "hard sensors" from different modalities such as radar, acoustic sensors, LiDAR, combined with optical, thermal cameras, and wireless sensor networks but also from soft sensing modalities (Internet of Things, social networks such as Twitter, Facebook, and others). Moreover, data arrives with different time rates and levels of accuracy. Making sense of such multiple heterogeneous data is a challenging task that has been extensively studied, but the provision of reliable solutions for autonomous and semi-autonomous systems is a task that remains only partially solved. Fusion of data from multiple heterogeneous sensors of this type is part of the challenge; even more so when the autonomous decisions have to be performed sequentially and in real-time. This is especially important for safety critical tasks such as with unmanned aerial vehicles (UAVs), aircraft flight control systems, the Future Combat Air System, digital health systems, and many others.

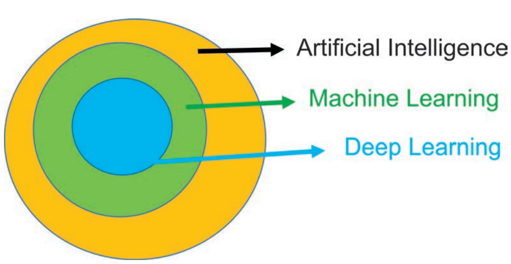

Figure 1
Artificial intelligence, machine learning, and deep learning [1].

ML methods can be subdivided into *model based* and *data-driven*. There is an increasing interest especially in reinforcement learning with many applied areas, one of them is for smart cities [4]. There is a trend towards data driven methods in which mathematical models are not necessarily present. Instead, patterns from data are autonomously learned and captured to represent these patterns and work without mathematical models that have many parameters and are difficult to calculate in short time scales. At the same time, AI methods need to be able to deal not only with big data, but with missing or incomplete data. Representing confidence levels and uncertainty from the integration of heterogeneous large-scale data still remains a challenging task. This leads to the next question about the level of trust in the developed AI methods.

TRUST, TRUSTWORTHY SOLUTIONS, AND EXPLAINABLE AI

ML methods as a branch of AI have been actively developed in the past decades to address the tasks of trustworthiness. We need to know where the strengths of AI methods are and when we can rely on them. AI provides a range of useful tools, but these can work well under certain conditions; for instance, different environmental or methods related constraints. An example of important environmental conditions for ML and computer vision methods are lighting conditions and other weather conditions or intentional adversarial changes (called adversarial attacks aimed to modify the data and to mislead the overall solution, e.g., in image classification and segmentation). Awareness of such challenges, constraints and other limitations needs further theoretical results and their practical validation before having AI algorithms as part of a UAV or an airplane, used without the presence of a human.

TRUST

To answer this question about trustworthiness of the developed solution, the first step is to characterise what we understand by "trust" in this context. The word trust means: "firm belief in the reliability, truth, or ability of someone or something". Being aware of this, the next question is how to characterise it numerically and have it as part of the learning process in AI solutions. The level of trust can be specified by a probabilistic measure, such as a variance of a Gaussian distribution, a score, a fuzzy logic rule, or by other ways. Under a Gaussian assumption about the considered noises, by propagating the mean and the variance could be a way to answer such questions. The variance as a tool of uncertainty quantification has been proven to be very powerful, especially in image classification, segmentation object tracking, and other inference tasks and can be represented within the Gaussian process methods framework and other upper bounds [5].

THE USER'S PERSPECTIVE

The development of methodological foundations during the past decades was linked with areas such as image fusion, time series analysis, reinforcement learning for robotics, transport systems, communications, and many others. The level of trust in the AI solution needs to be communicated quickly and in the best way to the users.

The user needs to trust the AI systems and be able to operate easily with them. The users need to understand what the AI system is offering, how to use it, and its advantages and limitations. However, the user may not necessarily need to know how exactly the AI system is designed and what methods are embedded.

EXPLAINABLE AI

Explainable AI has a big potential to find the main factors and inherent causes of events and occurrences. Explainable, responsible AI are concerned with questions like: "What is happening and what are the consequences of it?". Heat maps can be especially useful to understand where the objects of interest are, how to interpret them in the context of the overall task. and decision making. Heat maps could be seen also as a tool of quantifying uncertainties and understanding where things work and where deficiencies are present. An example is a heat map for the solution from image classification or for localisation with fingerprinting (Gaussian process methods).

Trustworthy, explainable, and resilient AI solutions need to be modular and to afford further development of all their components during the whole cycle of life. These could be achieved with efficient fusion at the different levels of sensor data, information, knowledge, and ontologies. Scalability adds another level of requirement and it is needed not only with respect to data, states (objects of interest), but is linked with communication constraints, especially for real-time tasks.

DEEP LEARNING FOR DATA FUSION

Data fusion methods have received a lot of developments over the past decades. Well-established methods for tracking such as the interacting multiple model filters, multiple hypothesis tracking [6], [7], or other fusion approaches based on the Dempster–Shafer theory have reached a high level of maturity. In the past, mainly high-level fusion algorithms were developed—for decision making, command and control, knowledge fusion, and fusion of ontologies, whereas the past 10 years witnessed the development of low-level fusion methods—such as for centralised, decentralised tracking, navigation, localisation, situation awareness, and related areas gained a momentum.

Current trends include developments of multiple types of sensor data fusion with convolutional neural networks (CNNs), transformers, kernel methods such as Gaussian process regression, and combinations between them, variational inference, to name a few and many others. New results were reported with deep learning methods, reinforcement learning for image classification, image segmentation, and others. ML methods are also core methods for cyber-security and cyber-physical systems.

Still fusion of multiple types of sensor data with deep learning methods for object detection, multiple target tracking, and localisation is an open area of research. How to fuse data from different modalities, such as images with inertial measurement unit data with data from social networks and other data, needs further attention.

CONCLUDING REMARKS

AI and ML methods are capable of providing efficient solutions and these are valuable to support human decisions, e.g., a pilot of an aircraft operating in difficult weather conditions or autonomous landing of a UAV. The development of trustworthy, resilient ML methods for cyber-physical systems is a big area of research that needs further attention and explainable results.

ACKNOWLEDGMENTS

L. Mihaylova is grateful to the support of EPSRC through Project NSF-EPSRC: ShiRAS; Towards Safe and Reliable Autonomy in Sensor Driven Systems, under Grant EP/T013265/1, and the ShiRAS part, co-funded by the US National Science Foundation under Grant NSF ECCS 1903466.

REFERENCES

1. Garbade, M. J. Clearing the confusion: ai vs machine learning vs deep learning differences, https://educationecosystem.com/blog/ai-vs-machine-learning-vs-deep-learning-differences/, last access 20 Feb. 2023.
2. Vagia, M., Transeth, A. A., and Fjerdingen, S. A. A literature review on the levels of automation during the years. What are the different taxonomies that have been proposed? *Applied Ergonomics*, Vol. 53, Part A (2016), 190–202.
3. Wang, J., Huang, H., Li, K., and Li, J. Towards the unified principles for level 5 autonomous vehicles. *Engineering*, Vol. 7, 9 (2021), 1313–1325.
4. Yan, Z., Kreidieh, A. R., Vinitsky, E., Bayen, A. M., and Wu, C. Unified automatic control of vehicular systems with reinforcement learning. *IEEE Trans. Automation Science and Engineering,* (2022), 1–16.
5. Liu, X., Mihaylova, L., George, J., and Pham, T. Gaussian process upper confidence bounds in distributed point target tracking over wireless sensor networks. *IEEE Journal of Selected Topics in Signal Processing*, Vol. 17, 1 (Jan. 2023), 295–310.
6. Bar-Shalom, Y., Li, X. R., and Kirubarajan, T. *Estimation with Applications to Tracking and Navigation*. New York, NY: Wiley-Interscience, 2001.
7. Bar-Shalom, Y., Willett, P. K., and Tian, X. *Tracking and Data Fusion: A Handbook of Algorithms*. Storrs, Connecticut: YBS Publishing, 2011.

Context-Enhanced Information Fusion

Jesus García
Universidad Carlos III de Madrid
Madrid, Spain
jgherrer@inf.uc3m.es

Lauro Snidaro
University of Udine
Udine, Italy
lauro.snidaro@uniud.it

The concept of *context* has been used in computer science for several decades, with early work dating back to the 1980s. However, the popularity of context-aware computing has significantly increased in recent years due to advances in mobile and ubiquitous computing, the Internet of Things, and the proliferation of data from various sources. The term "context" has been used in various subfields of computer science, including artificial intelligence (AI), human–computer interaction, information retrieval, and data management. The importance of context has been recognized in various application areas, including healthcare, transportation, and smart cities. However, a clear definition is still lacking due to the diversity of applications.

Following the positive response that special sessions on context-enhanced information fusion (IF) have received at the International Conference on Information Fusion, this short paper aims at providing an overview of current research, presenting works covering aspects that include contextual elements in the fusion process. The reader is referred to [1], [2], a survey and collection of works on context-enhanced IF.

DEFINITION OF CONTEXT

Context refers to the circumstances or situation in which something exists or occurs and can affect its meaning, interpretation, or significance. Context can include various factors, such as the physical environment, cultural background, social norms, historical events, prior experiences, and related factors.

In communication, context plays a crucial role in understanding the meaning of a message, as it provides the necessary background information and clues for interpreting the message correctly. For example, the meaning of a word or phrase may change depending on the context in which it is used.

In broader terms, context is also used to describe the overall framework or perspective that shapes how we view and interpret information, events, or situations. Understanding the context of a particular situation can help in making more informed decisions, forming more accurate judgments, and communicating more effectively with others.

Many definitions have been proposed in the literature; here we report the one proposed by Steinberg [3] that highlights the relational nature of context: a *context* is a *situation*. More specifically, if a situation is a set of relationships, then a context c could be understood as the subset of a situation s that can be used to resolve (estimate or infer) a set of random variables X.

THE PROBLEM

Over the past few years, it has become increasingly clear that simply combining data from multiple sources may not be enough to improve fusion systems' performance. Even with a potentially large number of sources, unexpected results may occur if the value being estimated, the error characteristics of the sources, and the fusion process itself are not properly contextualized. For example, the state of a target may depend on various factors such as the environment, nearby entities, time of day, and weather conditions. When estimating the position and speed of a car in city traffic, various factors such as the bends and turns of the road, condition of the asphalt, traffic signs, and overall traffic conditions can affect the state of the car.

Generally, context awareness involves considering and using information and knowledge about the environment or current situation surrounding the focal element of interest. However, the understanding and application of context in IF systems are still limited, with domain knowledge being traditionally acquired ad hoc and applied to stovepiped solutions. To improve adaptability and performance, context should play a crucial role at any level of a modern fusion system.

EXAMPLES OF APPLICATIONS

The possibilities of applications for context-based IF are diverse (improved estimation and classification, sensor characterization and management, decision making, situation sense-making, etc.). Some examples of recent applications in different fields are reported in the following. In [4], an example in the domain of electronic combat is used to illustrate context formalization through ontologies. It shows how context can facilitate the representation of entities at different fusion levels to make possible inferencing among levels, with a consistent representation of entities, the states at different levels and relationships. An example in the domain of environment perception for automated vehicles is presented in [5], where accurate detection models are required for safe operation. The authors address the evaluation of false object hypotheses in complex scenarios with high density in order to verify the existence of a tracked object with a probabilistic model, considering the influences of multiple digital map elements on each track's existence for every track in urban scenarios. In [6], a probabilistic approach is used for forecasting vessel trajectories. Context is exploited through a discrete probabilistic model with typical vessel behaviors using dynamic Bayesian networks to predict the speed and orientation of a vessel with a discretized representation of the space. In [7], context-aware data fusion is used in the design of personalized monitoring systems. In this domain, the specific vocabulary, facilities, and events of interest are modeled in order to develop customized monitoring solutions. In [8], an example in the domain of airborne passive localization of stationary ground-based emitters is used, including roadmap-assisted target tracking and integration of terrain map data for target localization. The authors show the integration of contextual knowledge into target tracking algorithms, exploiting the constraints on the target state both in the prediction step and in the measurement update step of a tracking filter. In [9], pretrained word embeddings, typically used for natural language processing, are fused to estimate word concreteness. The authors analyzed how much contextual information can affect final results and how to properly fuse different word embeddings in order to maximize their performance for a word concreteness task. Finally, the example in [10] shows how geographic datasets of roads and buildings can be enhanced with more contextual information by means of automatic processes exploiting available sensors onboard vehicles, like lidar and 360-degree cameras.

