

Approaches to information fusion with spatiotemporal aspects for standoff and other biodefense information sources¹

Jerome J. Braun^a, Austin Hess^a, Yan Glina^a, Edward C. Wack^a, Karianne Bergen^a,
Timothy J. Dasey^a, Robert M. Mays^b, John Strawbridge^b

^a Lincoln Laboratory, Massachusetts Institute of Technology
244 Wood Street, Lexington, Massachusetts 02420-9185, USA

^b JBSDS, Joint Product Manager, Biological Detection Systems,
ATTN: SFAE-CBD-BD-BDS, Bldg. E3548,
Aberdeen Proving Ground, MD 21010-5424, USA

Officially Approved
for Public Release
by JPEO-CBD

ABSTRACT

This paper discusses some of the techniques developed at MIT Lincoln Laboratory for information fusion of lidar-based biological standoff sensors, meteorology, point sensors, and potentially other information sources, for biodefense applications. The developed Spatiotemporal Coherence (STC) fusion approach includes phenomenology aspects and approximate uncertainty measures for information corroboration quantification. A supervised machine-learning approach was also developed. Computational experiments involved ground-truth data generated from measurements and by simulation techniques that were developed. The fusion results include performance measures that focus explicitly on the fusion algorithms' effectiveness. Both fusion approaches enable significant false-alarm reduction. Their respective advantages and tradeoffs are examined.

Keywords: Information fusion, spatiotemporal coherence, biodefense, detection, machine learning, disparity, uncertainty.

1 INTRODUCTION

Protection against aerosolized biological attacks is an important objective in biodefense, and therefore an important goal for biodefense decision support systems. Prominent among their objectives is the detection of bioattacks. Aerosolized bioattack detection in an outdoor area involves detecting a plume of bioagent that has been released into the air, an event referred to as a *release*. Natural biological substances ubiquitously present in the air are commonly referred to in biodefense context as the *background*. The background levels differ for different settings and are typically dynamic, i.e., they vary for numerous reasons such as environmental and weather conditions, activities within a given setting, and more. From the bioattack detection perspective the background constitutes clutter. The bioattack detection task formulated in this fashion amounts to detecting the release amidst a dynamic background clutter.

¹ This work was sponsored by the Department of the Army under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

* Email: jbraun@LL.mit.edu

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the 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 this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 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 a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2010		2. REPORT TYPE		3. DATES COVERED 00-00-2010 to 00-00-2010	
4. TITLE AND SUBTITLE Approaches to information fusion with spatiotemporal aspects for standoff and other biodefense information sources			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MIT Lincoln Laboratory, 244 Wood St, Lexington, MA, 02420-9108			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This paper discusses some of the techniques developed at MIT Lincoln Laboratory for information fusion of lidar-based biological standoff sensors, meteorology, point sensors, and potentially other information sources, for biodefense applications. The developed Spatiotemporal Coherence (STC) fusion approach includes phenomenology aspects and approximate uncertainty measures for information corroboration quantification. A supervised machine-learning approach was also developed. Computational experiments involved ground-truth data generated from measurements and by simulation techniques that were developed. The fusion results include performance measures that focus explicitly on the fusion algorithms' effectiveness. Both fusion approaches enable significant false-alarm reduction. Their respective advantages and tradeoffs are examined.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Detecting aerosolized bioattack is the first in a progression of decisions that must be made as quickly as possible following the attack. However, the commanders or other decision-makers involved in biodefense are faced with various decisions prior to the attack as well. Examples include resource planning and management decisions such as sensor placement and utilization, various options and contingencies, and more. Once the bioattack has been detected, additional decisions need to be made. Plume mapping involves determination of the plume location, shape and extent. The time progression of mapping results up to and including the most current map can be viewed as plume tracking. The next level of decision-making includes predicting the expected progression of the plume. Plume propagation prediction is also known as forecasting. Progressively higher-level decisions follow, e.g., courses-of-action (CoA) and consequence management. Decision-support systems can assist the commanders and other biodefense decision-makers, whether in the military or homeland-protection realms, in making decisions such as those outlined above.

Information sources exploitable for biodefense decision-support include biological standoff sensors, biological point sensors, as well as a range of other information sources. Point sensors determine the presence and concentration of bioagents at the sensor location. Standoff sensors determine bioagent presence remotely. Biological standoff sensors are often lidar (*light detection and ranging*) based. Their principle of operation involves scanning an area with a laser beam and measuring the backscatter or fluorescence as the beam encounters bioaerosol cloud. In this paper the term standoff and lidar will be used interchangeably to signify a biological standoff sensor.

