



AFRL-AFOSR-VA-TR-2022-0067

Reliable Inference in Dynamic Data Fusion

Peter Willett
UNIVERSITY OF CONNECTICUT
438 WHITNEY RD EXTENSION UNIT 1133
STORRS, CT, 06269
USA

12/15/2021
Final Technical Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
Air Force Office of Scientific Research
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 15-12-2021		2. REPORT TYPE Final		3. DATES COVERED (From - To) 01 Aug 2018 - 31 Dec 2021	
4. TITLE AND SUBTITLE Reliable Inference in Dynamic Data Fusion				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER FA9550-18-1-0463	
				5c. PROGRAM ELEMENT NUMBER 61102F	
6. AUTHOR(S) Peter Willett				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) UNIVERSITY OF CONNECTICUT 438 WHITNEY RD EXTENSION UNIT 1133 STORRS, CT 06269 USA				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AF Office of Scientific Research 875 N. Randolph St. Room 3112 Arlington, VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR RTA2	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-AFOSR-VA-TR-2022-0067	
12. DISTRIBUTION/AVAILABILITY STATEMENT A Distribution Unlimited: PB Public Release					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Proper inference in a decision or inference network requires that the commander (technically: the fusion center) have an understanding of the relative weight that he / she should place on the inputs each subordinate. Recent works have addressed the problem of estimating agents' behaviors in complex networks, of which social networks are a prominent example. These works are especially promising and would seem to be of considerable practical importance in a wide variety of command & control venues. However, these works are perhaps limited by their somewhat idealized assumptions: that the commander (fusion center) possess full information of all subordinates' histories, and that conditional statistical independence between these histories can be assumed. In the proposed project we intend to explore more general situations: of dependent sensors, of unknown structure of that (possible) dependence, of missing data and of subordinate identities that are either obscured / adulterated / entirely missing. For such dynamic fused inference problems we propose to extend results in a number of directions: exploring dependency amongst data sources (physical proximity or "group-think"), in term of useful communication strategies when the inference task and quantization are not necessarily matched, and even the unlabeled case in which the identity of each measurement's source is unknown – this is a form of the data association problem					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON ERIK BLASCH
a. REPORT	b. ABSTRACT	c. THIS PAGE			
U	U	U	UU	21	19b. TELEPHONE NUMBER (Include area code) 426-7311

Standard Form 298 (Rev.8/98)
Prescribed by ANSI Std. Z39.18

Reliable Inference in Dynamic Data Fusion

Peter Willett
ECE Department
U-4157, University of Connecticut
Storrs, CT 06269
phone: (860) 486-2195
fax: (860) 486-5585
email: peter.willett@uconn.edu

Award Number: FA9550-18-1-0463

Sponsor: AFOSR
Program Manager: Dr. Erik Blasch

LONG-TERM GOALS

Proper inference in a decision or inference network requires that the commander (technically: the fusion center) have an understanding of the relative weight that he / she should place on the inputs each subordinate. Recent works have addressed the problem of estimating agents' behaviors in complex networks, of which social networks are a prominent example. These works are especially promising and would seem to be of considerable practical importance in a wide variety of command & control venues. However, these works are perhaps limited by their somewhat idealized assumptions: that the commander (fusion center) possess full information of all subordinates' histories, and that conditional statistical independence between these histories can be assumed. In the proposed project we intend to explore more general situations: of dependent sensors, of unknown structure of that (possible) dependence, of missing data and of subordinate identities that are either obscured / adulterated / entirely missing. For such dynamic fused inference problems we propose to extend results in a number of directions: exploring dependency amongst data sources (physical proximity or "group-think"), in term of useful communication strategies when the inference task and quantization are not necessarily matched, and even the unlabeled case in which the identity of each measurement's source is unknown – this is a form of the data association problem

We also recognize that inference of dynamic underlying situations is of key interest. Given a traditional framework involving measurements and physical "targets" this is a familiar problem of tracking. But can techniques from target tracking and multi-sensor data association be applied to extract states that are not physical (physical as would be, say. an aircraft observed by radar) but are instead at a higher level? An example might be a terrorist threat or a battle plan – these would be observed from multiple sources via measurements such as intelligence reports and telemetry, and may even be thought to encompass civilian sources such as news or financial transactions. These are not

standard data, and neither are the dynamic systems that are of interest here the usual kinematic ones. Nonetheless we note that there is much commonality with (and thus opportunity for application of mature and emerging tools from) traditional target tracking: there can be multiple “targets”, there is clutter, and there is behavior that we might reasonably model via statistics. For such fused inference of dynamic systems we have the goal of extraction of unusual dynamic patterns that are evolving and would seem to merit closer attention. We specifically propose to ingest feature (identity) information, via modeling the clutter as a rich collection of similar activities and by adapting modern multi-sensor data association techniques for the task.

OBJECTIVES

The focus of the research is reliable inference in dynamic systems with fused observations. We begin with iterative fused decision-making (decisions need not be binary) from sensors / subordinates whose operating points (quality of data) are initially unknown. With time it should become possible to infer, jointly, that states of nature both of sensor suite and of the underlying situation (i.e., the real inference goal). We extend this in a number of directions as well: exploring dependency amongst data sources (physical proximity or “group-think”) and even the unlabeled case in which the identity of the measurement’s source is unknown. In the proposed project we additionally intend to explore situations: of dependent sensors, of unknown structure of that (possible) dependence of missing data and of subordinate identities that are either obscured / adulterated / entirely missing. We further intend to extend these ideas to fused estimation of dynamic threats, by ingesting feature (identity) information, via modeling the clutter as a rich collection of similar activities and by adapting modern multi-sensor data association techniques for the task.

APPROACH

In the first year of this effort we concentrated on the following two objectives.

1. Unknown Identities of Decision-Makers. It is likely that in an operational situation the fusion center (commander) receives sensor reports from his/her subordinates that is disordered: their identities can be mixed up or even completely missing. Such a situation can be a concern in a “big data” application in which data pedigree can be lost or for reasons of storage gets discarded. The former situation suggests an interesting twist on Task #1: the identity information has a strong prior to be correct but the locations of errors of identity must be inferred; again, the EM algorithm is suggested. However, it may be so that all identity information is lost, and the commander is presented with what might be thought of as a “bag of decisions” each time. The method of types is in this case proposed to accomplish joint inference of local (unlabeled) belief levels and ongoing optimal decisions.
2. Operating Points for Fused Inference of Dynamic Systems. Under previous support we have explored dynamic event extraction: we have developed a reasonable hidden Markov model, we have learnt to ingest (identity) features, we have a multi- Bernoulli filter-inspired extraction approach – and we have even provided some theoretical analysis. As part of the proposed work we will extend this in two ways. First, we intend to cast the measurements as a fused stream of data from sources of unknown credibility that must be estimated. Second, each such information source must be assumed to be cluttered with “ambient” events (such as the

financial and travel footprint of a family going on vacation) that, while benign and likely uncomplicated, are dynamic and in some sense similar to the threats sought. These must be modeled (from data) and suppressed (by a multi-target tracker).

Work in the second year focused on the following objectives.