FUTURE DIRECTIONS

In the area of AI, building "explainable" solutions has become an emergent research trend, as indicated in surveys like [11]. A fundamental challenge for next-generation AI systems will be the ability to adapt to contextual conditions. In this sense, a "context-aware" system is a paradigm in the AI community where the interaction with users improves significantly when high-level concepts are used by the system to explain outputs, in opposition to black-box solutions. Therefore, it seems reasonable to expect that next waves of AI will put more focus on how to perform inferences, also considering contextual factors and incorporating contextual models over time in the learning process. For instance, the model of "rational rules" [12] combines the inferential power of Bayesian induction with the representational power of mathematical logic and generative grammars for concept generalization. Similarly, Markov logic networks integrate probabilistic models with first-order logic to enable inferences under uncertainty [13]. In both cases, the possibility of using a symbolic representation of the concepts learned allows the system both to generalize and to adapt to specific conditions for each domain.

A fundamental challenge identified in both AI and IF future systems is "understanding" context, the capability to represent and relate how relevant the context is to the inference problems addressed, along with mechanisms to adapt the inference processes to this context. In this parallelism, the challenges include perception, reasoning, and context adaptation toward deploying AI and IF systems to support knowledge representation and situation understanding. A key objective is the capability to learn interpretable models from contextual data to bind observations with knowledge and use the semantics provided by context. In conclusion, recent developments in research show a convergence in IF and AI systems for situation understanding, where efforts are being made to develop representations and models that allow automatic adaptation to domain conditions.

ACKNOWLEDGMENT

J. García was partially supported by the Spanish Ministry of Science and Innovation in projects PID2020-118249RB-C22, PDC2021-121567-C22 and TED2021-131520B-C22.

REFERENCES

1. Snidaro, L., García, J., and Llinas, J., Context-based information fusion: a survey and discussion. *Information Fusion*. [Online] Vol. 25 (2015), 16–31.
2. Snidaro, L., García, J., Llinas, J., and Blasch, E. (Eds.) *Context-Enhanced Information Fusion—Boosting Real-World Performance with Domain Knowledge*. Cham, Switzerland: Springer, 2016.
3. Steinberg, A. N. Situations and contexts. *Perspectives Magazine, International Society on Information Fusion,* Vol. 1, 1 (2016).
4. Steinberg, A. N., and Bowman, C. L. Synergy in multi-level reasoning. In *Proceedings of the* 2022 *IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA),* 2022, 23–30.
5. Gies, F., Posselt, J., Buchholz, M., and Dietmayer, K. Extended existence probability using digital maps for object verification. In *Proc. of the IEEE 23rd International Conference on Information Fusion,* 2020.
6. Jansen, L., Pavlin, G., Atamas, A., and Mignet, F. Context-based vessel trajectory forecasting: a probabilistic approach combining dynamic Bayesian networks with an auxiliary position determination process. In *Proc. of the IEEE 23rd International Conference on Information Fusion,* 2020.
7. Iglesias, C., Miceli, C., and Silva, D. A domain model for personalized monitoring system based on context-aware data fusion. In *Proc. of the 22nd International Conference on Information Fusion,* 2019.
8. Ulmke, M., and Govaers, F. Gaussian mixture based target tracking combining bearing-only measurements and contextual information. In *Proc. of the 21st International Conference on Information Fusion,* 2018.
9. Incitti, F., and Snidaro, L. Fusing contextual word embeddings for concreteness estimation. In *Proc. of the IEEE 24th International Conference on Information Fusion,* 2021.
10. Amigo, D., Pedroche, D. S., García, J., and Molina, J. M. Automatic context learning based on 360 imageries triangulation and 3D lidar validation. In *Proc. of the IEEE 24th International Conference on Information Fusion,* 2021.
11. Biran, O., and Cotton, C. Explanation and justification in machine learning: a survey. *IJCAI-17 Workshop on Explainable AI,* Vol. 8, 1 (2017), 8–13.
12. Goodman, N. D., Tenenbaum, J. B., Feldman, J., and Griffiths, T. L. A rational analysis of rule-based concept learning. *Cognitive Science*, Vol. 32, 1 (2008), 108–154.
13. Domingos, P., and Lowd, D. Unifying logical and statistical AI with Markov logic. *Communications of the ACM,* Vol. 62, 7 (2019), 74–83.

Interoperability of Information Fusion Systems

Mieczyslaw M. Kokar
Jakub J. Moskal
VIStology, Inc.
Framingham, MA, USA
mkokar@vistology.com
jmoskal@vistology.com

Information Fusion (IF) studies solutions for combining information from multiple sources into one representation that is more concise, sound, and informative than any source individually or a disjointed union of all sources. In this paper, we assess three issues that are currently hampering the interoperability of IF systems: 1) Ontology of IF, 2) Formal Theory of Information Integration and Fusion, and 3) Situational Awareness.

We show a generic information fusion (IF) system in Figure 1. It inputs two (or more) sets of information (Info 1, Info *n*) from the sources (Source 1, Source *n*), *represented* in some data structure, and inserts them (Integrate) into one data store (Info 1 ⊕ Info *n*). This is just a disjoint union of the two data sets. The next step (Compose) aligns the data items that represent *the same* information, followed by the integration and fusion part of the algorithm that combines those data items into new assertions about the world. This process may also involve assigning degrees of "belief" to each assertion. Figure 1 also shows the operation of *inference,* indicated by the Derive labels. The sizes of the rectangles representing inferred sets of facts indicate how much new information is inferred by the specific Derive engines. The intent here is to show that the amount of information after the operation of Compose is larger than K1, K*n*, and K1+K*n*, while Info is more concise than Info 1 ⊕ Info *n*.

Since IF involves multiple information sources, each represented in different schemes, there is a need to first align the schemes. Usually, IF is concerned with resolving the coordinates of an object that is detected by a sensor. IF normally starts with a concrete representation of all the variables that are relevant to the problem, and the data are collected as values of such variables. The variables, typically, are of the type Real, and the names of the variables are provided by the designer of an IF system. The mappings between two sets of real numbers are obvious: 5 = 5 and 3 ≠ 5. This is because the language used does not have any ambiguities. However, if the information is presented in a language that does not have types (like Real numbers) but rather uses linguistic terms, it is necessary to align them [1]. The reason for this is that nonmathematical languages do not follow the principle of unique name assumption (UNA), in which two items with the same identifier (ID) are the same and two items that have different IDs are not necessarily different (e.g., Tom ≠ Bob), might need to be explicitly asserted.

In addition, the open-world assumption (OWA), often used in logic-based languages, posits that an observer can never have complete knowledge and therefore cannot deem something false merely because it lacks evidence of its truth. In such a case, the agent can only say that the status of a statement is "unknown". An opposite to OWA is the closed-world assumption (CWA). In CWA, inherent in database systems, if a piece of information is absent from the database, it is automatically false. Consequently, fusion agents that accept UNA, OWA, or CWA may arrive at vastly different decisions.

These aspects result in a requirement for alignment of types and then composition of multiple information sets into one in a consistent way. Thus, not only do the items from two information sources (Info 1, Info *n*) need to be treated as one, but the implications of such associations also need to be analyzed. For instance, if one source has an object Alice, the other has an object Bob, and the alignment states that (Alice = Bob), then all instances of Alice and Bob in each of the two information sets need to be replaced with a term that is not used for something else. However, since both Alice and Bob may occur in different

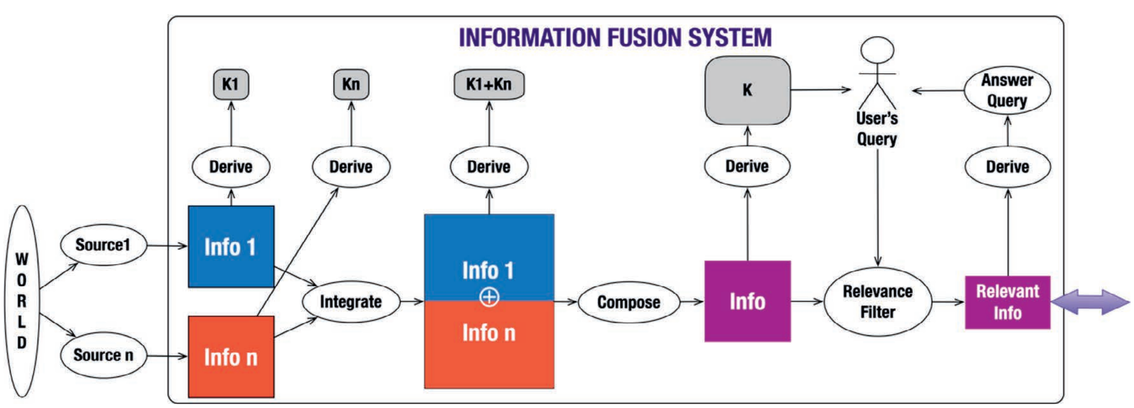

Figure 1
A conceptual view of an information fusion system.

relations with other objects, e.g., <Bob friend Tom> and <Alice enemy Tom>, this may cause a logical inconsistency, provided that the system knows that (friend ≠ enemy). Then, such an alignment is not admissible. However, if we want to align Bob with Tom and if the system knows that <Bob friend Alice> and <Tom friend Carl>, the system should infer that the object called Bob is friends with both Alice and Carl. In summary, alignment should be admissible if it does not imply inconsistencies, and alignment can result in the inference of relations between the aligned objects and other objects.

ONTOLOGY OF IF

If IF systems, like the one shown in Figure 1, were interoperable, they would be able to exchange information and knowledge shown in the figure (subject to the policy of the owner of the system)—information interoperability. The systems would need to use the language that they can understand, and the language should be standardized—a valuable goal for the IF society. Such a language is called *ontology*. Ontologies have been used in IF for a long time; see, e.g., [2], [3]. The use of ontologies has become quite popular. However, so far, we have not seen an ontology of IF that follows the ontology standards of the World Wide Web Consortium.

Such an ontology would need to provide for: (1) representation of object *types* (e.g., sensors, vehicles, and people), processes/functions, and their *instances*; (2) classifications of objects and processes, e.g., fusion functions; (3) representation of the Joint Directors of Laboratories (JDL) and levels; (4) representation of relationships between them; and (5) capability to specify IF systems and apply formal inference to *deduce* other facts that are only implicit in the ontology describing the systems. Ontologies then could be used to represent specific IF systems, and the representations could be used to infer *relations between the systems*. An example used in our paper was a *data fusion* system modeled as a subclass of a *decision fusion* system [4].

We are not aware of the existence of a generic ontology for IF. Although there are papers that refer to "information ontology" and "data fusion ontology", they do not satisfy all five of the above criteria. On the other hand, there are many papers that use specific ontologies inside of their IF systems.

THEORY OF INFORMATION INTEGRATION AND IF

The next level of interoperability would be the reuse of algorithms from one system in another. However, any algorithm interacts with many other algorithms within the system, so when we plug an algorithm into a system, it needs to be integrated with the whole system. To achieve this level of interoperability, IF needs methods of *algorithm composition* and *integration*. In our early searches for a theory of IF, we identified *category theory* as the most appropriate tool to be used as a most general model of integration of information; we introduced it to the IF community in [4]. We showed that algorithms cannot be combined by set-theoretic operations like union or product, but a more abstract category theory provides means for doing this—*morphisms, limits,* and *colimits*. We have used category theory in some of our work. A deeper theoretical investigation of fusion using category theory in general, and sheaves in particular, can be found in [5].

SITUATIONAL AWARENESS

We consider an IF system agent (or user) to be aware when they can answer queries about specific situations. Thus, a higher level of interoperability is to exchange information about situations among the IF systems and let the information receiver infer answers to the user queries (Figure 1). Within the JDL Data Fusion Model, which is widely used within the IF community, situational awareness is positioned at level 2 as a process and labeled Situation Assessment. The basis for this process is Endsley's work, e.g., [6]. In this conceptualization, "situation assessment" is understood as a process that can be measured and evaluated. If the process is efficient, then a high level of "situation awareness" is achieved. However, e.g., [3], [7] consider situations as objects, which have their own existence and thus can be described, and their descriptions can be exchanged, learned, and so on. The awareness of an agent of a specific situation is assessed then as the capability of answering queries *about the situations*. As shown in Figure 1, the process of Relevance Filter is part of an IF system. Its objective is to identify which parts of information and knowledge are relevant to a specific user's query. Relevant Info is then the description of the situation the user is inquiring about. It can be conveyed to another agent—either human or computer—for their use. In this scheme, Relevant Info is much smaller in size than all the information in the IF system; thus, it is less demanding on both the bandwidth of the communication channels and the user's cognitive load.