Both biological point and standoff sensing modalities are intended to sense biological substances, although each modality has its own special strengths. Point sensors may be more sensitive and more specific. Standoff sensors strengths include an ability to provide area coverage and data for all points in that area, and therefore they can be particularly useful for early warning and tracking of bioagent plumes.

Other information sources not directly related to biological sensing realm can be useful for biodefense decision support as well. From the biological sensing viewpoint, both the biological standoff and point sensing technologies have an obvious advantage over non-biological information sources in terms of specificity since the non-biological sources are not intended to provide information regarding biological substances. However, especially in the context of information fusion, non-biological information sources can be quite relevant. They can contribute significantly in terms of the overall system performance improvement. Prominent amongst the non-biological information sources are meteorological data, in particular wind speed and direction, because wind conditions directly affect the transport and dispersion of aerosol plumes. Examples of many other non-biological sensing modalities include various ISR sensors such as radars or electro-optical, acoustic and seismic sensors, and many more. Contextual information such sources may provide can be valuable.

Appropriate information-fusion algorithmic approaches and techniques are a key enabler of synergistic exploitation of multiple information sources.^{3,4,5,6,7,9} As such they offer a great opportunity and promise for biodefense decision-support systems. MIT Lincoln Laboratory is developing and investigating information fusion approaches and techniques for a networked sensing environment that includes biological lidar-based standoff sensors, biological point sensors, and other information sources. This paper discusses selected aspects of that effort, focusing in particular on the aspects related to the detection task.

In particular, the paper discusses fusion-based detection techniques developed with the goal of false alarm reduction. Since low system-level false-alarm rates enable avoiding unnecessary protective measures, such as donning a protective gear, false-alarm reduction is an important goal for biodefense systems. The paper discusses the approaches developed in this effort towards that goal. This includes an approach referred to as Spatiotemporal Coherence (STC) fusion, its adaptive extensions, and a supervised machine-learning approach. The paper discusses their respective advantages and tradeoffs. Measures of performance introduced for use in the approaches developed and for the algorithmic efficacy studies are discussed as well.

While development of information fusion algorithms is obviously a key aspect of this effort, work and advances in several other areas were needed to support that development. In addition to the measure of performance mentioned above, those supporting areas included sensor modeling, threat scenarios, sensing architectures, and techniques for generating ground-truth and multisource data. The ground-truth and multisource data generation techniques introduced and developed in this work are covered briefly in section 2.

This paper has a companion paper². The present paper focuses mainly on the detection task, while the companion paper² extends the discussion to the higher-level fusion tasks of plume mapping (tracking) and propagation prediction (forecasting). Some of the techniques discussed in this paper are also exploited within the approaches we developed for plume mapping and propagation prediction.² The two papers complement one another, in part overlapping and in part covering different elements and aspects of the effort. The present paper assumes readers with information-fusion expertise, including advanced topics such as machine learning, and focuses in a greater depth on the detection-task aspects of this effort. However, since those readers may not necessarily be familiar with the biological and chemical defense domain, the present paper includes some of the basics related to biological and chemical defense and sensing. The companion paper², on the other hand, assumes that the readers are familiar with the CBRNE defense and sensing domain, including areas such as biodefense-related phenomenology and transport and dispersion models, but are not necessarily focused or involved in algorithmic aspects or information fusion. As such, the companion paper² covers the techniques developed for information-fusion based bioattack detection, plume mapping (tracking), and plume propagation prediction (forecasting), essentially at an overview level. There are some overlaps between the two papers. In many cases the overlaps are in terms of topics, but may differ in coverage level of depth. For example, the data-generation techniques are covered in both papers, but the level of coverage is somewhat different. Also, both papers discuss the detection task and the Spatiotemporal Coherence (STC) fusion approach introduced in this effort, but the present paper includes more detail on these topics. While the present paper discusses in more detail the measures of performance introduced in this effort for the detection task, the companion paper² elaborates more on the measures of performance introduced for mapping and propagation prediction. Nevertheless, in the interest of readability and clarity, in some instances the discussion, including the language of some paragraphs (e.g., this paragraph), may be similar in the two papers.

The paper is organized as follows. A brief discussion of the ground-truth and multi-sensor data issue and the techniques to developed to address it is discussed in section 2. Measures of performance proposed in this work are outlined in section 3. Section 4 of the paper focuses on the Spatiotemporal Coherence (STC) approach, its adaptive extensions, and a machine-learning approach. Section 4 also includes example results of these approaches and a comparative discussion, which is followed by the conclusion section.