3. Identity Uncertainty in Data Fusion. When data are to be fused from multiple sources, and when this data refers to multiple truth objects, a key concern is to determine which data from one sensor go with which data from another: the “data association” problem. Actually the means for such fusion – and even good approaches for the association process – are fairly well-known. What is lacking is an understanding of the quality of the associations made. We attempt to provide this, and we intend to explore the effect of sensor bias and positioning.
4. Sensor Networks with Extreme Communication Constraints. Consider inference by a network of sensors whose positions are unknown and whose locations are subject to drift and diffusion – a Poisson field. Further, assume that in such a network the sensors, while cognizant of their identities and other such relevant data choose not to transmit that to the fusion center, in order to preserve bandwidth. What can be done? And what is lost? We examine these questions, as well as evaluating the role (in the information theoretic sense) of identity versus observation. That is, suppose two bandwidth-equal networks are compared; one with n sensors that transmit only observation; and the other with $n/2$ sensors that transmit both data and identity. Which is preferable, and when?
5. Tracking of the COVID-19 Epidemic Status. Admittedly epidemiology is not in the direct line of the proposed research, but given the skills represented and the pressing need for them during the current health emergency, it seems reasonable to be opportunistic. With a joint team of US and Italian researchers we have shown that we can reliably estimate and forecast the evolution of the infections from daily – and possibly uncertain – publicly available information provided by authorities, e.g., daily numbers of infected and recovered individuals. The proposed method is able to estimate infection and recovery parameters, and to track and predict the epidemiological curve with good accuracy when applied to real data from Lombardy region in Italy, and from the USA. We are presently extending our approaches to data segmentation, change detection (as in an increase/decrease in the infected numbers) and regional clustering.

Work in the third year focused on the following objectives.

SELECTED WORK COMPLETED

In the following subsections we briefly describe some of the work we have performed under this contract. More papers (and more details therein) are given in the following section “Publications”.




Figure 1: Error probabilities for $n=100$ samples and observations of cardinality $m=3$ and $m=10$. Note the difference between labeled and unlabeled cases, and that the GLR can, surprisingly, be relatively poor.

1. “Algorithms and Fundamental Limits for Unlabeled Detection using Types”. This paper was authored by S. Marano and P. Willett, and published in the *IEEE Transactions on Signal Processing*, vol. 67, no. 8, pp. 2022-2035, April 2019. We consider a canonical binary hypothesis test with independent data under both hypotheses. Motivated by modern applications of sensor networks engaged in big data analysis, we assume that the observation vector $X=[x_1, x_2, \dots, x_n]$ collected by the peripheral units is delivered to the fusion center in the form of a random set $X_u=\{x_1, x_2, \dots, x_n\}$ rather than a random vector – the distinction is that the former is labeled (it is known the source of x_i) whereas in the latter the provenance of x_i is not known, only that one of the sensors has communicated this x_i . The theoretical question addressed is how much information for detection is carried by X_u , as opposed to X . We provide the asymptotic (n diverging) characterization of the performance of the optimal test in terms of an error exponent rate $\Omega_u(\alpha)$, which replaces the canonical rate $\Omega(\alpha)$ of the labeled case – α is determined by the false alarm rate. It is proven that, when type I error tends to zero as $e^{-n\alpha}$ with the data size n , type II error may converge to zero as $\exp(-n\Omega_u(\alpha))$ but not faster. The rate difference $\Omega(\alpha) - \Omega_u(\alpha)$ quantifies the loss of information induced by the loss of data labels. The second part of this paper addresses the practical question of how to solve the test by algorithms of affordable computational complexity and good performance. The ULR detector makes no attempts to estimate the labels and is very efficient computationally. The GLRT solution for unlabeled data boils down to an assignment problem, for which a tailored form of the auction algorithm can be exploited. We also propose two alternative detection algorithms

with good trade-off between performance and complexity, as we show by computer experiments. Interestingly this is one of the few cases in which the GLRT (auction!) is very suboptimal, the type is actually a better indicator for detection.




Figure 2: Performance for situation that under \mathcal{H}_1 half the decisions are “always right” and half are “always wrong”. Note that the GLRT is quiet poor. The notional situation is described in the text, below.

2. **“Making Decisions by Unlabeled Bits”**. This paper was authored by S. Marano and P. Willett, and is to appear in the *IEEE Transactions on Signal Processing*. The error exponent for unlabeled detection has been completely characterized in the above paper, but left in the implicit form of a convex optimization problem from which limited insight and intuition can be gained. Here we focus on the case in which observations are binary and show that the structure and the properties of error exponent becomes self-evident, also allowing straightforward numerical solution that does not require specialized convex optimization tools. In the challenging scenario of low-detectability regime, we provide simple closed-form analytical solutions for the error exponent and related quantities, for which we obtain much insight and intuition. From a theoretical point of view, these are the main contributions of the present study. From a practical perspective in several decision statistics have been proposed and here we show that in the case of binary observations their properties and relative merits becomes very clear, but also unpleasant. The decision algorithm based on the GLRT principle should be used with care because its performance may be quite poor and possibly biased. Figure 2 deserves some explanation. Suppose that \mathcal{H}_0 refers to the situation that all n sensors deliver their data as “coin flips”. An further suppose that \mathcal{H}_1 refers to the case that the first $n/2$ sensors are all-zero while the remainder are all-one. Testing in the labeled case is easy. Unwary testing in the unlabeled case can be disastrous, since the count of the number of ones has mean $n/2$ in either case. The GLRT (which attempts under \mathcal{H}_1 to match the 1’s to a sensor indexed above $n/2$) is can be very poor. The best unlabeled strategy is actually to test the deviation of the number of 1’s from $n/2$.




Figure 3: An example HMM used to model a terrorist event. This is much less complex than would be of interest but illustrates the sequential and multi-optional behavior seen. The “states” of the attack (for example, “Gather chemical fertilizer”) are represented as X’s; the transactional observations are the Z’s. Naturally, there is a great deal of ambient clutter (such as people buying gardening supplies) amongst the Z’s. The key is to use modern target tracking ideas to extract the low-observable “targets”.




Figure 4: Quickest detection results over 2000 MC runs for the two filters. Vertical axis is in log scale since false alarms are rare. Mean target duration is marked since filter is useless if delay to detection is past this point. improves detection time and suppresses false alarms. The performance for threshold value of 0.8 is noted on both plots.




Figure 5: A graph of probability of the average time to detection vs. the logarithm of time between false alarms for 200 activations (Monte Carlo runs). The vertical orange line indicates the expected duration for the respective process. The plot demonstrates the advantage to using the underlying “clique” (group-membership and social connection) prior knowledge.