REFERENCES

1. Dapoigny, R., and Barlatier, P. Formal foundations for situation awareness based on dependent type theory. *Information Fusion*, Vol. 14, 1 (2013), 87–107.
2. Boury-Brisset, A. C. Ontology-based approach for information fusion. In *Proc. 6th Int. Conf. Information Fusion,* Vol. 1 (July 2003), 522–529.
3. Kokar, M. M., Matheus, C. J., and Baclawski, K. Ontology-based situation awareness. *Information Fusion,* Vol. 10, 1 (2009), 83–98.
4. Kokar, M. M., Tomasik, J. A., and Weyman, J. Formalizing classes of information fusion systems. *Information Fusion,* Vol. 5, 3 (2004), 189–202.
5. Robinson, M. Sheaves are the canonical data structure for sensor integration. *Information Fusion,* Vol. 36 (2017), 208–224.
6. Endsley, M. R. Toward a theory of situation awareness in dynamic systems. *Human Factors,* Vol. 37, 1 (1995), 32–64.
7. Lu, S., and Kokar, M. M. Inference-based information relevance reasoning method in situation assessment, 2023, under review.

High-Level Information Fusion Situational Developments

Erik Blasch
MOVEJ Analytics
Fairborn, OH, USA
erik.blasch@gmail.com

High-level information fusion (HLIF) was coined in the 1990s following the original Joint Directors of Laboratories (JDL) working group model [1] and updated by the same working group in 2004 as the Data Fusion Information Group/JDL model [2]. One key discussion in the 1990s was the debate about awareness and assessment. *Awareness,* such as from Boyd's control loop [3] and popularized by Endsley et al. [4], is a humancentric concept. *Assessment* was the machine counterpart to awareness that includes the developments from sensing that afforded that ability to collect and analyze data. Figure 1 showcases the alignment of awareness and assessment between low-level information fusion (LLIF) and HLIF.

While the information fusion (IF) community focused on the LLIF of data preprocessing (level 0) and object filtering, estimation, and prediction (level 1), a sensor data fusion system product requires HLIF constructs of situation (level 2) and impact (level 3) assessment, along with sensor (level 4), user (level 5), and mission (level 6) refinement. Hence, systems compose the duality of fusion (machine assessment) and control (human refinement). It is noted here that the control of sensor and information in level 4 process refinement is designed by humans to control data processing.

HLIF CHALLENGES

HLIF has design challenges of (1) assessment, (2) design, and (3) cognition. Notions of situation and impact assessment are overlapping with "awareness". Awareness is rooted in human perceptions of the world; hence, the aggregation of LLIF information from a system is a *producer* that provides data to situation assessment for situation awareness. Similarly, level 3 threat or impact assessment is a function of the needs by a *consumer* of the sensor data fusion *design*.

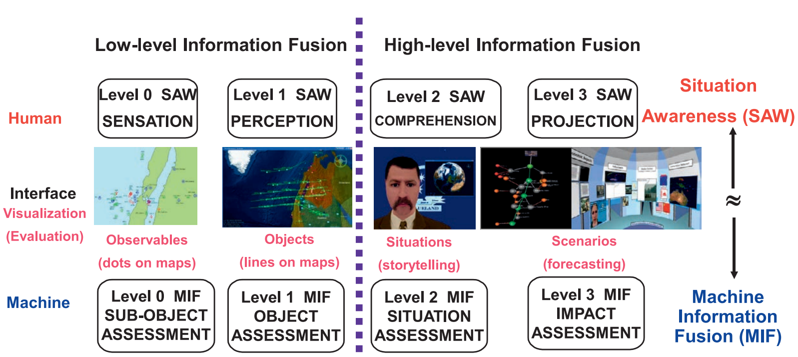

Figure 1
Awareness versus assessment.

The second area of HLIF is sensor, user, and mission (SUM) *refinement*. Since the IF system is a function of the product (e.g., architecture design, human user, and business operations), management of the information is conducted by refining the needs, such as sensor pointing, user satisfaction, and mission achievement. Essentially, HLIF allows stakeholders to tailor the system performance by exploiting the information they desire [5].

The term *cognition* is used often to seek elements of awareness that include reasoning, imagination, and understanding. Cognition is also a recent construct in machine intelligence (e.g., cognitive radar), from which the machine is aware of its own processes through self-assessment. Hence, data fusion system design requires physics-based and human-derived information fusion (PHIF) in translating machine sensor data results to human cognitive semantic meaning [6].

As from the HLIF taxonomy, issues of SUM management are more prescribed than those of situation and impact assessment. The assessment function can be that of the "situation", driven by context from data (assessment) and knowledge (awareness) [7].

HLIF DEVELOPMENTS

Over the last 25 years, even though it is well discussed in conceptual theory, there is no answer or solution to situation assessment. From the many papers referenced in published reviews [2], HLIF discussions emphasized the need for situation analysis, evaluation metrics, and realizable architectures [8]. One problem is that *scale* of the situation also depends on the SUM perspective on whether it is from the sensor (machines), user (organizations), or mission (purpose). As much as spatial, temporal, and spectrum scales, the same exists for HLIF. A PHIF system design as a user-defined operating picture (UDOP) is not the same as a mission common operating picture (COP). For example, the UDOP for the pilot must interface with the COP for an air traffic controller.

A key representation of the discussion through the 2000s was utilizing the advances in text analytics toward semantic meaning. Evaluations sought ways to measure a "situation". An example was that of hard (physics-based sensing) and soft (human-derived) *semantic constructs* [9] for UDOP surveillance systems, helping the user control sensing and supporting a narrative output.

The 2010s ushered in artificial intelligence (AI) as machine and deep learning (DL). The explosion of DL was applied to all forms of LLIF while discussions began toward using DL for situation, user, and mission analysis.

Since data-driven learning methods do not reason, there is still considerable need for HLIF research, especially for machine-supported situational assessment. With ever-increasing computing power to support HLIF, the current themes are in *multidomain operations*. Multidomain refers to data collection and application (space, air, ground, subsea, cyber, etc.). Currently, there are many multidomain paradigms for cloud, fog, and edge computing for IF. Hence, the recent decade focused on situation assessments such as from PHIF semantic cognition (reasoning), data utilization (context), and domain prediction (control).

APPLICATIONS

The design, development, and deployment of IF systems require many stakeholders coordinating the governance (policies), people (users), acquisition (buyers), and design (developers). Three examples include (1) surveillance, (2) logistic, and (3) infrastructure systems, all of which include various data, sensors, and IF. Most commonly discussed at the International Conference on Information Fusion are *surveillance systems* for ground-space management, such as underwater, battlefield, air, and space awareness.

HLIF for *logistics* includes the medical, aviation, and information communities. For example, the medical community includes patient care (diagnosis) and drug delivery (prognosis) from which LLIF data processing and analysis feed human–machine control of parts and supplies.

The *infrastructure* is critical for all systems, such as supporting the energy and supply grid. Utilizing LLIF signal processing, HLIF supports users of plant operations and maintenance engineers of components determined for failure and repair to enhance performance, safety, and security.

Operating pictures (Traffic management)—The development of a global positioning system (GPS) for ground systems (adopted in 1990), automatic identification system (AIS) for shipping (adopted in 2000), and automatic dependent surveillance–broadcast for aviation (adopted in 2020) have accelerated the constructs of HLIF SUM for traffic management, wit h future versions for space traffic management. Because the position, navigation, and timing data project the platform location, registration of many sensors affords the ability to sense (e.g., a camera on an unmanned aerial vehicle), use (e.g., displays), and task (e.g., route) systems for safety purposes. Hence, many COPs/UDOPs have been designed for space, air, ground (e.g., autonomous car displays), and maritime domain awareness through HLIF assessment.

Distributed human-machine teaming (Medical)—To acknowledge where HLIF added to decision speed for society improvement, the coronavirus disease 2019 pandemic provides a good use case. The many distributed researchers shared multimodal data, measurements, polices, and results.

LLIF data fusion (e.g., temperature and face detection) tools monitored the number of people complying with various polices to support government, commercial, and individual users through displays and apps. The detection and subsequent classification (who was wearing a mask) were correlated with the outbreak and spread. Diagnostic systems reported on the level of severity of individuals, with a widely available "heat map" of outbreaks to assess the "situation" and provide "awareness". At the same time, the introduction of vaccines required the distribution of these products to meet the mission need to get as many people vaccinated as possible. The medical community used these tools to aid the surge of where to place medical professionals, which policies to adopt, etc., which provided HLIF SUM refinements as to opportunistic placement of people, supplies, and polices to reduce the spread.

As a way forward, challenges of situation assessment for situation awareness still require IF researchers to characterize the "situation" uncertainty, determine the certification of systems, and support societal stakeholder needs [10].

REFERENCES

1. Waltz, E., and Llinas, J. *Multisensor Data Fusion*. Norwood, MA: Artech House, 1990.
2. Blasch, E., Kadar, I., Salerno, J., Kokar, M. M., Das, S., Powell, G. M., Corkill, D. and Ruspini, E. Issues and challenges in situation assessment (level 2 fusion). *J. of Advances in Information Fusion,* Vol. 1, 2 (Dec. 2006), 122–139.
3. Boyd, J. R. Destruction and creation. U.S. Army Command and General Staff College, Leavenworth, KS, Sept. 1976.
4. Endsley, M. R., Bolte, M., and Jones, D. *Design for Situation Awareness*. Boca Raton, FL: Taylor and Francis, 2003.
5. Blasch, E., Steinberg, A., Das, S., Llinas, J., Chong, C.-Y., Kessler, O., Waltz, E. and White, F. Revisiting the JDL model for information exploitation. In *Proc. Int. Conf. Information Fusion*, (Istanbul, Turkey), 2013.
6. Blasch, E., and Aved, A. J. Physics-based and human-derived information fusion video activity analysis. In *Proc. Int. Conf. Information Fusion*, (Cambridge, UK), 2018.
7. Garcia, J., Snidaro, L., Context-Enhanced Information Fusion, *Perspectives on Information Fusion*, Vol. 6, 1 (Jun. 2023), 51–52.
8. Blasch, E. P., Lambert, D. A., Valin, P., Kokar, M. M., Llinas, J., Das, S., Chong, C.-Y. and Shahbazian, E. High level information fusion (HLIF): survey of models, issues, and grand challenges. *IEEE Aerospace and Electronic Systems Magazine,* Vol. 27 9 (Sept. 2012), 4–20.
9. Pravia, M. A., Prasanth, R. K., Arambel, P. O., Sidner, C., and Chong, C.-Y. Generation of a fundamental data set for hard/soft information fusion. In *Proc. Int. Conf. Information Fusion*, (Cologne, Germany), 2001.
10. Blasch, E. P., de Villiers, J. P., Pavin, G., Jousselme, A.-L., Costa, P. C. G., Laskey, K. B., and Zeigler, J. Use of the URREF towards information fusion accountability evaluation. In *Proc. Int. Conf. Information Fusion*, (Virtual), 2021.

Presidents' Memories

Presidents of professional societies such as the International Society of Information Fusion (ISIF) are, by definition, in a unique position. Their obvious overarching task is to try to advance the goals of the whole community, but they also must deal with routine chores, e.g., managing ISIF board meetings. Over the course of their tenure, they get to see it all. After leaving office, they are free to muse about their experience and their service, and especially about ISIF. These reflections may ripen (like a fine wine?) into unique perspectives. And, as with a fine wine, thoughtful perspectives are best served when they are shared with friends and colleagues. This section of *Perspectives* is intended to be a place for sharing these views.

JIM LLINAS (1998–1999)

1998 seems—and is—a long time ago. But the memories and honor of being selected as the first president of ISIF remain very present in my mind—what a humbling experience for this guy from Brooklyn. Among various fond memories is the spirit of collegiality of this society. As the society grew and the conferences became memorable events both professionally and socially (I will never forget the banquet in Florence), the society became evermore grounded in the professional sense, with the evolution of the *Journal of Advances in Information Fusion* (JAIF) and *Perspectives* publications. In reviewing the proceedings of the 1998 conference, I am staggered by the continuing productivity of our membership, as well as the diversity of applications and research thrusts. As my Joint Directors of Laboratories colleague Frank White once said, "Data fusion is a way of thinking", pointing to its far-ranging application potential. Yaakov Bar-Shalom has pointed to the need for societal growth, and I agree. One possible path is to expand into more systemic applications rather than focused algorithmic research; system issues are inherently multidimensional. Continuing the path to diversity in our leadership and the consequent new ideas that would come can be another path. "Fusion" remains such an important technology/area of study even in these days of the explosion of artificial intelligence (AI), etc., so I remain confident that the society's future path is secure.