2 DATA FOR BIODEFENSE INFORMATION FUSION ALGORITHMS DEVELOPMENT

The data needed for this effort included ample amounts of temporal sequences representing two categories – the background category and release-in-background category. The first refers to ambient aerosol conditions, while the second refers to a release plume embedded in a background clutter. For either category, each data exemplar is a time sequence over some duration. These data sequences are referred to in this work as *cases*. Thus the data consist of a set of background cases and release cases. For each time instance within each case, the multisource data in this effort include standoff sensors and/or point sensors as needed for a given part of the work. The standoff sensor data are two-dimensional entities representing the area observed by the standoff sensor referred to as the *field-of-regard*. These two-dimensional entities are needed for all those time-points within the case sequence at which the standoff data are available. Point sensor data are the outputs from those sensors at the times within the case sequence at which the point sensors provide them. Furthermore, each case needs to include data for other required information sources, in particular the meteorology and specifically wind direction and speed. As mentioned, a significant multitude of both background and release-in-background cases are needed. These cases need to be of sufficient diversity, e.g., in terms of release plume and meteorology.

Multisource data are not the only type of data required for this work. Obviously, for detection-task algorithm performance studies, the identities of cases (background or release) are required. However, in this effort the ground-truth concentration data for all points within the spatial area under consideration as a function of time within each case sequence were also needed. There were two main reasons for that additional need. First, these data were required for the work related to higher-level information fusion tasks of plume mapping and propagation prediction, which are discussed in the companion paper². Second, they were needed as input to simulations, discussed below.

In many information-fusion application domains there exist large amounts of experimentally measured and collected data that can be used in the development and studies of information-fusion algorithms. However, in the biodefense realm obtaining data by measurements is a challenge. Backgrounds data can be acquired by placing appropriate sensors

in a given environment and collecting the data. Release data can be obtained by simulant release experiments at test sites, and acquiring sensor data for the test-site simulant releases. In practice, both the background data collections and the simulant release experiments are complex endeavors. Their logistics is far from trivial and they can be effort-intensive and costly. The need for appropriate suites of sensors, including additional sensors that could be deemed to reasonably represent the ground truth, is just one of the many technical challenges. Background diversity is another. Finally, however, even when all these challenges are addressed, for information fusion work the following additional issue needs to be addressed. Experimentally obtained simulant release data represent releases amidst the background clutter conditions prevailing at a test site during the simulant release experiments. For information-fusion studies, however, data representing background clutter conditions other than those of test sites are also desired.

Computational simulations are an alternative to the above experimental approach. In principle, if the ground-truth data representing backgrounds and releases-in-background in terms of concentrations over the area being considered were generated by simulation, computational models of sensors could then be applied to those data, yielding the multisource data corresponding to that ground truth. However, the simulation paradigm presents its own set of challenges.

Computationally efficient plume transport and dispersion models are large-ensemble oriented, i.e., they provide a statistical representation of large ensembles of putative plumes. We argue that the use of such models in the context of information fusion is inappropriate. This is because the multiple sensors observe a specific plume rather than a large-ensemble average. A major strength of the information-fusion approaches developed in this effort is the exploitation of spatiotemporal aspects within the multisource observations of the plume. Large-ensemble representations do not represent those aspects adequately, since those representations are *not* individual plumes.

Simulations of individual plumes, referred to in this paper as *single realization* simulations or models, can be performed using the computational fluid dynamics (CFD) methods. However, CFD simulations are extremely computationally intensive. Therefore, assuming the use of the present standard commercial computing technologies, generating amounts of data large enough to support information fusion work is problematic.

Two alternative solutions to alleviate the multisource and ground-truth data issue were developed in this effort. One of them, referred to as *Release Augmentation* computationally combines release data from a test site with background data collected in another setting.¹ This results in release-in-background data which include the background clutter characteristic of that particular setting rather than that of the test-site where the simulant release was conducted. The Release Augmentation makes significant use of image-processing techniques to determine hard targets and the location of the plume within the scene. Furthermore, it performs transformations that reconcile the differences in wind direction and speed to correctly embed the plume data within the background clutter data.^{2,1} This allows combining various test-site release data with various background data, significantly increasing the utility of available experimental data. Release Augmentation outcomes can be used as ground-truth from which the sensor models can generate the corresponding multisource data.