3. **“Target Tracking Applied to Extraction of Multiple Evolving Threats from a Stream of Surveillance Data”**. This paper was co-authored by Z. Sutton, P. Willett and Y. Bar-Shalom, and has been submitted to the IEEE Transactions on Computational Social Systems. Threats are composed of some process or plan being carried out by a group of people with an end goal that is generally to cause harm. Some examples of these kinds of threats are terrorist attacks, military actions, or stock fraud. These threats can be modeled stochastically with help from experts within the relevant field. We model these threats with a hypothesis as to how these events will unfold along with a method for observing the unfolding threat. We use this model to detect the threat before its completion and theoretically allow for preemptive action against the threat’s perpetrators. The models used for threats in this paper are variations of Hidden Markov Models (HMMs) with sparse observation emission (compared to the expected process length), see Figure 3. There is a rich target tracking literature, with many methods to deal with dynamic estimation, target extraction, data association and multiple objects. Here we co-opt this literature, and offer significant refinement of earlier estimation procedures. Specifically, we now allow for *multiple* threats to exist and be extracted. This has necessitated a *data association* step, since we now “frame” the observation stream into sets of transactional data (as opposed to one-at-a-time ingestion) to facilitate efficient operation. We further have augmented our model to admit “identity” information: transactional data often involves actors and places, etc., and the continual reappearance of these can offer a significant clue in the data association phase. The improvement through the use of identity information is readily seen in Figure 4.
4. **“Taking Advantage of Group Behavior When Tracking Multiple Threats in Cluttered Surveillance Data”**. This paper was authored by A. Finelli, Z. Sutton, P. Willett, and Y. Bar-

Shalom, and was presented at the IEEE Aerospace Conference, Big Sky MT, in March 2019. The paper was recognized as **Best Paper** in Track 6. The population observed is assumed to be organized into groups called “cliques”. Rather than tracking an individual’s involvement probability, we track a clique’s (group’s) involvement probability across all threats using a Bayesian update equation and conditioning on association events between the observations and the set of measurement generating HMMs (threat and clutter processes). We assign an individual’s probability conditionally based on their group’s and the state of each threat process then its state is estimated using a bank of Bernoulli filters. This allows us to accurately detect multiple threat processes within a single stream of observations (most of which will be clutter). Figure 5 shows that there is considerable advantage to the exploitation of clique behavior.



Figure 6: The targets present in the experiments. Top left: The Ocean Space Drone. Bottom left: Munkholmen II. Right: The seamount.

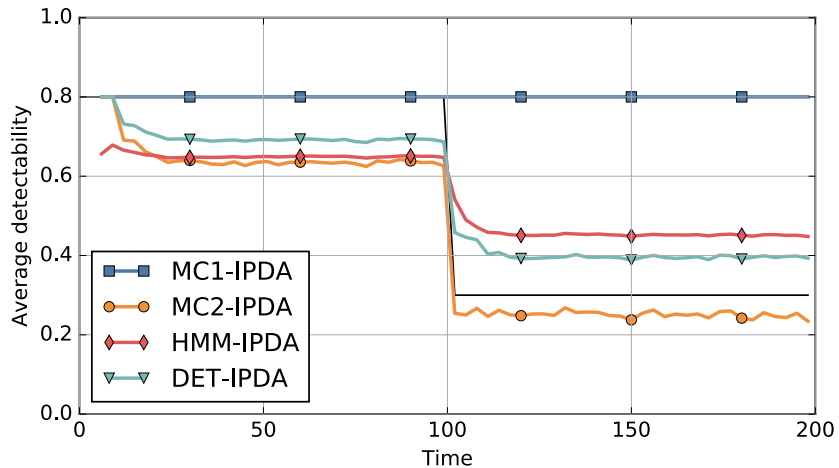


Figure 7: Average mode of the detectability estimate of confirmed tracks when the target is present. The black line shows the ground truth.

5. “**Estimation of Target Detectability for Maritime Target Tracking in the PDA Framework**”. This paper was authored by E. Wilthil, P. Willett, Y. Bar-Shalom, and E. Brekke was presented at the ISIF FUSION Conference, Ottawa Canada, in June 2019. The paper was recognized as **Best Student Paper** runner-up. Accounting for varying target detectability can

significantly improve tracking performance when these issues are present. The detectability can be estimated with a HMM based on the number of validated measurements, or the probability of the joint detectability and target existence may be jointly evaluated using the based on the likelihood ratio of a target vs. clutter. Simulations shows that both of these methods are able to maintain the track when the detectability is lowered, and terminates lost tracks significantly faster than a Markov chain 2-IPDA. Tests on real data shows that the joint estimation of target detectability and existence probabilities reduces the number of false tracks, at the cost of slightly higher track confirmation time. The paper uses real data from radar and some friendly targets, and Figure 6 shows these. Figure 7 illustrates the point of the work: when the continuous-valued target-detection probability is integrated to the tracker and is itself tracked, the target existence probability can be far more effectively estimated.

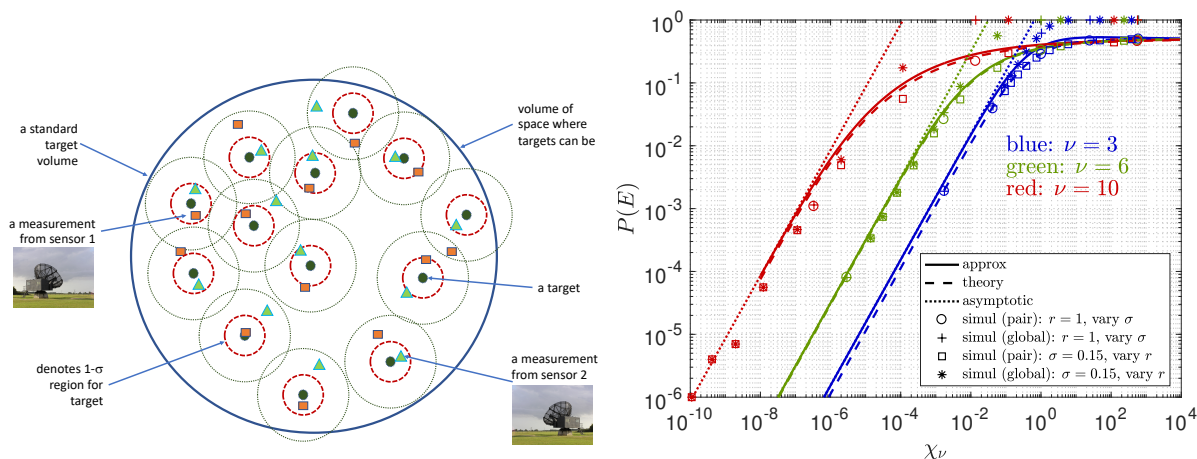


Figure 8. Left: Notional sketch of the problem and our posing in two dimensions. There are $N = 13$ targets randomly located in a ball whose area (volume) is $N = 13$ times that of a standard target volume. A standard volume is a ball whose radius is one half the nominal target separation; but of course some pairs or targets are close together and some are more widely separated. Each target is represented by a measurement at radar sensor 1 (a small red square) and also at radar sensor 2 (a small green triangle). These are assumed generated by adding independent Gaussian noise with the indicated standard deviation. Right: Solid lines are the approximated analytical probability of single pairwise switch error $P(E)$ shown, as a function of the scene-difficulty parameter, for three values of dimension = 3, 6, 10. Dashed lines show the value of accurate $P(E)$ via numerical integration. Dotted lines are the asymptotic approximation, valid for small $P(E)$. Symbols show the results of 10^6 Monte Carlo simulations with $N = 25$.