YAAKOV BAR-SHALOM (2000, 2002)

During my terms as president of ISIF—2000 and 2002, the society's "toddler years"—we experienced growth and maturing. The International Conference on Information Fusion (FUSION) 2000, a.k.a. F2K, held in Paris, confirmed the international nature of ISIF. At this conference, I wore two hats—a top hat when I gave the State of the Society report and a French Navy sailor's hat (modified to indicate FUSION 2000) when giving the General Chair's report—both courtesy of Jean Dezert, who single-handedly organized this conference. In 2002, I had the privilege of announcing the launch of the JAIF. This journal, after some delays, had its first issue in 2006, under the most competent Editor-in-Chief Dale Blair. The special features of this journal were (and are) that it be free to members of ISIF and have no page limits or charges to authors. JAIF publishes only (thoroughly) peer-reviewed submissions with the same standards as the Institute of Electrical and Electronics Engineers (IEEE) *Transactions* collection. This is unlike another journal in the field, which while claiming to be peer reviewed, published several papers with absolutely no peer review, including one with the claim that the probabilities of detection and false alarm have to sum up to unity.

As a society, we still have one problem—our size is small. Since information fusion is used in numerous areas, many authors of papers with applications in specific areas prefer to present and publish them in conferences and journals of those areas. While the FUSION conferences have had sessions with applications in different areas, I would like to encourage our "core members" to solicit or organize more such sessions in order to increase our membership, which is automatic and at no cost to FUSION participants. More such sessions would lead to increased interaction with other areas, as well as mutual learning opportunities. To conclude, I wish all of our members (past, present, and future) continued success after 25 years of ISIF.

PRAMOD VARSHNEY (2001)

My data/information fusion journey started in 1981 after listening to the plenary lecture titled "Distributed Sensor Networks," given by R. E. Kahn of the Advanced Research Projects Agency, at the IEEE International Symposium on Information Theory. I started working on detection problems in a distributed setting. Several of my doctoral students wrote their dissertations on this topic. The goal was to develop theory for collaborative inference by a handful of sensors. The field matured sufficiently so that I completed my book, titled *Distributed Detection and Data Fusion*. I also guest edited a special issue of the *Proceedings of the IEEE* on data fusion in January 1997, which contained articles on different aspects of information fusion written by authorities in the field. This was followed by the founding of ISIF at a meeting that I remember vividly. I served on the founding ISIF Board of Directors (BoD). It was a privilege to serve as the third president of ISIF in 2001. The FUSION conference was held in Montreal that year. I could sense a lot of excitement in the emerging field of information fusion. Application domains were expanding beyond defense. With the emergence of low-cost wireless sensors, there was a lot of interest in wireless sensor networks that employed a very large number of sensors. These research efforts have subsequently led to the Internet of Things (IoT) and similar applications. Another topic that generated significant interest was fusion of hard and soft data.

With the ever-growing amount of data available in various application domains, attention has shifted to fusion via machine learning. This approach has resulted in a lot of success, but naturally, there are limitations. One major limitation is explainability, as a fusion engine based on machine learning is a black box and may not inherently provide the reasons behind the results that it yields. Other issues include inherent biases (potentially due to biased training data) and fairness. Information fusion methods need to ensure that they minimize the impact of the above issues. Another important area is decision making by human–machine teams. New fusion paradigms need to be developed that consider the characteristics of human behavior while fusing their input with machine data, with which the fusion community has more experience.

More recently, AI and machine learning (AI/ML) have given rise to fairly accurate language models and the ability to mimic human action and behavior, such as producing humanlike art and written material. Significant challenges to our fusion community are to coexist with AI-bots and to formulate research problems and challenges that go beyond the capabilities of modern AI/ML tools and advance the state of knowledge. Research questions include determination of whether a specific output is generated by a human or an AI-bot; who has the authority to make critical decisions, a human or an AI-bot; and who is responsible if something goes wrong due to decisions made by an AI-bot. Naturally, these questions are beyond just the technology that many of us are familiar with and lie in the realm of policy, social science, and law. I am sure our fusion community will rise to the challenges and continue to contribute to the advancement of society and human quality of life.

XIAO-RONG LI (2003)

Here is an interesting story from when I served as ISIF president. As Per Svensson and Johan Schubert wrote in their part of the article "25 Years of FUSION Conferences: Collection of Memories" about FUSION 2004 in Stockholm, Sweden, they had a concerted effort to make the review process more rigorous and to raise the standard for accepting papers. Indeed, as the chart for FUSION conference acceptance rate shows, its acceptance rate was significantly lower than that of every other FUSION conference that provided a reliable rate.

Actually, there were quite a lot of outcries then, including from ISIF board members, about how harsh and biased the review process was, because their papers were rejected. The outcry was so great that some radical guy—forgive me for not disclosing his name—even advocated a boycott of the conference; almost all board members who voiced their opinions were negative and said the board should (or at least consider whether to) intervene. But as the ISIF president then in charge of such issues, I insisted that we had heard only one side of the story. Then, at my request, Svensson explained the review process in detail; told me about his frustration with such a request at a time when they were so busy organizing the conference, probably because I didn't tell him about the crisis in order to avoid their overreaction; and gave me an "unnatural" explanation of the process. After receiving the explanation, I argued as follows: Being biased or not is for a process (just like an estimator) and is not easily judged by only a few resulting sample points that appear to be somewhat abnormal; since the process seems unbiased, I don't think we should do anything. So, that's the end of it.

In fact, personally there was more to me. Earlier that year, I accumulated six papers with my students and submitted them all in one batch in one envelope (via mail at that time!), but they were all returned without review for being a couple of days late. This left me with only one FUSION 2004 paper (coauthored with others), much fewer than other years. As a conference organizing team, to be successful, we all want to receive as many papers as possible, and so why did they do this? Shocked as I was, I learned the hard way that some Swedes, like many other northern Europeans, are not as flexible as Americans or southern Europeans. It appears to me that the colder the weather is, the more rigorous the people are likely to be. Thinking back, I am sure that this impression may have played a role in my reasoning toward the above decision, because to me, they are more likely to be somewhat more rigorous than biased, hence the outcries.

CHEE-YEE CHONG (2004)

When I went to the attorney's office in the summer of 1998 to start the paperwork for the incorporation of ISIF, I didn't expect the FUSION conferences to gain stature so quickly. So I was awed when the reception for FUSION 2004 was held in Stockholm City Hall, the location of the Nobel banquets. For the president's message in the FUSION 2004 proceedings, I wrote, "As I sat down to write this message, the biggest news in the internet world and here in Silicon Valley is the upcoming initial public offering (IPO) of Google, a service that many of us use daily. Google has made its name and a lot of money by rapidly searching the web for relevant data. However, the user still has to go through the search results to extract the information that he or she is looking for. Imagine the day when intelligent algorithms will automatically fuse all this data into useful information".

That day has arrived with chat generative pretrained transformer (ChatGPT) and other large language models that provide well-written responses. I also wrote, "At the other end of the problem spectrum, networks of inexpensive sensors are being deployed for civilian and military applications. All this data has to be fused to be useful". IoT is now ubiquitous in both civilian and military systems. I then wrote, "Between these two extremes of fusing textual data over the internet and fusing signals from small sensors, there are many other challenging and important fusion problems. The future for the information fusion community is certainly bright and exciting".

The future is now, and that presents challenges for ISIF. Information fusion is now an important part of many applications, not just traditional defense and aerospace systems. As an example, sensor fusion is a prerequisite for driver assistance and autonomous driving. Each area has its own conferences and journals that compete with FUSION and JAIF. ISIF has to find its niche in the expanded landscape of information fusion. At the same time, recent advances in machine learning provide opportunities for integration with traditional model-based fusion.

I have been fortunate and privileged to be involved with ISIF from the beginning. It is the most rewarding experience of my professional career. I have learned the latest advances in information and made many lifelong friends. I am confident that ISIF will meet the challenges and seize the opportunities.

ERIK BLASCH (2007)

I was honored to be the ISIF president in 2007, having been with ISIF since its inception. ISIF continues to extend professional content to the research community, industrial practice, and student education. Among the many discussions the ISIF BoD had in the 2000s, several proposals have come to pass, such as the JAIF, a professional website for members and the extended community, and awards inspiring the next generation that recognize past members. It is hoped that new members will engage and enhance the society service through working groups, paper submissions, conference leadership, and contemporary tutorials.

STEFANO CORALUPPI (2010)

I have been an ISIF member since 2000, ever since my wonderful first FUSION experience in Paris. I later served as ISIF president in 2010. These were both important times in my professional career: in 2000, I was soon to transition to a research position with the North Atlantic Treaty Organization in Italy, and in 2010 I transitioned back to a research position in industry in the United States. I find that it is at moments of career transition that professional societies play a particularly important role. For me, as for many of you, both ISIF and its much larger sister IEEE provide a professional home, a community of people with shared technical interests and distinct, complementary perspectives. Like any healthy community, a professional society offers unity, not uniformity: we learn from one another's approaches to similar challenges. An international technical family is particularly important for those of us working in the defense community, as our work often and invariably introduces limits to what can be shared. ISIF (like IEEE) helps to provide a broader perspective in which we recognize that our professional identity need not be tied exclusively to our employer. Furthermore, our professional identity and professional relationships span the full extent of our working life. In this sense, the importance of ISIF remains the same today as it was for me in 2000 and 2010. We must continue to offer our members a vibrant community while continuing to encourage excellence in JAIF, *Perspectives,* FUSION, and all ISIF endeavors. How we approach complex challenges, with a desire for personal and professional growth, while supporting the growth of those around us, is as important as the technical solutions that we develop.

What challenges does the future hold? I am perhaps less concerned than some of my fellow past presidents with how ISIF positions itself to capture parts of high-interest technical spaces, say, the growing world of data-driven, model-free inferencing. I am still working out for myself to what extent these methods represent true scientific progress. In the end, ISIF will grow according to the interest of its members, and that is how it should be. I am a bit more concerned about encouraging our younger research colleagues to step forward and to share their talents actively in ISIF, to the benefit of themselves and of all of us. There are so many leadership and service roles to be filled, including roles that don't exist until one thinks to propose them! A recent demonstration of this initiative is that of Felix Govaers, with his work as our new vice president of social media. So, as the Italians say, *Largo ai giovani!* "Make way for the young!" I look forward to learning and sharing with all of you for the next 25 years in the life of ISIF, and beyond.

JOACHIM BIERMANN (2011)

The motto of my presidency in 2011 was "To strengthen the role of ISIF as 'the' society in the information fusion community and to keep and improve the public awareness and interest. Therefore, we should: find a better balance between high- and low-level fusion; widen the scope of applications represented in the conference; and invite new areas of research where information fusion is a relevant factor". Much of this had been a general concern of the BoD for some time, which I shared and wanted to support. In particular, it was close to my heart to bring the topic of "high-level fusion" more into focus than before and to represent it in the annual conferences. However, this has only been achieved to a modest extent.

One aspect that had already become apparent was the BoD's endeavour to contribute by our activities to the further development of the subject area "information fusion". This was to be realised by encouraging and supporting the foundation and work of the so-called ISIF working groups. The Working Group on Multistatic Tracking already existed, and a new Working Group on Fusion Models and Frameworks had just started. In addition to such scientific impulses, the organisational work of the BoD should also be improved. The work of the BoD should be better structured and become more professional.

Among other things, a major contribution to this became possible from the fact that for the July 2011 board meeting, there were two proposals to host the 2013 FUSION conference, both of which were sufficiently qualified to be accepted. In order not to discourage any of the competitors, it was decided not only to decide on the organisation of the 2013 conference but also to accept the other application for 2014. This procedure of deciding on the venue and the organisation of the conference three years in advance became the established way of doing things and ensured that the previously usual short preparation time for the conference was extended by one year. This allowed for more relaxed planning, which benefited all parties involved.

Happy 25th birthday to ISIF as an organisation and promoter of the increasingly important topic of "information fusion". All good wishes for the future development of the society, in the hope that our common scientific work and goals may contribute to the good of all. May the Force be with you!

ROY STREIT (2012)

To be candid, possibly too candid, in 1998, when I first heard of the upcoming meeting on information fusion in Las Vegas, I was underwhelmed. It struck me as a continuation of the earlier Joint Service Data Fusion Symposia, and it was held in a location that offered little of interest to me. My views moderated somewhat over the next two years. After the Paris conference in 2000, I read the final program with interest and decided, finally, to submit papers to the Montreal conference in 2001. It was a good decision. After 2005, when I came to Metron (another good decision), I was able to support ISIF in new ways. I was even president in 2012 and cochaired two FUSION conferences. I feel fortunate to have been able to participate in the ISIF community for so many years. I have gained much, both professionally and personally. I especially value the many personal relationships that have grown over time—they would have been hard to sustain without shared commitments to common goals. Professionally, I gained new perspectives about the field that I might not have encountered otherwise, and I learned new methods. I have a growing sense of the importance of information fusion and the kind of role ISIF can play in the coming years.