To generate additional data beyond what can be accomplished with Release Augmentation, a computationally-efficient simulation approach was also developed in this effort, building on and expanding the work reported in our previous publications.^{5,6,3} In comparison to that previous work, the developed approach, referred to as *Meandering Plume and Background Simulation* (MPBS) includes significant additions in the area of turbulence effects and spatial and temporal correlations. Unlike large-ensemble models, MPBS generates single-realization plume-in-background cases. In contrast to its CFD counterparts, however, MPBS is computationally efficient. In particular, its transport and dispersion components involve a mix of basic phenomenology effects and stochastic aspects, similar to the approach we described previously⁵. MPBS, shown schematically in Figure 1a, allows an efficient generation of substantial amounts of cases, and it grants the user a significant flexibility in terms of defining various options such as the release types and parameters, scenery obstacles, and wind conditions. An example of MPBS output, shown in Figure 1b illustrates the effects of plume meandering, stretching and folding, as well as the effects of obstacles. The outputs produced by MPBS provide the desirable ground truth as discussed earlier in this section. These outputs can subsequently be used by sensor models to generate the multisource data counterpart of that ground truth.




Figure 1: Meandering Plume and Background Simulation (MPBS) and simulation example

While an additional discussion of the above methods can be found in the companion paper², the above brief discussion is intended to provide at an overview level the kind of supporting efforts that were essential for the main effort in the area of information-fusion algorithms. Both the Release Augmentation approach and the MPBS approach are advances that can be viewed as significant enablers for information fusion development and studies.

3 EFFICACY REPRESENTATION AND MEASURES OF PERFORMANCE

Development and studies of information-fusion algorithms require appropriate measures of performance. Receiver Operating Characteristics (ROC) curve is a well-established approach to performance evaluation in the detection task. Detection vs. false-alarm rate characteristics portrayed by a standard ROC curve includes the sensor aspects as well as information-processing, i.e., algorithm performance aspects. The ROC curve representation is certainly appropriate and desired when the objective is to evaluate the performance of an entire system including its sensors, algorithms, and other aspects – in those situations all of these need to be taken into account simultaneously.

When comparative evaluation of different algorithmic techniques is an objective, a common approach involves comparing ROC curves generated under the same conditions and on the same data by the different algorithms. While technically correct, this may sometimes lead to misinterpretations. For example, processing very high-quality, easy to discriminate, sensor data with different algorithmic approaches may blur the differences in algorithm performance and make different algorithms erroneously appear similarly promising. Perhaps even more importantly, processing lower-quality sensor data with different algorithms can lead to ROC curves that appear unsatisfactory for either of the algorithms, and that impression may overshadow the potentially important performance differences and benefits of one algorithmic approach vs. another. While such possible misinterpretations can be resolved by a careful analysis within the ROC space confines, erroneous impressions can be an issue especially when communicating across the technical communities.

As this effort is focused on developing and investigating the value of information-fusion algorithmic techniques, it was desirable to decouple algorithmic aspects from others. This was accomplished by introducing an *efficacy curve* representation, shown notionally in Figure 2. Obtaining an insight into algorithm efficacy while sensors may be under continuing development or enhancements is one example of circumstances in which the efficacy curve construct may be particularly useful.

A key notion in the context of the efficacy curve view is that of a reference. The performance of a given algorithm is viewed with respect to that reference in terms of its detection/false-alarm characteristics. Different types of reference are possible. In some applications, not all information sources (sensors) are deemed of equal quality. Application requirements or procedures may, for example, call for a specific “baseline” sensor declaring a detection as a condition for the fusion-based detection (alarm) declaration. The reference in this case is the performance of the baseline sensor. Under this condition, fusing additional sources cannot improve detection beyond that of the baseline sensor, which we refer to in this effort as the *baseline restriction*. However, fusion techniques can in this situation reduce the false-alarm rate. The efficacy curve shows the tradeoff in terms of detection loss for the false-alarm reduction.

For a decision-level fusion scheme, another reference point that yields a similar form of efficacy curve is that of trivial “OR” fusion. Since for a decision-level fusion scheme the “OR” fusion constitutes the upper bounds of the best

detection level and the worst false-alarm level, the tradeoff between the detection loss and false-alarm reduction with respect to the OR fusion bound can be insightful. The efficacy curve corresponding to both of these situations is shown notionally in Figure 2a. In both situations described above the efficacy curve cannot extend beyond the reference point.

Without the baseline restriction, and for feature-level fusion when the “OR” fusion is the reference, the efficacy curve can extend beyond the reference point, as shown notionally in Figure 2b. In this effort we found efficacy curves useful as a representation of the algorithm performance.

Another concept related to measures of performance introduced in this effort is referred to as *elliptical distance*. In the development of information fusion algorithms for the detection task, in particular when attempting to optimize an algorithm, both the detection and false-alarm aspects must be taken into account. This leads to a multi-objective optimization problem. In some applications specific constraints with respect to the detection/false-alarm trade-off may be known, however, in many instances that is not the case. In the ROC space, the equal-error-rate (EER) is the notion that corresponds to an equal value of detections and false-alarms. However, EER is not sufficient for performing optimization in the detection/false-alarm spaces.