6. **“On the Probability of Cross-Radar Assignment Error”.** This paper, by P. Braca, P. Willett and W.D. Blair, will be presented at the 2020 IEEE Radar Conference. If two radar sensors observe the same target their measurements can be combined to produce a *fused* target-state estimate that is of higher quality than that from one radar alone. If there are multiple targets whose information is shared, a necessary first step to fusion is to “assign” each measurement from the first sensor to that at the other in such a way that both refer to the same underlying object, a task generally accomplished by minimizing a global cost involving distance. An assignment error occurs when the measurement originated by target i at the first radar is wrongly associated to a measurement originated by target j (not i) at the second radar, see the left panel of Figure 8. Naturally, when such an error occurs the result is fusion of information describing disparate objects, resulting in degraded estimation performance and poor self-assessment in terms of posterior uncertainty. Here we address the issue, and derive approximate

assignment error probability. Remarkably, performance (see the right panel of Figure 8) depends only upon the parameters combined to a single scalar constant.

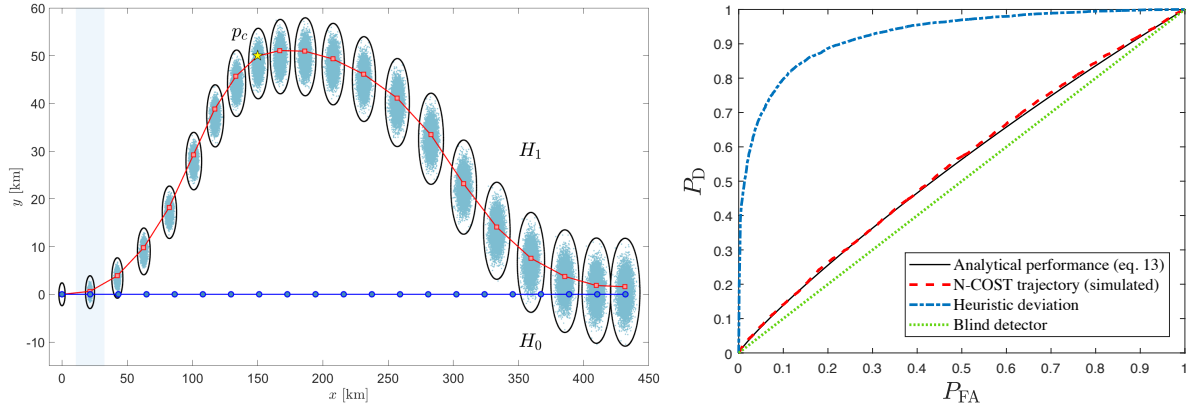


Figure 9. Comparison of the optimized positions (red squares) versus the positions drawn from the OU process evolution (cerulean points) and the positions of the nominal trajectory (blue dots). The ellipses (black solid lines) represent the 95%-confidence covariance of the OU/IOU process given the initial point of the trajectory. The point p_c near which the vessel lingers is marked with a yellow star. On the left, a single radar contact, indicated by light blue shade, located in $p_{m1} = 1$ is taken into account. The right plot shows ROC curves describing the performance of an anomaly-detector in terms of track detection probability P_D versus false alarm probability P_{FA} when the optimal deviation (when it exists) is an OU process (black solid line) and is the deterministic output of the N-COST algorithm (red dashed line). The first ROC curve is the predicted performance of the optimal anomaly detector operating on the pessimal trajectory, performance provided by equation (13) in the paper, while the second is simulated with 1000 Monte Carlo runs, a close match. The blue dash-dotted ROC curve describes a sub-optimal (unwary) deviation, and the other plot is the chance line.

7. **“Optimal Opponent Stealth Trajectory Planning based on an Efficient Optimization Technique.”** This paper, by A. Aubry, P. Braca, E. d’Afflisio, A. De Maio, L. Millefiori and Peter Willett, has been submitted to IEEE Transactions on Signal Processing. This work proposed a computationally efficient technique, called Non-Convex Optimized Stealth Trajectory (N-COST) algorithm, to solve the route planning problem with the goal to make a vessel’s trajectory as stealthy as possible to an anomaly detector, so as to hide a deviation from a nominal traffic route to accomplish a specific mission. Previous research has discussed tracking of an object whose velocity evolves according to an Ornstein-Uhlenbeck mean-reverting stochastic process, while proper kinematic and practical constraints are taken into account – IOU processes (integrated because velocity is the derivative of position) have been shown to be excellent empirical matches to commercial traffic in many modalities (air, sea, etc.) – and the model is key to tracking when the observation stream has gaps such as between fusion hand-offs. In this paper we look at the reverse problem: How can a target plan its most effective trajectory such that a goal is met (say: a rendezvous for smuggling) yet detection of the accompanying “diversion” is made most problematic; presumably, a game-theoretic approach would incorporate both perspectives. From this “red-team” viewpoint, the optimization problem minimizes the Kullback-Leibler divergence between the statistical hypotheses of the nominal and the anomalous trajectories. Interesting case studies concerning both synthetic and real-world scenarios are reported to prove the effectiveness of the proposed N-COST algorithm. In other words, as illustrated in **Figure 9**, we considered the worst condition case from the detection point of view, by minimizing its performance with the final

goal of highlighting the detector limitations and opening the door to possible future works aiming at improving the anomaly detector capabilities by determining the optimal surveillance asset.

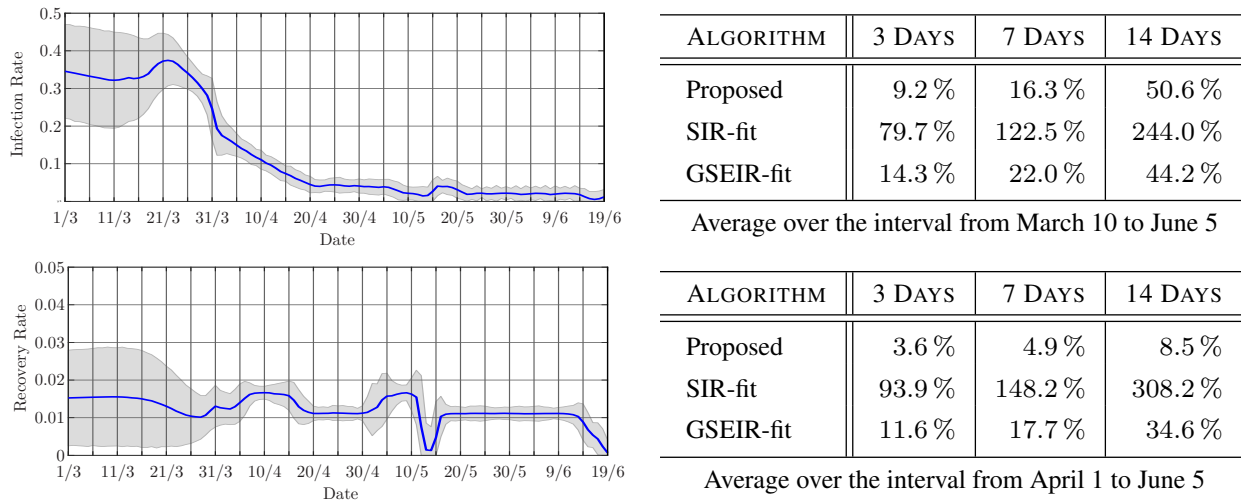


Figure 10. Left: Estimated infection rate (“beta”) and recovery rate (“gamma”) for USA, in days. The shaded areas represent the 90 % confidence interval. Right: Average mean absolute percentage errors (MAPEs) of the forecasts of the epidemic evolution in USA, via the proposed algorithm, and with the fixed-parameter SIR-fit and GSEIR-fit (GSEIR is a more-sophisticated version of SIR) curve-fitting approaches, for different forecast horizons: 3, 7, and 14 days. The upper table reports the average MAPEs computed over the interval from March 10 to June 5; the lower table reports the average MAPEs computed over the interval from April 1 to June 5.