As an organization, ISIF must do what all healthy organizations must do—broadcast to all who will listen that the field of information fusion is important in the modern world, encourage new talent to join in the search for responsible solutions, provide an intellectually safe forum to engage in debate, encourage members to take on leadership roles, and publish results that matter.

Those who work in information fusion are not blessed with an abundance of what I term "canonical models". One consequence is that the field is increasingly data driven. AI/ML systems produce "deep fakes" of many different kinds, and to me that suggests that a degree of order can be "discovered" in haphazardly gathered data sets. Discovered order can lead to unintended outcomes, e.g., the ugly ethical failures of AI/ML, and it can be used deliberately for nefarious purposes. Researchers who study effective ways to mitigate the defects of data-driven technologies will find a home in ISIF.

The gavel pictured was a gift from the local host organization involved in arranging FUSION 2013 in Istanbul. Murat Efe intended it to be passed from one ISIF President to the next. An ISIF tradition is born.

DARIN DUNHAM (2014–2015)

ISIF will continue to be challenged to be viable and relevant in our technology areas. There are scores of other conferences that try to fill a similar need, but ISIF is unique. We need to continue to market our unique combination of fusion research and technologies to a diverse set of current and future members. My best memories are all of the people from around the world that I have worked with in ISIF. Chairing the conference in Chicago in 2011 was fulfilling—so much so, that I am helping Terry Ogle chair the conference this year in Charleston. The board meeting in Salamanca was memorable due to its location (it felt like I was presiding over a court).

JEAN DEZERT (2016)

Taking over from Darin Dunham (ISIF president, 2015), I had the honor and privilege of serving ISIF as president during 2016. At that time, ISIF was already running smoothly and the FUSION 2016 conference in Heidelberg was perfectly organized by the team of Uwe Hanebeck and Wolfgang Koch, and I must, on behalf of ISIF, salute them. FUSION 2016 was a great success at all levels, and it enabled ISIF to end 2016 in good financial balance. Thus, at the end of 2016, I was able to hand over my hand (and ISIF gavel) to Luydmila Mihaylova, newly elected to ISIF presidency for 2017 and later reelected for 2018 as well. Maintaining a balanced budget for ISIF in 2016 was important, since the FUSION 2017 conference was scheduled for the first time in China and was to celebrate FUSION's 20th anniversary. Although Xiao-Rong Li and Roy Streit had already been successfully involved in ISIF for years, and the organizing team of Xi'an Jiaotong University was serious, ISIF could not estimate precisely the number of local participants in FUSION 2017, as well as its rate of international participants. Its success was therefore not guaranteed a priori, and ISIF was taking a real financial risk by sponsoring this conference. In 2016, ISIF was aware of the risk it was taking, and its good financial health protected it to a certain extent from the possible hazards of a loss linked to low participation in FUSION 2017. All the participants in this 2017 conference will remember, I believe, an excellent conference in terms of both the scientific content and the various actions to celebrate FUSION's 20th anniversary (see ISIF *Per-*

spectives on Information Fusion, vol. 2, No. 1, March 2019), despite the 2017 summer heat wave in Xi'an.

Despite the international geopolitical situation, global warming, and successive health crises, ISIF has, I believe, been able to adapt to the current context by maintaining, despite the difficulties, the organization of its annual international conference in virtual and hybrid modes. ISIF will continue to serve the scientific community through its various actions and media (*JAIF, Perspectives*, working groups, etc.). The current strong enthusiasm for AI techniques based on deep learning, as well as generative pretrained transformer techniques (such as ChatGPT and others) must be used with great caution and, in particular, for the fusion of information, since the results produced depend strongly on the quantity, diversity, and quality of training data used. These AI techniques have no capacity for imagination and creativity, which remain specific to human understanding. Recent small tests conducted with ChatGPT to assess its ability to produce a reliable and correct answer to a relatively simple mathematical problem clearly show the current limits of AI and fortunately augur well for good days and good decades for good mathematicians and researchers.

LYUDMILA MIHAYLOVA (2017–2018)

After Jean Dezert, I had the great pleasure and honor to become the ISIF president for 2017, and later this was extended with a second term for 2018. I am grateful for this opportunity to lead in this period and for the many opportunities for insightful discussions, collaboration, and knowledge exchange that led to cocreation. This was a period when the ISIF community was flourishing and continued its expansion. The two FUSION conferences—in Xi'an, China, 2017, and in Cambridge, United Kingdom, 2018—were remarkable in different ways and witnessed a significant interest in nearly all sensor data fusion areas at both methodological and application levels. This was a time of significant expansion of the community and collaborations with industry. Irrespective of the hot weather in Xi'an, we had excellent discussions and analysis of the historic ISIF activities—given by Xiao-Rong Li with the support of the local committee.

This was a period when we were rethinking particle filtering methods and expanding them to high dimensional spaces. We celebrated 25 years with particles and other random point methods, and we looked toward fusion of multiband images, AI, and machine learning methods, such as Gaussian processes and variational inference, distributed data-driven methods, transfer learning, and other learning methods.

AI-enabled fusion for federated environments; trust, uncertainty, and deception in information fusion; and big data fusion were some of the topics that dominated at FUSION 2018 in Cambridge, and that indicates a starting trend toward AI topics and model-free methods. I personally would like to see an expansion of the ISIF toward AI, machine learning areas, and autonomous systems. They are subject to rapid development, and the achievements can make a big difference in our lives. Together with the beneficial aspects, ethics, data, and privacy protection are only some of the many questions that need careful answers.

At the end of 2018, I handed over the ISIF gavel to Paolo Costa, who continued the mission to unite, inspire, and support ISIF in all its activities. I wish all ISIF members success, and I am looking forward to continuing working with you.

PAULO COSTA (2019–2020)

More than just choosing a person, the election of a new ISIF president by the BoD is really about investing on a vision. In 2019, the ISIF BoD chose a vision focused on preparing its infrastructure to the natural next steps for the society, centering on making is processes more efficient and transparent, defining our strategic goals and using them as the main criteria to allocate our budget, doubling down on our strengths while tackling the difficult issues we had, and leveraging the amazing work already done by our predecessors while paving the way for those who would follow.

In that same year, the ISIF BoD and its executive committee did just that, having more frequent but time-constrained meetings in which we implemented new ideas while tackling difficult issues from the past. It was a period in which those not involved hardly noticed any changes, as most of it was foundational groundwork that did not bring flashy, attention-grabbing results. Yet it worked well, so the ISIF BoD decided to continue investing in that vision for 2020. We entered the year featuring the same leadership, with all of us being rather excited with the future prospects for a stronger and more efficient society.

Then coronavirus disease 2019 (COVID-19) struck… As much as anyone else inhabiting planet Earth, we were suddenly faced with hard, future-defining decisions for which no previous experience was available to serve as a guide. The months of February to April were by far the most difficult ones, as the 3-year preparation cycle of FUSION 2020 was at its last phase and suddenly had to be drastically changed. Many options surfaced, including canceling the conference; postponing to a later, hard-to-predict date; choosing some of the many untested virtual solutions; and others that reflected the chaos everyone was facing in their professional and personal lives.

Tough times require resolute leadership and no hesitation in making risky decisions that would certainly not please everyone but must be made regardless—status quo was not an option. We had various meetings involving the ISIF BoD, its executive committee, and the amazing South Africa 2020 team, eventually opting for a virtual format for which details had to be defined in a rush as well. As we already know, FUSION 2020 was considered successful by most of the same metrics we use to evaluate regular conferences while being unequivocally extremely successful given the conditions and the expectations we had during the difficult months preceding it. It also set the society's expectations to how FUSION conferences should be run in this new world of ubiquitous virtual meetings, travel restrictions, and other factors that clearly show we are in a completely different environment than the one in which we ran 22 successful FUSION conferences. From a rearview mirror perspective, it might be easy to underestimate the challenges and tensions of that period, but those who lived through it know that we, as a society, not only dodged a wrecking ball but also left the society in a comfortable position to continue its inspiring path.

SIMON MASKELL (2021–2022)

As part of becoming a new member of ISIF's BoD in 2019, I wrote a mission statement. I focused on three issues: making ISIF's finances more transparent, encouraging practices (e.g., sharing of code) that will ensure ISIF's members are justifiably respected by other communities, and working to put a stop to unethical conduct (which felt like a painful scar on the community in 2018). I hoped that delivering on that mandate would become easier when I became president of ISIF in 2021.

Much of my work as president became focused on improving operational aspects of the ISIF BoD's activity. I don't feel I had the impact I would have liked on the respect that ISIF's members receive from each other or other communities: that is not to say it isn't great to see the Stone Soup project maturing, or that I don't welcome the return to a more friendly ISIF, just that I think these things happened while I was president, not because I was president. However, I did manage to present slides as part of the president's welcome to the FUSION 2021 conference that explained ISIF's financial position. As I explained then, information about ISIF's finances is freely available from http://guidestar.org: looking now, I can see that the most recent Form 990 I can download was submitted by ISIF's treasurer on 22 September 2020 and stated that ISIF had $622,284 in the bank.

I have a distinct memory of presenting those slides remotely from the comfort of my desk in Liverpool: FUSION 2021 took place in South Africa and had been postponed from July to take place in early November 2021. Even so, physical attendance was significantly reduced as a result of concerns related to COVID-19. In fact, the Omicron variant emerged from South Africa in late November 2021, and that was arguably the start of the end of the pandemic.

I also remember seeing the physical attendees relishing the experience of the first FUSION conference to take place (physically) in Africa: I particularly remember being a remote predinner attendee and thinking that a banquet is not a good experience to join via Zoom! Of course, the organisational team for FUSION 2021 were the same people who had previously found themselves having to work so hard to enable FUSION 2020 to happen in the height of the pandemic: COVID-19 had forced FUSION 2020 to be a fully remote experience. The team for FUSION 2021 wanted to, and deserved the right to, run an in-person event. They did succeed to make the hybrid event a success but also solidified the ISIF BoD's view that future FUSION conferences should be in-person events, if at all possible.

FUSION 2022 was then the overt return to normality (albeit in Linköping!). There were fewer papers submitted than were anticipated, but the number of attendees was close to pre-COVID levels, and the buzz at the venue was palpable. The community seemed at ease with itself and the world. ISIF had navigated COVID-19 and was now ready for the future.

Tributes to Fusion Minds

Communities form around commitments to shared interests, goals, values, and more. In time, given need and opportunity, they develop institutions to support the community as a whole and to encourage its members to grow and develop. All of this requires significant investments of time and effort by many people from diverse backgrounds, all working together, each doing what they can, when they are able. The International Society of Information Fusion (ISIF) would not have endured for 25 years without their dedication. We stand on their shoulders, and whether they be giants or not, we are bound by a duty to acknowledge their freely given gifts of time and labor, and yes a kind of love, to our field and our community. Some are no longer with us, sadly, and we miss them in many ways. This section of *Perspectives* is intended to be a tribute to, and an acknowledgement of, the contributions they made to the ISIF community.

CHRIS BOWMAN (1948–2023)

Dr. Christopher Bowman unexpectedly passed away on February 11, 2023, due to complications following a medical emergency. We honor his memory and his contributions to the data fusion community.

Chris is survived by his loving wife, Julie Hunter Bowman, and nine children: Aaron, Mark, Sarah, Adam, Laura, Jeff, Quinn, Rebecca, and Bailee. In 1971, Chris joined the Church of Jesus Christ of Latter-Day Saints and served in the church in many different callings. He was an avid racquet sport player, including badminton, tennis, platform or paddle tennis, and pickleball. He loved singing along to old musicals. Despite his many passions and professional dedication, spending time with his family was his greatest joy in life. Chris was wonderfully optimistic. His positive attitude and open-minded appreciation of people and of life were immediately evident to everyone who met him.

Chris Bowman (right) with Alan Steinberg at Workshop on Critical Issues in Information Fusion in Beaver Hollow, NY, 2008 (courtesy of Alan Steinberg).

Chris graduated in 1966 from Garden Grove High School in Garden Grove, CA. In 1970, he completed his bachelor's degree from the University of California, Riverside, with a double major in mathematics and physics. In 1973, he completed his master's degree in mathematics, followed by his Ph.D. in mathematics in 1977 from the University of California, Irvine.