Figure 2: Notional efficacy curves

The elliptical distance measure introduced in this work can be viewed as a generalization of the EER. Conceptually, the elliptical contours with varying parameters represent an equivalent performance, i.e., each point on the ellipse is assigned the same cost function value. As shown in Figure 3, the elliptical distance tends to the lowest-performance corresponding to the maximum elliptical distance value of one as the elliptical contours converge to the infinitely large elliptical contour limit coinciding with the main diagonal of the diagram. On the highest-performance end of elliptical distance scale, the elliptical contours converge to the infinitely small contour of elliptical distance limit of zero.




Figure 3: Elliptical distance

3.1 Measures of performance for higher-level fusion tasks

In context of higher-level information fusion tasks beyond the detection realm, additional and different measures of performance are needed. The effort from which this paper stems, in addition to the detection task, also addressed two higher-level information fusion tasks: plume mapping and plume propagation prediction (both can qualify as level-three tasks in terms of the JDL fusion model). For each of these two tasks new measures of performance were required and

introduced in the course of the development effort. The plume mapping and plume propagation prediction tasks and the information fusion methods developed for them are discussed in the companion paper². In this short subsection we briefly outline the mapping task and the mapping measure of performance introduced in this effort, as an example of the considerations and solutions related to measures of performance for higher level information-fusion tasks.

Plume mapping includes determination of the extent and shape of the plume. While biological standoff sensors provide spatial data that can be used for that purpose, our objective was to develop techniques that could improve the plume map quality beyond that offered by a single standoff sensor. The fusion-based plume mapping approach developed in this effort involved multiple standoff sensors and meteorology information. Fusing of multiple standoff sensors can compensate for individual sensor field-of-regard limitations and performance limitations. The developed approach exploits field-of-regard differences, spatial and temporal aspects, and the detection and discrimination events.²

Considering the effectiveness of plume mapping, two possible error types arise: a) missing parts of the plume, and b) including within the map areas that are plume-free. These two aspects are not independent. For instance, accepting an arbitrarily large map-excess error can trivially lead to arbitrarily small plume-miss error as the map becomes sufficiently large.

The measure of performance introduced and used in this effort in conjunction with the fusion-based plume mapping is a two-component vector measure.² One component is the percentage of the plume contained within the map area; the other is the percentage of the map area that is outside the plume area. The first component represents the mapping correctness. The second represents the degree of over-estimation. More details on this mapping measure and its use can be found in the companion paper². The brief discussion provided here is intended to exemplify the new measures of performance in higher-level information-fusion tasks, for which standard measures such as ROC curves are inapplicable. As illustrated by the solution introduced for the plume mapping task, the measures of performance for higher-level information fusion tasks are likely to be problem-domain specific. Tradeoffs that are specific to a given task need to be taken into account, leading to measure constructs that are more complex in terms of dimensionality and degrees of freedom than the traditional detection and identification task measures.

4 INFORMATION FUSION FOR BIOATTACK DETECTION

This section describes the algorithmic approaches developed and investigated in this effort. This includes the Spatiotemporal Coherence (STC) fusion approach,¹ its adaptive extensions, and a machine-learning approach. All of these approaches involve the underlying concept of exploiting spatiotemporal aspects of multisource data for plume detection.^{3,5,6}

The essence of STC is the combining of phenomenology and uncertainty aspects with the goal of quantifying the level of corroboration between the different available elements of disparate information, such as biological standoff, point, and meteorology sensors. In the context of decision-level fusion for the detection task this involves corroborating an alarm issued by a sensor with the information from other sensors or sources. When a sufficiently high level of corroboration in terms of spatial and temporal aspects is not reached, the individual sensor alarm can be assigned lower confidence, resulting in a better false-alarm rejection by the fusion system as compared to the individual sensor.

In Figure 4a the principle of STC is illustrated on an example architecture that includes standoff, point, and meteorology sensors. It is emphasized, however, that the principle discussed in that context is applicable to other sensing architectures. For the multiple-standoff and meteorology architecture, it is sufficient in the discussion below to rename point-sensor alarms as standoff-sensor alarm indications at the locations corresponding to point-sensor locations in the discussion below.

Suppose that the standoff sensor reports a detection event at location x_1 at time t_1 , and that subsequently a point sensor located at $x_2 \neq x_1$ issues an alarm at $t_2 > t_1$. Using the wind velocity (direction and magnitude) and information for other alarms, the STC modulates the confidence that the point sensor alarm represents a true release rather than a false alarm.