8. “Adaptive Bayesian Learning and Forecasting of Epidemic Evolution – Data Analysis of the COVID-19 Outbreak.” This paper, by D. Gaglione, P. Braca, L. Millefiori, G. Soldi, N. Forti, S. Marano, P. Willett and K. Pattipati, will appear online shortly in IEEE Access. The recent worldwide epidemic outbreak, due to a new strain of Coronavirus, has intensified research into novel mathematical models and algorithms that are able to reliably estimate and predict the epidemiological curve of the infection. The signal processing community has in its arsenal many tools to track, to track models that “switch”, and to detect such changes; and it is good, in the context of this project, to apply these tools to a current crisis. Hence, in this paper, we proposed a Bayesian sequential estimation and forecasting algorithm that, based on the information that authorities provide on a daily basis, that is able to estimate the state of the epidemic and the parameters of the underlying model, as well as to forecast the evolution of the epidemiological curve. We developed an efficient implementation specifically tailored to the stochastic SIR (susceptible/infected/recovered) model of pandemic evolution. The proposed algorithm is validated using synthetic data simulating two epidemic scenarios, and on real data acquired during the recent COVID-19 outbreak both in the Lombardy region of Italy and in the USA. The model “switches” mentioned generally reflect changes in policy, specifically a “lockdown”. Results (see Figure 10) show that the mean absolute percentage error computed after the lockdown is on average below 5% when the forecast is at 7 days, and approximately 10% when the forecast horizon is 14 days. Moreover, the described Bayesian framework outperforms curve-fitting approaches that use deterministic epidemiological models, particularly when a clear change of model parameters occur, e.g., a decrease of the infection rate following the lockdown. Finally, accurate and timely data collection, especially on recovered individuals, hospitalizations, intensive care unit admissions, and intubations, is

essential for reliable model based decisions. There exists an enormous amount of very recent literature related to the forecast of COVID-19 pandemic evolution, and the analysis of this literature makes clear the effectiveness of model-based approaches, over less structured data-centric methodologies. In this respect, one lesson learned by the present study is that accurate epidemic modeling requires accurate estimation of time-varying key parameters, such as the infection rate “beta” and recovery rate “gamma” (the celebrated R0 parameter – that which is exponentiated by time to determine the pandemic’s evolution – is beta divided by gamma). This is obviously true in the presence of abrupt changes of the underlying physical situation (e.g. adoption of drastic countermeasures) but, more interestingly, it is by no means limited to these extreme situations. One consequence is that, once the epidemic is under control, small variations in the estimated beta may be used as a sensible proxy for incipient detection of possible pandemic recurrence.




Figure 11: Left is an example trajectory and ML-PMHT’s tracks. On the right we show the benefits of data fusion, comparing the joint ML-PMHT to one that works individually on the 2 FPAs.

9. **“ML-PMH Tracking in 3 Dimensions Using Cluttered Measurements From Multiple 2-Dimensional Sensors.”** This paper, by Z. Sutton, P. Willett, T. Fair and Y. Bar-Shalom extends the maximum likelihood probabilistic multi-hypothesis tracker (ML-PMHT) to track targets in a 3-dimensional “global” space with observations provided by multiple 2-dimensional sensors placed throughout the global space. ML-PMHT is a tracking method whose flexibility and scalability derive from relinquishing the assumption that each target emits at most one “hit” per scan of the sensor. It is a maximum likelihood method that essentially reduces to an optimization problem—recursively maximizing a likelihood function that is simple to evaluate given a batch of observations. Unlike maximum a posteriori or MMSE trackers, this likelihood maximization tracker requires neither prior knowledge about target motion nor measurement association, making it conceptually easy to work with. Since the observation model is nonlinear, the likelihood maximization is done via hill climbing. For this purpose, we also address the issue of “hill finding”. Due to the presence of clutter in the measurement model, the likelihood is a multi-modal function of the parameter space. That is, there are multiple hills in the likelihood function and it is of great advantage to the tracker to initialize the hill climber close to the right hill—the one whose peak is the global maximum. In this work, we present a data-driven method of initializing the hill climber based on the received observations. Figure 11 shows the improvement possible using our ML-PMHT approach; the paper also shows similar comparisons against other tracking/fusion approaches.

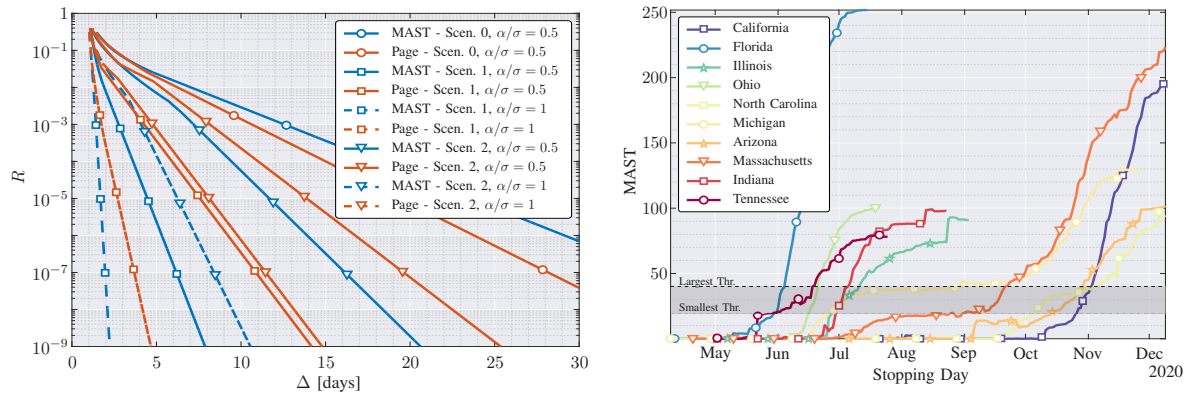


Figure 12: Left shows the operational characteristic (risk R versus decision delay Δ) of the MAST quickest detection test, compared to the benchmark Page's test. Three scenarios are considered, as described in the main text. In scenario 0, Page's test is optimal. MAST outperforms Page's test in scenarios 1 and 2, in which the sequences $\{\mu_{0,n}\}$ and $\{\mu_{1,n}\}$ are time-varying. Scenario 2, in particular, mimics the actual behavior of the sequences, as observed in COVID-19 pandemic data. MAST decision statistic computed for 10 US states and used to detect the onset of the COVID-19 second wave. The dashed horizontal lines represent the smallest and largest thresholds corresponding to $R = 10^{-9}$, for the ensemble of the ten states. Curves are prolonged beyond threshold crossing for clarity.