Chris's impact on data fusion, neural networks, and systems engineering has been extensive and profound. As Edward Waltz noted, "Chris Bowman is a legend… truly a pioneer in the fusion field and his depth of understanding (from concept to code) was impressive". Few scientists have such breadth; many can formulate and offer design and algorithmic concepts, and others are strong in implementation, but few have such a breath of skills, vision, and innovation.

Chris was a major pioneer in data fusion, having been the first to define the fundamental functional and architectural design principles of data fusion system engineering. He was first to extend these concepts to resource management processes, thereby enabling the cost-efficient development and operation of responsive information exploitation systems. He is known internationally for his development of the widely used data fusion and resource management dual node network technical architecture that supports affordable system design, development, and comparative analyses.

Chris was more than a theoretician; he had a stellar career applying these theoretical and engineering principles in designing and developing innovative, advanced data fusion and resource management systems for integrated avionics, cooperative tactical operations, missile defense and space situation awareness, and diverse other applications.

From 1978 to 1991, he managed the data fusion and neural networks programs at Verac, Inc., where his team delivered software for multispectral integration, weapon/sensor management, and nonlinear adaptive control with applications in tactical avionics, missile defense, and tactical surveillance. From 1992 to 1995, Chris supported Ball Aerospace in Strategic Defense Initiative efforts.

In 1995, Chris founded and led Data Fusion and Neural Networks, LLC (DF&NN) as president. Prominently, he led DF&NN's ground-breaking development operational software for remote diagnostics of satellite health and operations. He led the Information Fusion Working Group that developed a data fusion and resource management roadmap for the Air Force Research Laboratory. He served on the Signals Intelligence

Science and Technology Advisory Board to oversee development of a large data fusion system, and he was a member of the Hercules Blue Technology Team in support of the Ballistic Missile Defense Organization. Under Project Correlation for the U.S. Air Force Space Warfare Center in the 1990s, Chris was instrumental in developing a set of data fusion engineering guidelines and an evaluation of several major combat information/intelligence correlation systems.

Most recently, he served as cotechnical lead for the U.S. Department of Defense (DoD) high-level data fusion project, developing advanced concepts and methods for exploiting high-level data fusion in a critical tactical application.

In 2018, Chris was honored by the U.S. Military Sensing Symposium Specialty Committee on Sensor and Data Fusion with the prestigious Joseph Mignogna award for his contribution to the development and use of data fusion.

We have lost a great person and a very good one.

—Alan Steinberg and James Llinas

CRAIG CARTHEL (1964–2022)

Craig Carthel, a dear friend and professional colleague, died unexpected in July 2022, shortly before I was to attend FUSION in Linköping. This was a devastating loss, for me as for all of Craig's family and friends. On the professional side, Craig and I collaborated closely in the development of advanced data fusion solutions, with a particular focus on multiple hypothesis tracking (MHT) methods. Our collaboration spanned multiple professional employers and 24 years of joint efforts. He was admired by so many for his technical brilliance (in both mathematics and software solutions), his kindness, and his good humor. In 2006, we hosted FUSION in Florence, where Craig served as technical chair. Beyond his leadership and participation in FUSION, those who supported the ISIF Multistatic Tracking Working Group will remember his contributions and collegiality over many years. In August, during an online tribute that was held in remembrance of Craig, Peter Willett remarked that, when faced with a tough tracking problem, "When all else failed, that was when we turned to Craig". Indeed, how many times I experienced that over the years! I am so fortunate to have had Craig as a friend and colleague for many years.

—Stefano Coraluppi

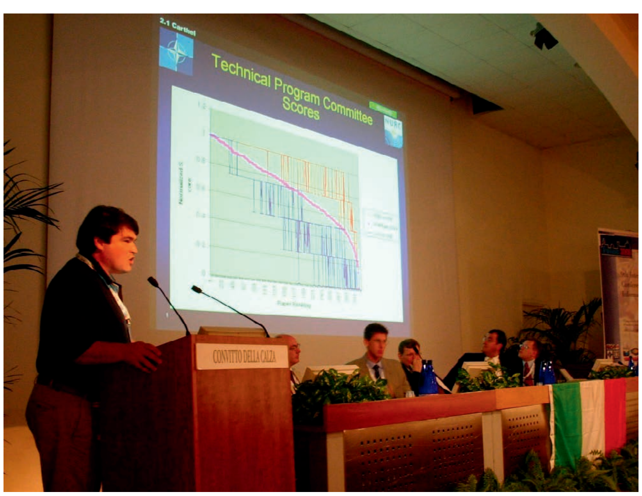

Craig Carthel explaining the paper review and acceptance process during FUSION 2006 in Florence, Italy (courtesy of Stefano Coraluppi).

SAMUEL BLACKMAN (1938–2019)

Sam Blackman is famous for creating industrial strength MHT–interacting multiple model (IMM) software that is used in dozens of important radar, passive optics, and multisensor data fusion systems. It works robustly in the real world and runs in real time on low size, weight, and power computers for stressing multiple target scenarios. It is clearly the algorithm of choice. Sam emphasized how lucky he was to have Bob Dempster help him designing, coding, testing, and tuning the algorithm. What a team! Sam and Bob decided to create this wonderful software all by themselves, on their own time, without any internal or customer research and development support. I vividly recall one big contract that we won specifically because the customer insisted that they wanted Sam's MHT-IMM for data fusion. The customer was thrilled to meet Sam and learn more about his work. Sam was enormously generous with sharing his ideas and time to mentor other engineers. After Sam became too sick to drive into work, many engineers would visit the guru at his house to drink from the fountain of wisdom. Sam loved it. Sam's books and papers are treasured by engineers looking for practical tracking algorithms described clearly by a real expert. But Sam was pretty feisty when he thought that you had grasped for undeserved laurels. This got him in hot water once, but the worldwide tracking community pitched in to help him out.

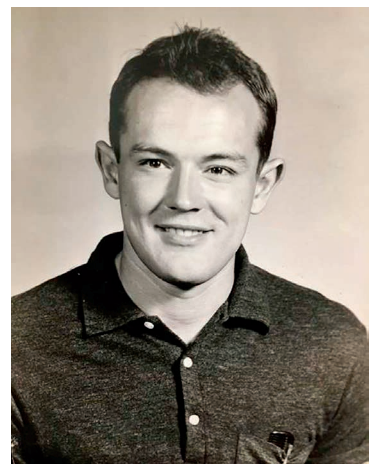

Sam Blackman, the Wizard of multiple hypothesis tracking dreaming about playing basketball (courtesy of Dale Blair).

—Fred Daum

RUDOLF KALMAN (1930–2016)

"Kalman filter is a beacon in the vastness of non-linear, non-Gaussian phenomena". A historical perspective on the role of the Kalman filter in aerospace is reported in [1]. My appreciation of the contribution given to humanity by Professor Rudolph Emil Kalman's inventions is described in [2]. On 12 September 2016, Professor Sergio Bittanti presented a posthumous

Rudy Kalman (right) with Alfonso Farina in Gainesville, Florida, April 2016 (courtesy of Alfonso Farina).

Ph.D. honoris causa in information technology to Mrs. Dina Kalman in a ceremony held at the Department of Electronics, Information Technology and Bioengineering at the Polytechnic of Milan [3]. I was present at the ceremony together with Professor Luigi Chisci of the University of Florence.

In recent years, I greatly valued my friendship with Rudy and Dina Kalman and our periodic meetings at their residences in Zurich and Gainesville, FL. We used to consult and browse Rudy's extensive library spanning from technical books—in several languages—to his many famous publications. I also had the chance to hear Rudy's explanations of his investigations on long-standing electrical network synthesis and look at his authoritatively handwritten notes.

It was quite an experience to jointly consult on philosophical books like *De Consolatione Philosophiae* by Severino Boezio [4], of which Rudy had a precious English copy together with the original Latin version. He also introduced me to "The Relation of Sense-Data to Physics" by Bertrand Russell, and in particular, he pointed out to me the last paragraph of section 1 at page 114 [5], which I have transcribed: "Thus, if physics is to be verifiable we are faced with the following problem: Physics exhibits sense-data as a functions of physical objects, but verification is only possible if physical objects can be exhibited as functions of sense-data. We have therefore to solve the equations giving sense-data in terms of physical objects, so as to make them instead give physical objects in terms of sense-data".

I keep it in my study room not far from my gaze so that I can ponder it time to time. In the summer, our talks were intertwined with a pleasant lunch or dinner in their garden at the Zurich house or in restaurants in the fresh hills around.

REFERENCES

1. Grewal, M. S., and Andrews, A. P. Applications of Kalman filtering in aerospace 1960 to the present. *IEEE Control Systems Magazine* (June 2010), 69–78.
2. Farina, A. July 27, 2016: Celebrating the professional life of Professor Rudolf Emil Kalman. *IEEE A&E Systems Magazine,* Vol. 31, 12 (December 2016), 47–49.
3. Bittanti, S. Ph.D. honoris causa in information technology to R. E. Kalman. Presented at Department of Electronics, Information Technology and Bioengineering, Politecnico di Milano, Milan, Italy, Sept. 12, 2016. [Online] https://bittanti.faculty.polimi.it/rudolf%20Kalman.html.
4. Boëthius, A. M. T. S. *De Consolatione Philosophiae*. 524 dC, in Latin.
5. Russell, B. The relation of sense-data to physics. *Scientia,* Vol. 16, 16 (1914), 113–141. Available: https://www.princeton.edu/~hhalvors/teaching/phi520_s2015/Russell_SenseDataPhysics.pdf.

—Alfonso Farina

OLIVER DRUMMOND (1928–2016)

Oliver Drummond was a unique character in the target tracking community, which is full of characters (e.g., Dale Blair and Yaakov Bar-Shalom). I had the privilege of knowing Oliver for many years and actually worked for him for a few years. There are a couple of things that stand out in my memories of Oliver. First, he was obsessed with avoiding being constantly connected to the internet. I think he used dial-up internet and would disconnect the phone line from the modem when he wasn't using it. In addition, because he only used dial-up internet, he couldn't download updates the Ballistic Missile Defense Benchmark simulation. So, every couple of months, I would download the updates and mail a CD to him. Second, Oliver was a classic late-night worker. It was best not to call Oliver until after 11 am, California time! On a serious note, Oliver was a very nice person, who always had time to explain something. We all remember some of his famous sayings, and still quote them, and I think all of us truly miss him as a colleague, researcher, and friend.

Oliver Drummond as a young naval officer with a burning ambition to make tracking algorithms work in the real world (courtesy of Dale Blair).

—Darin Dunham

DAVID L. HALL (1946–2015)

Dr. David L. Hall certainly made his mark on the multisensor data fusion and information science community. An Air Force veteran, member of industry, and finally Dean of the College of Information Sciences and Technology at the Pennsylvania State University, where he established the Center for Network Centric Cognition and Information Fusion, Dr. Hall was a teacher and researcher known for integrity and vision. His contributions to information fusion include multiple books, over 200 papers, and key input to the development of the Joint Directors of Laboratories (JDL) Data Fusion Model. His vision for fusion extended beyond defense and intelligence to

David Hall with his daughters, Dr. Sonya A. Hall McMullen and Dr. Cristin M. Hall (courtesy of Sonya McMullen).

the realm of all things possible, such as applications for smart and safe transportation to health care, to smart homes, and many others. Dr. Hall was quick to address the potential pitfalls of fusion and technology as the coauthor of the publication *Dirty Secrets of Multisensor Data Fusion* that provides enduring guidance for data fusion system development. Dr. Hall was an Institute of Electrical and Electronics Engineers fellow for his contributions to data fusion, and he was awarded the U.S. DoD Joseph Mignogna national data fusion career award for his leadership in data fusion. Although he was a huge proponent of technology, Dr. Hall's philosophy was one of pursuing technology to enhance humanity and human capability for information processing and decision making. Most of all, Dr. Hall would want to be remembered as a son, husband, father, and twin.

—Jim Llinas and Sonya Hall McMullen

ROBERT LYNCH (1960–2015)

Thanks to *Perspectives* for letting me write a few words about Robert S. (Bob) Lynch. Bob worked with me for his Ph.D. (please read his stuff, it's good) while holding down a full-time research job at the Naval Undersea Warfare Center, where he was responsible for all kinds of data fusion products.

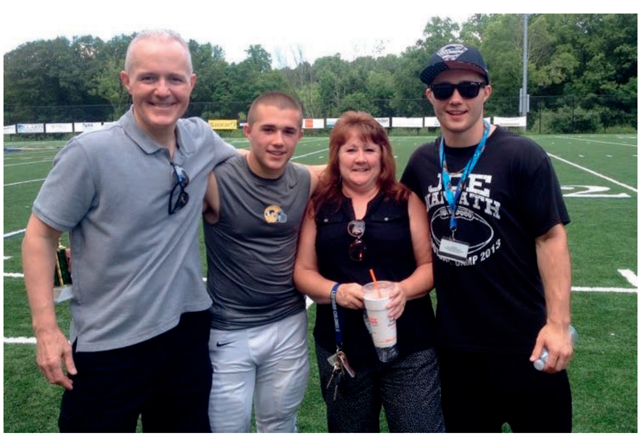

Bob Lynch and his family (courtesy of Peter Willett).