STC uncertainty profiles are shown notionally in Figure 4b. The confidence of corroboration between, say, a lidar detection event at time t and a point sensor detection event at some later time $t+\Delta$ is the highest when Δ is small, and it decreases sigmoidally. The spatial coherence functions are Gaussians with varying parameter values. A projected alarm is defined as the alarm point determined from prior actual alarm location and wind information. As the distance between a projected alarm and the new actual alarm increases, the spatial coherence between the current and previous actual alarms decreases. Three instances of spatial coherence function are shown in Figure 4b. The spatial coherence function amplitudes decrease as the time between the detection events increases.




Figure 4: Spatiotemporal Coherence approach: principle and example

The principle of STC approach is further illustrated in Figure 5. Here the STC processes two separate toy-problem situations that are similar in terms of sensor indications but a bioattack occurs only in the first situation. They may be thought of as occurring on different dates. Each of the two situations is observed by a single standoff sensor and a *single* point sensor. In the first situation the point sensor is at location 1, in the second at location 2. The wind conditions are similar for both situations. The lidar standoff sensor is located at the point where its field-of-regard boundary lines in the figure connect. Two point sensor locations are depicted as rectangles.

Suppose that the standoff lidar reports bioattack detection for location marked by the diamond symbol, same in both situations. Also, in each of the two situations the respective point sensor (at location 1 or 2, respectively), after a certain time after the standoff detection event, issues an alarm.

The illustration images in Figure 5 represent four snapshots starting with Figure 5a, which corresponds to the time at which the standoff lidar alarm occurs. Figure 5b and Figure 5c show the STC uncertainty profile progressing consistent with the wind conditions, evolving and expanding during the progression. Figure 5d corresponds to the time at which a subsequent point-sensor alarm occurs. The magnitude of the STC coherence, i.e., the level of point-sensor alarm corroboration by the preceding standoff alarm and meteorology (wind) information is shown in Figure 5d as the grayscale intensity of the point-sensor symbols. It is evident that the bioattack confidence is much higher for situation 1 than for situation 2. The confidence results are shown as bars in Figure 5d. As indicated, the STC issues the fused alarm in situation 1, but suppresses the sensor alarms in situation 2 considering them as sensor-level false alarms. Thus, regardless of the fact that standoff and point sensors issued the same indications in both situations, STC correctly determines a bioattack occurrence in situation 1 and correctly discards sensor alarms as false in situation 2.

The STC approach was tested for different sensing architecture, including multiple standoff and meteorology fusion as well as standoff, point, and meteorology fusion. It is noted as a sideline that in the earlier stage of this effort a limited preliminary study of adding other "contextual" non-biological sources such as seismic/acoustic or ISR, which were simulated at a rudimentary decision level, was also carried out.




Figure 5: STC for two hypothetical situations

The STC false-alarm reduction potential is shown in Figure 6 for a two biological standoff and meteorology sensing architecture. Significant gains were also observed for standoff/point/meteorology sensing architecture. Meandering Plume and Background Simulation (MPBS) data were used as the ground truth for the computational experiment shown in Figure 6. They included multiple release and background cases, with varying wind conditions, plume and background conditions. These ground-truth data were subsequently processed by the sensor models yielding outputs for two non-collocated standoff sensors. Those outputs together with the meteorology data constituted the inputs to the fusion process.




Figure 6: False-alarm reduction using STC fusion approach

Figure 6 shows the STC performance in comparison to the simple OR fusion of the two standoff sensors. For the decision-level fusion, clearly the OR fusion constitutes the upper bound on detection performance since all sensor detections lead to the OR-fused system-level alarm. However, as is well known, the OR fusion also yields the worst case false-alarm performance, as each sensor-level false alarm results in the OR-fused system-level false alarm. As Figure 6 indicates, for this dataset the STC detection level, while somewhat lower than the OR-fusion detection upper bound, is not much lower than that upper bound. However, the STC false-alarm rate is decreased significantly. This shows the false-alarm reduction potential of the STC fusion approach developed in this effort.

The intuitive character of the STC fusion approach, and its foundation in the fundamental phenomenology aspects, make it attractive. Its performance, discussed in section 4, is encouraging. However, STC in its basic form presents a number of developmental challenges. In particular, various STC aspects such as the uncertainty profile parameters need to be carefully designed and tuned. Moreover, some parameters are likely sensitive to various environmental and application specifics. Therefore, they may need to be adjusted over time or for different conditions such as in case of different system-deployment settings.

Two solution paths were pursued to address the above challenges. The essence of the first, shown in Figure 7a, involves tuning and retuning STC parameters automatically from data acquired during the fusion system operation. In particular, false-alarm cases can be used to further improve the STC performance. This amounts to an optimization process. The swarm optimization approach was used, as shown schematically in Figure 7b. The objective function was the elliptical distance measure introduced in section 3.