10. **“Quickest Detection of COVID-19 Pandemic Onset.”** This Signal Processing Letter, by P. Braca, D. Gaglione, S. Marano, L. Millefiori, P. Willett and K. Pattipati, develops a novel version of Page's CUSUM quickest-detection test, designed to work in composite hypothesis scenarios with time-varying data statistics, specifically an unknown change in mean. The derived decision statistic can be cast in recursive form, particularly suited for on-line analysis. When applied to COVID-19 data, the developed test allows to predict the explosion of the infection on a large scale, by analyzing the publicly-available sequence of new positive individuals per day from different countries. It is envisioned that the developed tool might help to proactively supporting the political decision makers for the adoption of restrictive measures to contain the COVID-19 pandemic explosion. Figure 12 shows the performance of MAST as compared to the optimal (but somewhat fragile) Page test, and applies it to detect critical phases of the pandemic in several US states.




Figure 13: Left shows the growth rate of the hospitalized individuals (green solid line) — and its time-varying mean obtained through a moving average that uses a window of 21 days (magenta solid line) — compared to the growth rate of the daily new positives individuals (green area) — and its time-varying mean (magenta dotted line) — in Italy since February 21, 2020. On the right, we show the application of the MAST procedure on the growth rate sequence of the daily new positive individuals and on the growth rate sequence of the hospitalized individuals (dotted lines). On the left-side vertical axis we select a desired risk, e.g., $R = .0001$. Then, the blue curves indicate the stopping day (about July 18 if the growth rate sequence of the daily new positive individuals is used, and August 10 if the growth rate sequence of the hospitalized individuals is used) corresponding to the selected value of risk. Finally, the red curves referred to the right-side vertical axis show the mean delay Δ corresponding to the selected risk R (about 3 days if the growth rate sequence of the daily new positive individuals is used, and below 5 days if the growth rate sequence of the hospitalized individuals is used). For clarity, note that the right-side scale for the delay is split into two linear ranges, for a better rendering of the small- Δ range.

11. “Decision Support for the Quickest Detection of Critical COVID-19 Phases.” This Nature Communications paper, by P. Braca, D. Gaglione, S. Marano, L. Millefiori, P. Willett and K. Pattipati, leverages the MAST to rapidly detect the passage from a controlled regime to a critical one. The performance of MAST is investigated for the second pandemic wave, showing an effective trade-off between average decision delay Δ and risk R , where R is inversely proportional to the time required to declare the need to take unnecessary restrictive measures. The risk is determined by the average occurrence rate of false alarms, which could have unnecessary social and economic ramifications. Ideally, the decision mechanism should be as quick as possible for a given level of risk. We find that all the countries share the same behavior in terms of quickest detection, specifically the risk scales exponentially with the delay, $R \sim \exp(-\omega \Delta)$, where ω depends on the specific nation. For a reasonably small risk level, say, one possibility in ten thousand (i.e., unmotivated implementation of countermeasures every 27 years, on the average), the proposed algorithm detects the onset of the critical regime with delay between a few days to three weeks, much earlier than when the exponential growth becomes evident. Strictly from the quickest-detection perspective adopted in this paper, it turns out that countermeasures against the second epidemic wave have not always been taken in a timely manner. The developed tool can be used to support decisions at different geographic scales (regions, cities, local areas, etc.), levels of risk, instantiations of controlled/critical regime, and is general enough to be applied to different pandemic time-series. Additional analysis and applications of MAST are made available on a dedicated website. Figure 13 shows the MAST procedure applied to real data, respectively the day-over-day ratio of new positive reports and hospitalizations.




Figure 14: The left plot supports our contention that the distribution of phrase lengths output from the LZ77 algorithm before and after the change (H and K, respectively); the H and statistics of the original (before LZ)

data are HMMs with the same stationary distributions. On the right we see delay-to-detection and false-alarm rate performances for suggested two post-processing schemes – both are quite good.

12. **“Transient Detection with Unknown Statistics via Source Coding.”** This ICASSP submission, by A. Finelli, P. Willett, Y. Bar-Shalom and S. Marano, notes that quickest detection problems are fairly common in surveillance applications, as framing surveillance alerts as a change in an observation sequence’s statistics is often apt. It considers the scenario where an appropriate statistical description of our observations is *not* available, neither before nor after the transient to be detected. In this vein, the use of the database Lempel-Ziv, or LZ77, procedure, is applied to detect this transient in the observation data. This algorithm is known to produce code “phrase” lengths that are asymptotically distributed as Gaussian random variables, which allows us to form a quickest detection problem around statistics of the coded output. This work specifies procedures to perform source-agnostic transient detection using Locally Optimal (LO) statistic to augment a Page CUSUM test. The work also shows an application to acoustic data. Figure 14 justifies our claim of the phrase lengths following Gaussian distributions, and offers performance plots for two possible approaches. Further results (not yet submitted) show that an appropriate MAST procedure is even better. The approach has been applied very successfully to US, Italian and world data.

IMPACT/APPLICATIONS

Dynamic fusion of data from disparate and non-traditional sources may require design and analysis novel methodologies.

1. For reasons both of robustness and of bandwidth it may be advantageous to transmit data in an unlabeled format. For example, we might consider a sensor network sending coarsely-quantized observations of 1 or 2 bits: does it make sense to accompany such low-fidelity data by far more bits explaining the identity of the sensor? Why not explore a situation in which the identity is not transmitted, but instead inferred by the *fusion* engine? What is lost by doing so? The impact and application are to the design of sensor networks.
2. We explore the extraction of low-observable *dynamic* “targets” from transactional data. Both the targets (for example, a terrorist plot) and the observations (human activities noted) are highly nontraditional. However, target-tracking and low-observable target-extraction (track-before-detect, or TBD) techniques are to some extent mature as applied to classic applications (like radar), and are therefore ripe to be applied in wider venues. We do so, and especially make use of the emerging Bernoulli filter (MBF) paradigm for target extraction. The key in the work reported is to marry the MBF with feature-aided tracking and knowledge-aided tracking, the former of which here is identity (who purchased the chemical fertilizer?) and the latter prior relationship data (noted surveillance of an Air Force facility by a relative of the former actor is something to be wary of). The impact and application are to defense of the homeland and projected forces from nontraditional attacks.
3. Since the beginning of 2020, the outbreak of a new strain of Coronavirus has caused hundreds of thousands of deaths and put under heavy pressure the world’s most advanced healthcare systems. In order to slow down the spread of the disease, known as COVID-19, and reduce the stress on healthcare structures and intensive care units, many governments have taken drastic and unprecedented measures, such as closure of schools, shops and entire industries, and enforced drastic social distancing regulations, including local and national lockdowns. To effectively address such pandemics in a systematic and informed manner in the future, it is of fundamental importance to develop mathematical models and algorithms to predict the evolution of the spread of the disease to support policy and decision making at the governmental level. There is a strong literature describing the application of Bayesian sequential and adaptive dynamic estimation to surveillance (tracking and prediction) of objects such as missiles and ships; and in this paper, we transfer some of its key lessons to epidemiology. We show that we can reliably estimate and forecast the evolution of the infections from daily — and possibly uncertain — publicly available information provided by authorities, e.g., daily numbers of infected and recovered individuals. The proposed method is able to estimate infection and recovery parameters, and to track and predict the epidemiological curve with good accuracy when applied to real data from Lombardy region in Italy, and from the USA. The impact and application are to the health of the US (and Allied) civilian populations, via guidance to local, State and National authorities.
4. If two radar sensors observe the same target their measurements can be combined to produce a *fused* target-state estimate that is of higher quality than that from one radar alone. If there are multiple targets whose information is shared, a necessary first step to fusion is to “assign” each measurement from the first sensor to that at the other in such a way that both refer to the same underlying object, a task generally accomplished by minimizing a global cost involving distance. An assignment error occurs when the measurement originated by target i at the first radar is wrongly associated to a measurement originated by target j (not i) at the second radar.