Bob loved and lived research. But—second, of course, to his family—his passion was, I think, our ISIF. Bob was always involved in our conferences, and was general chair with Chee-Yee Chong for the successful FUSION 2009 in Seattle. Let's add that our *Journal of Advances in Information Fusion* is now "indexed" (i.e., has an impact factor)—according to my recollection, that can be traced to Bob and his persistence.

Bob was taken from us far too early, in 2015, from a horrible disease that he had been fighting—optimistically and with amazing humor—for more than 5 years. Almost up to the end, Bob was still involved, driving himself (very unwisely) to FUSION in Washington, DC. Bob is survived by his wife Sherry and sons Bobby and Ryan. Bob's memory is honored by our society with the endowed ISIF Robert Lynch Award for Distinguished Service. I miss him.

—Peter Willett

OTTO KESSLER (1942–2015)

Otto Kessler began his career at Naval Air Warfare Center Warminster. He developed advanced radars and sensors for the F14 and the P3. He also sought to fuse output from an aircraft's onboard sensors to improve mission outcome. Frustrated by the lack of investment in fusion, he became a program manager for Office of Naval Research in advanced sensors and fusion. He sponsored a Navy fusion meeting in 1983, and when the JDL was in its infancy in 1985, he became the first additional member (fourth overall) of the JDL Data Fusion Subpanel (DFS). Otto became a vital contributor, and when the DFS decided a symposium was needed, it was his dogged pursuit of funding and support that not only made the first symposium possible but also built it into an annual event.

Otto Kessler (right) and his wife Mary Ann at FUSION 2006 in Florence, Italy (courtesy of Frank White).

Otto was a talented engineer committed to the scientific method. He was a pugnacious believer in making decisions based on facts and careful analysis. Although his intensity could unnerve people, he had a warm, caring, and loyal personality. The JDL DFS was a team, leveraging the talents and vision of all members. Otto, with his brilliant mind and long experience, was an essential member of that team.

—Franklin White

PIERRE VALIN (1949–2014)

When I first met Pierre, he was working for Lockheed Martin Canada in Montreal. Later, he joined as a colleague at Defence Research and Development Canada, Valcartier, and pursued his research on target recognition with evidence theory as the mathematical framework for uncertainty representation. He was a brilliant and unpretentious researcher. Beyond his academic pursuits, Pierre was always eager to connect with others and build relationships. He was dedicated to ISIF, for which he served on the Board of Directors for several years and as President in 2006.

Pierre Valin enjoying a piece of dolce vita in Florence, Italy during FUSION 2006 (courtesy of Jean Dezert).

His contributions to international efforts in information fusion were unparalleled, and he was best known for his role as vice president of membership for ISIF. He established and maintained the society's membership database, analyzed conference participation, and worked tirelessly to manually collect data from past conferences' attendance. The uncounted numbers of hours spent building this colored Excel file provided ISIF with a solid basis of memory that grew and up to now can be queried.

He left too soon, in 2014, before he had the chance to receive in 2015 in Washington, D.C. the *IEEE Aerospace and Electronic Systems Magazine* award for the best paper in 2012.

—Anne-Laure Jousselme

JEAN-PIERRE LE CADRE (1953–2009)

Jean-Pierre Le Cadre was a friend and a brilliant researcher in signal processing and data fusion. He passed away too soon, in July 2009. He was well respected for his high professional standards, integrity, and contributions to the ISIF and the FUSION conferences over many years. Since 2011, the ISIF has recognized his contribution by naming the best paper award after him. Jean-Pierre made significant scientific contributions in the field of antenna signal processing, particularly in the area of passive sensor target tracking. He was a respected and dedicated mentor to his Ph.D. students, always willing to share his knowledge and opinions with a high sense of responsibility toward them. He will be remembered for his passionate interest in the hard mathematical problems of the field and his interest in problems of human society. The writings of Jean-Pierre will continue to live by themselves, and through them Jean-Pierre will remain with us: www.irisa.fr/vista/Publis/Auteur/Jean-Pierre.LeCadre.english.html.

Jean-Pierre Le Cadre in a collegial exchange at FUSION 2000 in Paris, France (courtesy of Claude Jauffrey).

—Claude Jauffrey

PHILIPPE SMETS (1938–2005)

Phillipe Smets was a highly recognised and universally respected scientist who made significant contributions to the field of uncertainty modelling and reasoning under uncertainty. He is best known for the transferrable belief model (TBM), a model for the representation of quantified beliefs, as a subjective and nonprobabilistic interpretation of the Dempster-Shafer theory of evidence. The TBM is based on the assumption that beliefs manifest themselves at two mental levels: the *credal* level, where beliefs are expressed and combined, and the *pignistic* level, where decisions are made. The TBM is equipped with many concepts and tools for handling belief functions, such as the conjunctive combination rule (i.e., the unnormalized Dempster rule), the refinement and specialisation of vacuous and ballooning extensions, the generalised Bayesian theorem, and the pignistic transform. Phillipe was also involved in the development of algorithms for fast computations (the fast Möbius transform, matrix calculus for belief functions, and algorithms for evidential networks), a comparison of TBM to alternative approaches to uncertainty reasoning (such as imprecise probability, random sets, possibility theory, default reasoning, and modal logics), and practical applications of TBM in medicine and engineering.

Philippe Smets, shining under uncertainty (from IEEE).

I had a great pleasure and privilege to be mentored by Phillipe during my study leave at Artificial Intelligence Research Laboratory of the Université Libre de Bruxelles in 2003–2004. He retired in 1999, so we spent many hours working together at his home. Philippe was not only a man of great eminence—he was also a charismatic, visionary, and kind person.

—Branko Ristic

GÜNTER VAN KEUK (1939–2003)

Twenty years ago, on October 17, 2003, the theoretical physicist Dr. Günter van Keuk succumbed to cancer. Since 1965, he had created essential methodological foundations of sensor data fusion and sensor control for the armed forces of the Federal Republic of Germany, the Bundeswehr. His scientific home was the Forschungsinstitut für Funk und Mathematik (FFM, Research Institute for Radio and Mathematics), founded in 1963, whose successor, the Fraunhofer Institute for Information Processing, Communication and Ergonomics (FKIE), can look back on 60 years in 2023. As an outstanding scientist, van Keuk founded the Sensor Data Processing and Control Methods Department in 1975, left his mark on it until 2001, and gave it international recognition. It was the nucleus of the Sensor Data and Information Fusion Department of FKIE.

Günter Karl Friedrich van Keuk as a young theoretical physicist, laying the foundations of data fusion and sensor management in Germany, 1973 (courtesy of Lars van Keuk).

A sense of their own history characterizes mature scientific communities. In this spirit, I dedicated my paper to my esteemed mentor for the special session at FUSION 2018 that recalled 40 years of MHT [1]. Two years after West Germany's entry into the North Atlantic Treaty Organization in 1955, the Society for the Promotion of Astrophysical Research began research in the interest of national defense. In 1963, the FFM became a member of this society, which was renamed Forschungsgesellschaft für Angewandte Naturwissenschaften (Research Society for Applied Natural Science) in 1975 until it was absorbed into the Fraunhofer Society for the Promotion of Applied Research in 2009.

How did an applied data fusion problem lead to the creation of the FFM? On the initiative of physicist Paul Kotowski (1904–1971), Telefunken, a predecessor of today's Hensoldt AG, and of high-frequency engineer Leo Brandt (1908–1971), then state secretary in the North Rhine–Westphalian Ministry of Economics and Transport, mathematician Wolfgang Haack (1902–1994), at the Technical University of Berlin, began an investigation in 1957 of the "use of computers in air traffic control". As early as 1959, Haack and his colleagues presented their first results, which can be seen as the beginning of computer-based, networked information fusion in Germany [2]. van Keuk, then a young physicist from the University of Hamburg, joined FFM in 1965. He was the Ph.D. student of Harry Lehmann (1924–1998), a pioneer of quantum field theory, and of Lothar Collatz (1910–1990), cofounder of numerical mathematics in Germany.

van Keuk was among the first who proposed and demonstrated a sequential track initiation scheme based on an optimal criterion related to state estimates. In this context, he developed a performance prediction model for phased-array radar, which has been called the van Keuk equation in the tracking literature. In Sam Blackman's monumental 1999 handbook *Design and Analysis of Modern Tracking Systems,* many of van Keuk's papers are referenced, indicating his growing international reputation.

REFERENCES

1. Koch, W. On anti-symmetry in multiple target tracking. Presented at the 21st International Conference on Information Fusion (FUSION), Cambridge, United Kingdom, 2018.
2. Haack, W., and Hildebrandt, W. "Die arbeitsvorgänge einer elektronischen rechenanlage für den flugsicherungsdienst, im besonderen die erkennung von kollisionsgefahren. *Telefunkenzeitung,* Vo. 32, 126 (1959), S. 3–10.

—Wolfgang Koch

FUSION 2024
8-11 July 2024, Venice, Italy

FUSION
VENICE | JULY 2024

Paper Submission Guidelines

All papers must be original and not simultaneously submitted to another journal or conference. Full Papers (between 5 and 8 pages) will be reviewed for acceptance in one or more of the following topics, typically involving some aspect of information, data or sensor fusion. Papers must conform to specifications in templates provided by IEEE at https://www.ieee.org/conferences/publishing/templates.html to be considered for acceptance and publication.

Topics of Interest

The topics below are to be read within the context of sensor fusion and/or data fusion and/or information fusion.

Theory and Representation:

Probability theory, Bayesian inference, argumentation, Dempster-Shafer theory, possibility and fuzzy set theory, rough sets, logic fusion, preference aggregation, decision theory, random sets, finite point processes and others.

Algorithms:

Cognitive methods, signal processing and localisation, recognition, classification, identification, nonlinear filtering, data association, tracking, prediction, situation/impact assessment, alignment and registration, pattern/behavioural analysis, image fusion, fusion architectures, resource management, machine learning and artificial intelligence, topic modelling, natural language processing, contextual adaptation, anomaly/change detection.

Applications:

Soft-hard fusion, autonomous systems, defence/security, robotics, intelligent transportation, mining/manufacturing, wireless sensor networks, economics, finance, fintech, environmental monitoring, medical care/e-health, bioinformatics, radio astronomy, critical infrastructure protection, condition monitoring, precision agriculture, video streaming, streaming and sketching and other emerging applications.

Methods/tools:

Sequential inference, data mining, graph analysis, ontologies/semantics, modelling/realisation/evaluation, target/sensor modelling, benchmarks/testbeds, trust in fusion systems, computational methods, cloud/edge computing/fusion, fusion performance.

Important Dates

February 1, 2024: Proposals for special sessions and demonstrations (Notification of Acceptance Feb. 15)
March 1, 2024: Full paper submission (Notification of Acceptance May 1)
March 1, 2024: Proposals for tutorials (Notification of acceptance March 15)
June 1, 2024: Final paper submission

Publication

FUSION 2024 proceedings will be published by the IEEE in IEEEXplore and will be Scopus and WoS indexed.

Contact

Chairs:
chairs@fusion2024.org

Tutorial submissions: tutorials@fusion2024.org
Special sessions: specialsessions@fusion2024.org
Program chairs: program@fusion2024.org
Technical chairs: technical@fusion2024.org
Awards chairs: awards@fusion2024.org

27th International Conference on Information Fusion

For more information on past FUSION Conferences, please visit: www.isif.org

ISIF IEEE

Book Review

Claire Laudy
Thales
Palaiseau, France
claire.laudy@thalesgroup.com

Misinformation and Disinformation—Detecting Fakes with the Eye and AI

Victoria L. Rubin

Springer International Publishing, 2022
ISBN 978-3-030-95655-4, $69.99 USD, Hardcover: 289 pages

INTRODUCTION

"How do we detect, deter and prevent the spread of mis- and disinformation with the human eye and AI?" This is the question Victoria L. Rubin tries to answer in this book. She provides a large, detailed, and complete overview of the question, divided into two parts. Part 1 of the book focuses on the human interaction with information in order to understand the nature of deception and how the human mind conceives it and falls for it. In Part 2, Rubin explores how the theoretical knowledge described in Part 1 can be applied to develop automated artificial intelligence (AI)–based fake detection systems.