While a careful manual tuning of STC parameters by the developer could be competitive with or outperform an automatic tuning approach, our goal in this part of the effort was to develop provisions that do not depend on the developers' availability and expertise. Therefore, comparing adaptive STC performance to that achievable via manual tuning would not be appropriate. Rather, some form of naïve tuning reference should be used. In the initial experiments this naïve reference was mimicked by a random selection of STC parameter values. This can be viewed as representative of a non-expert tuning, such as the reliance on some nominal parameter default values which may not be suitable for a specific application. In the initial experiments with the above tuning method, the available multisource data were split into optimization and test datasets. The measure used for estimating the level of STC performance was the elliptical distance over the dataset. Optimizing on the test dataset can in this case be viewed as the best-performance bound attainable on this dataset with the particular optimization algorithm. The tuning performance experiments, however, obviously involved tuning on the optimization dataset and testing on the test dataset. Those initial experiments indicated that in terms of minimizing the elliptical distance the proposed approach can offer performance closer to the lower bound optimum described above.

Initial adaptation-related experimentation involved withholding part of the optimization dataset and a subsequent re-optimization using all data including those previously withheld. That initial experimentation showed that re-adaptation with additional data can further improve STC performance. This indicated that, as new cases are assimilated during the fusion system operation, adaptive STC approach discussed above can indeed make use of those new data towards improving its performance.




Figure 7: Adaptive STC extensions

The second path pursued in this work involved supervised machine learning. The approach is shown schematically in Figure 8. Adaptation in this approach involves in essence a re-training as new data become available.




Figure 8: Machine learning approach

Designing robust features is one of the key aspects that enables or limits the performance of machine-learning algorithms. Some of the features we used were inspired by our STC approach. Machine learning methods included mainly feedforward neural networks, although some experiments involved Support Vector Machine techniques^{11,8} as well.

The results in Figure 9 show a comparative performance of the neural network based approach and STC, for a sensing architecture involving two biological standoff sensors and meteorology.

The data for this computational experiment were numerically and content-wise different from those used to generate the results shown in Figure 6. The parallels were in terms of the approach to generating the data needed for the construction of datasets. Meandering Plume and Background Simulation (MPBS) data were used as the ground truth. The data included multiple release and background cases with varying wind, plume and background conditions. These ground-truth data were processed with the sensor models yielding outputs corresponding to two non-collocated standoff sensors. Those outputs together and meteorology data were used to construct the training and testing datasets.

For the computational experiment that yielded the results shown in Figure 9 the training dataset was constructed from half of the data. A feedforward neural network was trained using the Levenberg-Marquardt procedure.¹⁰ The other half of the data were used for the test dataset construction on which the network was tested. The STC results shown in Figure 9 represent STC performance on the same test dataset.




Figure 9: STC vs. machine learning

The bars in Figure 9 show the false-alarm and detection performance of the simple OR fusion, the STC fusion, and the neural network based fusion approaches. It should be noted that results such as those in Figure 9 depend on various factors, including overall sensing configuration, sensor positions, hard targets, and other specifics that underlie the data from which a given dataset is constructed. For the dataset used to generate results in Figure 9, the STC performance is slightly different than for the one shown in Figure 6. It is evident that for this dataset, while STC reduces false alarms substantially in comparison to the OR-fusion reference, its detection level is somewhat reduced as well. On the other hand, the neural-network based approach reduces false alarms even below the significantly reduced STC level, while the detection level is appreciably closer to the OR fusion upper bound than that offered by STC for this particular dataset.

The benefits of machine learning approach over STC discussed above may appear moderate. It should be pointed out, however, that the role of machine learning extends beyond the results reported here. In the computational experiments reported here STC performed very well. It can be expected, however, that particularly difficult datasets would pose greater challenges to STC.

In particular, adaptability of the machine learning approaches is expected to be difficult to match with non-learning techniques. This includes a potential of significantly higher robustness to challenges such as unexpected data, changes in environmental conditions, and changes of deployment settings. Computational experiments that would demonstrate these expected advantages require datasets that are particularly oriented to represent such challenges, including datasets that represent significant changes in environmental conditions or settings. Development of such datasets was beyond the scope of this effort. This effort also did not exploit such advanced machine-intelligence based approaches as those we proposed in [3]. The fundamental characteristics of the machine learning paradigm strongly suggests that, beyond the gains reported in this paper, machine learning techniques constitute a particularly promising and essential path for decision-support systems.