Naturally, when such an error occurs the result is fusion of information describing disparate objects, resulting in degraded estimation performance and poor self-assessment in terms of posterior uncertainty. Here we address the issue, and derive approximate assignment error probability. The impact and application are to performance prediction for (and thence design of) multi-sensor inference systems.

5. In principle, IFF (Identification Friend or Foe) and AIS (Automatic Identification System) makes covert rendezvous at sea (smuggling, piracy, etc.) impossible; in practice, AIS can be spoofed or simply disabled. Previous work showed a means whereby such deviations can be spotted. Here we play the opponent's side, and describe the least-detectable trajectory that that the elusive vessel can take. The impact and application are to air defense and surveillance via a game-theoretic perspective.
6. Talk delivered to AFRL researchers, December 6th, 2020: "Distributed Detection and Data Fusion." Abstract: The initial paper on the subject of distributed detection, by Tenney and Sandell, showed that under a fixed fusion rule, for two sensors with one bit outputs, the optimal Bayes sensor decision rule is a likelihood ratio test. It has been shown that the optimal fusion rule for N sensors is a likelihood ratio test on the data received from the sensors. Reibman and Nolte and Hoballah and Varshney have generalized the results to N sensors with optimal fusion, again with the restriction of one bit sensor outputs; this has been relaxed later to multi-bit quantizations. In this "primer" talk we explore a number of issues in distributed detection, including some pathologies, the benefits of fusion, optimal design, structures for decision flow, consensus, sensor biases, feedback, deliberate obfuscation (i.e., security) and censoring. We also devote some time to distributed estimation (i.e., fusion for tracking): why is it difficult and what seems to work best?
7. The ML-PMHT has recently emerged as a go-to means to extract very low observable (VLO) target signatures. Previous implementations have worked with measurements that correspond to the tracking space: two-dimensional measurements for two-dimensional targets, and three for three. Under this contract, the technique has been extended to *fusion* of data from non-commensurate data streams. The complication is that ML-PMHT requires optimization of a multimodal likelihood surface, and clever initialization is vital. The key insight here is to trigger hill-climbing from closest approaches of lines-of-sight vectors. The performance is excellent.
8. Many important surveillance problems can be posed as of detection of a change: of network activity, of social and trust relationships, of acoustic signatures, of images, of target behavior, etc. The issue is that optimal procedures require that the statistics before and after the changes be modeled precisely; and that almost trivially limits the applicability of (say) a Page test. A clever approach has recently emerged in the literature, but seems to be little appreciated: use a Lempel-Ziv encoding of the data – whatever the data may be, and really the format matters little – and, since the LZ theme is to look for the longest data patterns ("phrases") that can be seen in the past, to alarm upon a sudden decrease in the coding efficiency (i.e., shorter phrases). Our contribution has been to address the automated alert of a change, and we have developed several procedures for this. First, we leverage results that show asymptotic Gaussianity of the data; and, second, we develop several novel approaches to detect unknown and unmodeled changes in a stream of Gaussian data. The future impact of this technology, if properly promulgated, is very wide: changes of any sort generally presage something of importance to examine more closely (say, an incursion of a network or of a defended space) and this technology delivers a key bell-ringer.

9. Reported as #3 above, we have applied some of our tracking expertise to help authorities predict the future trends of the COVID-19 pandemic, focusing specifically on the β and γ parameters whose ratio, if greater than unity, indicates an out-of-control situation. We have gone further, and applied our signal processing experience with detection of changes to the pandemic. The goal is to detect as quickly as possible any recrudescence of the pandemic so that authorities can institute mitigation policies in a timely manner; and conversely to detect as quickly as possible a return to a more benign state so that public patience can be maintained. There are two key developments we have pioneered. The first is to reformulate the data stream as a ratio of easily available observations such as hospitalizations or new cases; what is good about this is the empirical (and asymptotic, according to theory) Gaussianity of such ratios. The second is new quickest detection update statistic (MAST) that uses a GLR formulation to detect an unknown change in Gaussian data. The potential impact for policy makers is obvious, and our group has worked hard to offer this technology transfer (of ideas from dynamic estimation and from signal processing) to a wide audience of readers and audiences.

PUBLICATIONS AND OUTCOMES

Awards

1. E. Wilthil, Y. Bar-Shalom, P. Willett and E. Brekke, “Estimation of Target Detectability for Maritime Target Tracking in the PDA Framework,” Proceedings of FUSION, Ottawa Canada, July 2019. [Runner-up Best Student Paper.]
2. A. Finelli, Z. Sutton, P. Willett and Y. Bar-Shalom, “Taking Advantage of Group Behavior When Tracking Multiple Threats in Cluttered Surveillance Data,” IEEE Aerospace Conference, Big Sky MT, March 2019. [Best Paper in Conference Track.]

Talks

1. Talk: AFRL Lecture, December 6th, 2019. “Distributed Detection and Data Fusion”. TecEdge, Dayton OH.

Conference

1. A. Finelli, O. Willett, Y. Bar-Shalom and S. Marano, “Transient Detection with Unknown Statistics via Source Coding,” *submitted to ICASSP*, Singapore, May 2022.
2. Z. Sutton, P. Willett and S. Marano, “The Data/Identity Tradeoff with Censored Sensors,” *submitted to ICASSP*, Singapore, May 2022.
3. P. Willett, S. Marano, P. Braca, L. Millefiori, W.D. Blair, P. Miceli, M. Kowalski, T. Ogle, “Expression for the Probability of Correlation Error in Data Fusion,” *IEEE Aerospace Conference*, Big Sky MT, March 2022.
4. M. Brambilla, M. Nicoli, G. Soldi, D. Gaglione, L. Millefiori, P. Braca, P. Willett and M. Win, “Cooperative Localization and Target Tracking by Belief Propagation: A Real Data Analysis,” *MILCOM*, San Diego CA, November 2021.
5. E. d’Afflisio, P. Braca, L. Chisci, G. Battistelli and P. Willett, “Maritime Anomaly Detection of Malicious Data Spoofing and Stealth Deviation from Nominal Route Exploiting Heterogeneous Sources of Information,” *Proceedings of the ISIF Fusion Conference*, South Africa, November 2021.
6. T. Kropfreiter, F. Meyer, S. Coraluppi, C. Carthel, R. Mendrzik and P. Willett, “Track Coalescence and Repulsion: MHT, JPDA, and BP,” *Proceedings of the ISIF Fusion Conference*, South Africa, November 2021.
7. S. Capobianco, N. Forti, L. Millefiori, P. Braca and P. Willett, “Uncertainty-Aware Recurrent Encoder-Decoder Networks for Vessel Trajectory Prediction,” *Proceedings of the ISIF Fusion Conference*, South Africa, November 2021.
8. P. Braca, D. Gaglione, S. Marano, L. Millefiori, P. Willett and K. Pattipati, “When Should Strategic Decisions Be Taken? MAST: A Quickest Detection Procedure For COVID-19 Epidemiological Data,” *European Signal Processing Conference (EUSIPCO)*, Dublin, August 2021.
9. Z. Sutton, P. Willett and S. Marano, “Target Detection from Distributed Passive Sensors: Semi-Labeled Data Quantization,” *ICASSP*, Toronto, Canada, June 2021.
10. P. Willett, P. Braca, S. Marano, L. Millefiori and W.D. Blair, “On the Probability of Association Error in Data Fusion,” *IEEE Aerospace Conference*, Big Sky MT, March 2022.
11. P. Braca, S. Marano, L. Millefiori, P. Willett and W.D. Blair, “On the Probability of Cross-Radar Assignment Error,” *IEEE Radar Conference*, Florence, Italy, September 2020.