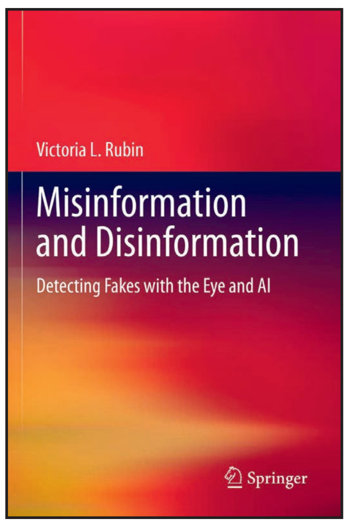

Rubin is multilingual, passionate about languages, and especially fascinated how language is used under challenging circumstances. As such, she has been particularly interested in studying how "lying and deception may be distinctly cultural, yet universal in the sense of their relevance to human condition". Through this book, she offers an overview of over 10 years of her studies of natural language processing (NLP) in the LiT.RL Lab [1]. The book combines aspects of previous publications, adds important details, and puts previous work into perspective, offering a framework for future research and development on fakes and deception understanding and detection.

The book is intended for a broad spectrum of readers. The author makes sure to present the different theories in a simple manner. It is thus intended for people dealing with large amounts of information and online presence who are looking for a primer on deception research. It is also intended for programmers and information retrieval experts whose aim is to develop fake detection systems.

PART 1: HUMAN NATURE OF DECEPTION AND PERCEPTION OF TRUTH

CHAPTER 1. THE PROBLEM OF MISINFORMATION AND DISINFORMATION ONLINE

Chapter 1 opens with useful definitions of the concepts of infodemic and infodemiology, as well as the distinction between mis- and disinformation. This chapter gives the reader the concepts necessary to understand mis- and disinformation and enjoy the following chapters. To define the concepts of infodemic and infodemiology, the author translates the triangle model for classical disease from epidemiology to the context of digital communication, identifying the three factors that allow the spread of diseases: compromised hosts, virulent pathogens, and conductive environments. This infodemiological model identifies the three interacting causal factors responsible for the spread of mis- and disinformation: automation, education, and regulation.

Here, the author sets the scene for the major importance of the subject she studies by reporting about the status of society about infodemic and mis- and disinformation. It appears there is a consensus among various organizations on the importance of the problem and the urge to tackle it. She states that assistance from AI to detect deception and fakes is inevitable and explains that the way we "accumulate knowledge from the past and the newest technological advancements can be combined […] to bring the current online infodemic under control". The author also makes an inventory of the different types of fakes that can be identified with AI. She concludes that even with AI-based solutions, there will still be a need for a human in the loop in the detection and management of fakes.

CHAPTER 2. PSYCHOLOGY OF MISINFORMATION AND LANGUAGE OF DECEIT

This chapter presents deception as an uncooperative communicative behavior. It describes the motivations for deception and disinformation. The different types of deceptions are studied, and a useful alignment is proposed for the varieties of deception described in various taxonomies in an earlier work in the field of psychology and communication research [2].

The author then enumerates the reasons deception is effective for us as humans. For instance, studies showed that the more information is repeated, the more it seems true. Furthermore, the less cognitive effort that is needed to understand the information, the more fluent (i.e., easy to integrate as one's own) this information is, and finally the more true it looks [2].

The chapter describes several cues to detect deception in natural language that may also be used within automatic systems. Finally, existing responses to mis- and disinformation are listed, such as fact-checking, educating, inoculating (i.e., preventing by making aware of the phenomenon, as for vaccines in medicine), or labeling the information.

CHAPTER 3. CREDIBILITY ASSESSMENT MODELS AND TRUST INDICATORS IN SOCIAL SCIENCES

This chapter proposes a wide survey and a deep analysis of the literature about credibility and trust. The credibility of a message depends on the characteristics of its source, contents, and medium of delivery. The author introduces many works that aim to describe these characteristics. Initiatives are presented, for instance, the common credibility assessment terminology definitions [3].

People tend to naturally trust others and the messages they receive. The author describes trust and distrust markers in language. These markers can be used in fake detection systems. Following that, she proposes ideas to develop AI systems may assist people in discriminating online information.

CHAPTER 4. PHILOSOPHIES OF TRUTH

This chapter discusses various philosophical perspective about truth. The author summarizes the different philosophies of truth and untangles the key concepts of truth, reality, facts, and knowledge that are often confused one for another. This perspective leads us to think about what exactly the automated AI-based fake detection systems should look for. Facts may be wanted more than truth for users of these systems.

> ***"The credibility of a message depends on the characteristics of its source, contents, and medium of delivery".***

PART 2: APPLIED PROFESSIONAL PRACTICES AND ARTIFICIAL INTELLIGENCE

CHAPTER 5. INVESTIGATION IN LAW ENFORCEMENT, JOURNALISM, AND SCIENCES

In this chapter, the author investigates the best practices of three expert domains: law enforcement, scientific inquiry, and investigative journalism. Her aim is to find insights that would be useful for informing automated systems on the detection of deception and supporting the process of fact-checking. From law enforcement experts, she suggests that established tools and checklists for statement validity analysis used by detectives during police interrogations are potential guides to develop automated lie detection systems. The five Ws (who, what, when, where, and why) of journalism form a useful framework to examine the credibility of some online stories. She also emphasizes the use of rational observation and systematic questions, linked to the scientific method.

As a conclusion to this chapter, the author states that what unites experts in these three domains is their inquisitive critical mindset. She then makes the point that even if we develop automated AI-based fake detection systems, humans should still validate their results: "Technology advises and assists us but never replaces human judgement in determining what is truthful and what is disinformative".

CHAPTER 6. MANIPULATION IN MARKETING, ADVERTISING, PROPAGANDA, AND PUBLIC RELATIONS

The chapter describes the different means used to propagate mis- and disinformation. It also describes AI techniques that mimic these means of propagation. Examples are taken from the marketing, advertising, and public relation domains, with specific and concrete examples from each. The author further explores what makes us—as humans—particularly vulnerable to viral conspiracy theories and how human biases are being exploited to propagate mis- and disinformation.

CHAPTER 7. ARTIFICIALLY INTELLIGENT SOLUTIONS: DETECTION, DEBUNKING, AND FACT-CHECKING

This is the longest chapter in the book (58 pages). In this chapter, the author presents an overview of five large families of AI-enabled applications that may support humans in detecting and managing fakes. AI-based systems aim to accomplish three different tasks: assist in the detection itself, alert, and filter fake information. The author emphasizes the importance of having a good human–machine interface. Humans should remain in the loop and use AI only to assist them.

Five application families are reviewed: deception detectors, click-bait detectors, satire detectors, rumor debunkers, and fact-checkers. For each of these families of applications, the author offers detailed examples, together with some technical details of how they work. She keeps technical vocabulary to a minimum, in "favor [of] explaining step-wise procedures in principles, and wherever possible, offer some examples".

CHAPTER 8. CONCLUSIONS: LESSONS FOR INFODEMIC CONTROL AND FUTURE OF DIGITAL VERIFICATION

This chapter concludes the book, giving arguments and claims about the use of automated ways of detecting online fakes. The author then gives recommendations for educational, AI-based, and regulatory interventions.

SUMMARY

The book gives an overview on many historical, technological, and psychological aspects of the subject. The author explores many AI approaches, all based on statistical machine learning solutions for detecting fakes.

The author emphasizes that deceptions and fakes form a diverse set. The existing—and future—automated solutions for fake detection should thus focus on a limited objective, in order to be both relevant and efficient. She presents an interesting alignment of taxonomies of deception varieties. This should be widely used by researchers and solution providers as a pivotal model that would enable the rigorous description of the specific fakes they aim at detecting.

The book presents existing work mostly using statistical NLP. It is thus oriented toward using pragmatics and content to detect fakes. It would be interesting to expand the analysis toward meeting semantic and knowledge-based solutions, using a semantic description of human motivations and processes, which are well detailed by the author. This would enable understanding and explaining AI-based fake detection.

Reading the book is inspiring. For instance, a direct idea that comes to mind after reading Chapter 3 is to extend the work to provide a synthesis of all the reviewed models into a single complete one. Existing ontologies of trust could be used to align the different models, for instance. This work of synthesis is nicely done for the alignment regarding deception taxonomies proposed in Chapter 2.

The reading also raises the possibility of attempting to build semantic models to inform AI (i.e., choose which model from those presented in the book to use to inform machine learning approaches on a specific task). With the broad literature review and the analysis of existing works on the perception and detection of fakes by humans that the author provides, this semantic modelling step may be envisioned.

Finally, I recommend reading this inspiring book to members of the International Society of Information Fusion community who would want to have a different perspective on the subject of detecting fakes. The ideas and explanations given by Rubin may help our community find innovative ways of managing mis- and disinformation, using our usual solutions with enriched background and understanding of the way mis- and disinformation are produced and spread.

REFERENCES

1. Victoria Rubin, prof., research lab director, writer, speaker, http://victoriarubin.fims.uwo.ca/research/.
2. Shane, T. The psychology of misinformation: why we're vulnerable, June 2020. [Online] https://firstdraftnews.org/articles/the-psychology-of-misinformation-why-were-vulnerable/.
3. Zhang, A. X., Ranganathan, A., Metz, S. E., Appling, S., Moon Sehat, C., Gilmore, N., et al. A structured response to misinformation: defining and annotating credibility indicators in news articles. In *Companion Proceedings of the Web Conference 2018*. Geneva, Switzerland: International World Wide Web Conferences Steering Committee, 2018, 603–612, doi:10.1145/3184558.3188731

Claire Laudy is a senior research engineer, at the research center of Thales, where she has worked for 22 years. She focuses on knowledge representation and modeling, (semantic networks, ontologies, conceptual graphs), semantic fusion and graph algorithms for high-level information management. She was involved in numerous research projects dealing with the new usage soft information (e.g., social media posts and citizen science) within fusion-based systems. Among others, she was involved in the promotion of the use of soft information towards operational experts used to work with sensor data only, such as marine biologists and Police and Safety organizations. Claire first worked as an engineer in the domain of Human-machine interaction and then obtained her Ph.D. from Sorbone Université in 2010 in the field of semantic information fusion.

Since 2007, Claire has been involved in the international research community on information fusion, through her presence at FUSION conferences. She published several book chapters on high level and soft information fusion. She has been a member of the ISIF Working Group on Evaluation Techniques of Uncertainty Representation (ETUR) since 2017 and was recently elected member of the Board of Directors of the International Society for Information Fusion (ISIF).

ISIF VISION STATEMENT

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.

ISIF Journal of Advances in Information Fusion (JAIF)

The Journal of Advances in Information Fusion (JAIF) is the flagship journal of ISIF. JAIF is an open-access, peer-reviewed, semi-annual, archival journal published electronically and distributed via the internet. JAIF was founded in July 2006. The journal is indexed at SCOPUS, free for authors, and freely available for readers at http://www.isif.org/journals/all. Authors are invited to submit both regular papers as well as short correspondences describing advances, applications, and new ideas in information fusion, both theory and application. Authors of papers presented at our annual International Conference on Information Fusion are strongly encouraged to consider submitting expanded versions of their papers to JAIF. Manuscripts can be submitted at http://jaif.msubmit.net.

Rio de Janeiro
Brasil

Rio de Janeiro, Brasil, will host the 28th ISIF Fusion Conference from July 7-10, 2025, the first ever in South America. Rio de Janeiro never fails to impress with its modern outlook that reflects its progression through the times of yore. The historic sites, sparkling beaches, green belts and cheerful attitude of the locals attract tourists where heaven meets earth. If Ronaldo played football at the Maracanã stadium, then Christ the Redeemer put Rio on the world map. Rio's carnival, with its energetic samba dancers , attracts thousands of tourists from around the world. The rain-forests, museums, beaches, gardens and the glitz of the city have made Rio what it is today Brasil's top-notch tourist destination.

The venue is at the center of the touristic zone of Rio de Janeiro in front of one of the most famous beaches - Copacabana. It provides excellent facilities, first-rate service levels, and all the additional amenities expected in major hotels, including easy access to various attractions - such as the Sugar Loaf, Christ the Redeemer and much more.

In addition to the numerous attractions and activities located in the South Zone, Rio has much more to offer. From the pitoresc neighbourhood of Santa Teresa to Maracanã Stadium, from Gigoia Island to beautiful art museums, from Samba Schools to Lapa passing through the Royal Gabinet of Literature (one of the most beautiful libraries in the world) and the beaches of Barra da Tijuca - Rio has so much to surprise you.

FUSION 2025
July 7-10, 2025

Rio is waiting for you!