Developing machine learning approaches does demand a high level of developer expertise.⁷ The ubiquitous availability of various machine-learning toolkits encourages their use in a "cookbook" manner. This can lead to disappointing

results and should be avoided. In-depth background and knowledge in machine-learning, as well as a significant experience with applying machine-learning techniques, are required of the developers. These prerequisites can lead to robust and powerful machine-learning based solutions.

A final important aspect that must be discussed in connection with the algorithmic approach tradeoffs returns us to the problem of data availability. In the context of machine learning this requires a particular attention. Conventional machine learning solutions often require large amounts of data for training. Algorithm developer expertise mentioned in the previous paragraph is crucial for that aspect as well. Data-related solutions such as those we developed and discussed in section 2, and appropriately designed machine-learning fusion approaches, can both mitigate data scarcity related issues. Nevertheless, in applications in which data scarcity remains a significant constraint, it may be desirable to consider alternatives. Developer-tuned STC, for example, may be competitive with a machine-learning based approach if data scarcity is significant. This is because in that case the human designer choices, e.g., regarding the STC uncertainty profile parameters, in a sense encode human insight. This may be another argument for hybrid approaches to information fusion such as the approach proposed by us previously in [3,4]. Based on the work reported here, the STC approach, its adaptive extensions, and the machine-learning based fusion approaches are all promising for the purposes discussed in this paper, as separate techniques and possibly also as components of hybrid algorithmic constructs.

5 CONCLUSION

This paper discussed some of the techniques developed at MIT Lincoln Laboratory for information fusion of lidar-based biological standoff sensors, meteorology, point sensors, and potentially other information sources, for biodefense applications. It focused on the detection task, in particular addressing the goal of false-alarm reduction. An approach developed for that purpose, referred to as the Spatiotemporal Coherence (STC) fusion, was discussed. Its highlights include combining phenomenology aspects and approximate uncertainty measures for information corroboration quantification. A supervised machine-learning approach that was also developed was discussed as well. Selected measures of performance introduced in this effort were presented. In addition, techniques for generation of ground-truth and multisource data developed as part of this effort were briefly discussed as well. Discussion of other areas of the effort, in particular the higher-level information fusion tasks of plume mapping and plume propagation prediction and the approaches developed for those tasks, can be found in the companion paper². Results of the computational experiments discussed in this paper indicate that STC and the supervised learning approach both enable a significant false-alarm rate reduction.

REFERENCES

1. J. J. Braun, Y. Glina, A. Hess, T.J. Dasey, E.C. Wack, R.M. Mays, J. Strawbridge, "Information fusion of biological standoff, point, and other Information sources for bioattack detection," *Proc. Chemical and Biological Defense Physical Science and Technology Conf.*, November, 2008.
2. J. J. Braun, A. Hess, Y. Glina, E. C. Wack, K. Bergen, T.J. Dasey, R.M. Mays, J. Strawbridge, "Information fusion of standoff and other information for biodefense decision support," *Proc. of SPIE*, vol. 7665, 2010.
3. J. J. Braun and Y. Glina, "Hybrid methods for multisource information fusion and decision support," *Proc. of SPIE*, vol. 6242, 2006.
4. J. J. Braun, Y. Glina, L. Brattain, "Fusion of disparate information sources in a hybrid decision-support architecture," *Proc. of SPIE*, vol. 6571, 2007.
5. J. J. Braun, Y. Glina, J. K. Su, "Urban Biodefense Oriented Aerosol Anomaly Detection using Sensor Data Fusion," *Proc. 1st Joint Conf. Battle Mgt. Nuclear, Chem., Biolog. and Radiological Defense*, November, 2002.
6. J. J. Braun, Y. Glina, J. K. Su, T. J. Dasey, "Computational Intelligence in Biological Sensing," *Proc. of the SPIE*, vol. 5416, 2004.
7. J. J. Braun, Y. Glina, A. Hess, T.J. Dasey, E.C. Wack, "Information fusion for CB defense applications: Challenge and opportunity," *IEEE Conf. on Technologies for Homeland Security*, May 2009.
8. J. J. Braun, "Sensor Data Fusion with Support Vector Machine Techniques," *Proc. of the SPIE*, vol. 4731, 2002.
9. T.J. Dasey and J.J. Braun, "Information Fusion and Response Guidance", *Lincoln Lab. J.*, vol. 17, no.1, 2007.
10. S. Haykin, *Neural networks: a comprehensive foundation*, Prentice-Hall, Inc., Upper Saddle River, N.J., 1999.
11. V. N. Vapnik, *The Nature of Statistical Learning Theory*, Springer-Verlag New York, Inc., New York, NY, 2000.