12. M. Lexa, S. Coraluppi, C. Carthel and P. Willett, "Distributed MHT and ML-PMHT Approaches to Multi-Sensor Passive Sonar Tracking," *IEEE Aerospace Conference*, Big Sky MT, March 2020.
13. Z. Sutton, P. Willett and S. Marano, "Sensor Network Target Detection with Unlabeled Observations," *IEEE Aerospace Conference*, March 2020.
14. P. Willett, T.L. Ogle, W.D. Blair, P. Miceli and S. Marano, "Testing Unbiasedness of Testing a Sequence of Tests," *CAMSAP*, December 2019.
15. S. Marano and P. Willett, "Shuffled Bits in the Low-Detectability Regime," *Proceedings of the European Signal Processing Conference (EUSIPCO)*, A Coruna Spain, September 2019.
16. T. Ogle and P. Willett, "An Alternative Derivation of Generalized Likelihood Tests for Track-to-Track Correlation," *Proceedings of FUSION*, Ottawa Canada, July 2019.
17. E. Wilthil, Y. Bar-Shalom, P. Willett and E. Brekke, "Estimation of Target Detectability for Maritime Target Tracking in the PDA Framework," *Proceedings of FUSION*, Ottawa Canada, July 2019.
18. S. Marano and P. Willett, "Making Decisions with Shuffled Bits," *ICASSP 2019*, Brighton UK, May 2019.
19. A. Finelli, Z. Sutton, P. Willett and Y. Bar-Shalom, "Taking Advantage of Group Behavior When Tracking Multiple Threats in Cluttered Surveillance Data," *IEEE Aerospace Conference*, Big Sky MT, March 2019.

Journal

1. S. Capobianco, L. Millefiori, N. Forti, P. Braca and P. Willett, "Deep Learning Methods for Vessel Trajectory Prediction Based on Recurrent Neural Networks," *IEEE Transactions on Aerospace and Electronic Systems* [early access].
2. G. Soldi, N. Forti, D. Gaglione, P. Braca, L. Millefiori, S. Marano, P. Willett and K. Pattipati, "Quickest Detection and Forecast of an Epidemic Outbreak: Analysis of the COVID-19 Second Waves," to appear in *IEEE Communications Magazine*.
3. P. Braca, D. Gaglione, S. Marano, L. Millefiori, P. Willett and K. Pattipati, "Quickest Detection of Critical COVID-19 Phases: When Should Restrictive Measures Be Taken?" to appear in *Nature Scientific Reports*.
4. G. Soldi, D. Gaglione, N. Forti, A. Di Simone, F.-C. Daffin`a, G. Bottini, D. Quattrociochi, L. Millefiori, P. Braca, P. Willett, A. Iodice, D. Riccio and A. Farina, "Space-based Global Maritime Surveillance – Part 1: Satellite Technologies," *IEEE AES Systems Magazine*, September 2021.
5. G. Soldi, D. Gaglione, N. Forti, A. Di Simone, F.-C. Daffin`a, G. Bottini, D. Quattrociochi, L. Millefiori, P. Braca, P. Willett, A. Iodice, D. Riccio and A. Farina, "Space-based Global Maritime Surveillance – Part 2: Artificial Intelligence and Data Fusion Techniques," *IEEE AES Systems Magazine*, September 2021.
6. Z. Sutton, P. Willett and Y. Bar-Shalom, "Target Tracking Applied to Extraction of Multiple Evolving Threats from a Stream of Surveillance Data," *IEEE Transactions on Computational Social Systems*, pp. 434-450, April 2021.
7. P. Braca, D. Gaglione, S. Marano, L. Millefiori, P. Willett and K. Pattipati, "Quickest Detection of COVID-19 Pandemic Onset," *Signal Processing Letters*, pp. 683-687, March 2021.

8. A. Aubry, P. Braca, E. d’Afflisio, A. De Maio, L. Millefiori and Peter Willett, “Optimal Opponent Stealth Trajectory Planning Based on an Efficient Optimization Technique,” *IEEE Transactions on Signal Processing*, pp. 270-283, December 2020.
9. D. Gaglione, P. Braca, L. Millefiori, G. Soldi, N. Forti, S. Marano, P. Willett, K. Pattipati, “Adaptive Bayesian Learning and Forecasting of Epidemic Evolution – Data Analysis of the COVID-19 Outbreak,” *IEEE Access*, pp. 175244-175264, October 2020.
10. S. Ye, Y. Bar-Shalom and P. Willett, “Estimation of the Support Parameters of a Uniform pdf and the Cramer-Rao-Leibniz Lower Bound,” *Signal Processing Letters*, pp. 1765-1768, August 2020.
11. S. Marano and P. Willett, “Making Decisions by Unlabeled Bits,” *IEEE Transactions on Signal Processing*, p. 2935-2947, April 2020.
12. Q. Lu, Y. Bar-Shalom, P. Willett F. Palmieri, R. Ben-Dov and B. Milgrom, “Measurement Extraction for Two Closely-Spaced Objects Using an Imaging Sensor,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 55, no. 6, pp. 2965-2977, December 2019.
13. S. Marano and P. Willett, “Algorithms and Fundamental Limits for Unlabeled Detection using Types,” *IEEE Transactions on Signal Processing*, vol. 67, no. 8, pp. 2022-2035, April 2019.
14. Z. Sutton, P. Willett, T. Fair and Y. Bar-Shalom, “MLPMH Tracking in 3 Dimensions Using Cluttered Measurements From Multiple 2-Dimensional Sensors,” *to appear in the Journal of Advances in Information Fusion*.
15. S. Schoenecker, P. Willett and Y. Bar-Shalom, “Feature-Aided Tracking: Does it Help or Hurt?” *submitted to IEEE Journal of Oceanic Engineering*, October 2018.
16. Z. Sutton, P. Willett and S. Marano, “Sensor Network Target Detection with Unlabeled Observations,” *in preparation*.
17. M. Kowalski, P. Willett and Y. Bar-Shalom, “Use of Sensor Opinions in Data Fusion,” *in preparation*.