

Decision Level Fusion: An Event Driven Approach

Siddharth Roheda

*Department of Electrical and Computer Engineering
North Carolina State University
Raleigh, NC, USA
sroheda@ncsu.edu*

Hamid Krim

*Department of Electrical and Computer Engineering
North Carolina State University
Raleigh, NC, USA
ahk@ncsu.edu*

Zhi-Quan Luo

*Department of Electrical and Computer Engineering
University of Minnesota
Minneapolis, MN, USA
luozq@ece.umn.edu*

Tianfu Wu

*Department of Electrical and Computer Engineering
North Carolina State University
Raleigh, NC, USA
tianfu_wu@ncsu.edu*

Abstract—This paper presents a technique that combines the occurrence of certain events, as observed by different sensors, in order to detect and classify objects. This technique explores the extent of dependence between features being observed by the sensors, and generates more informed probability distributions over the events. Provided some additional information about the features of the object, this fusion technique can outperform other existing decision level fusion approaches that may not take into account the relationship between different features.

Index Terms—Sensor Fusion, Decision Level Fusion, Event based Classification, Coupling

I. INTRODUCTION

One often resorts to fusing data from more than one sensors to successfully carry out target detection and classification. Data fusion has, over the years been recognized to distinguish three levels of fusion, namely data level, feature level, and decision level. Data level fusion generally processes raw data generated by each sensor, and performs the fusion of information according to some criterion before proceeding to inference. Feature level fusion, on the other hand, first gleans features from raw data (e.g., transformed data) from diverse sensors, to subsequently coherently merge them for inference. In decision level sensor fusion, each sensor reaches an individual decision, prior to optimal combination of the decisions to yield a more informed decision, e.g. target classification.

In this paper, we present on the decision level fusion strategy, and propose a principled approach to fusion for improved inference performance. A classification decision is reached by cataloguing sets of events along with their probabilistic characterization for each sensor, and following a joint probabilistic and coherent evaluation of these events. These events are formalized for each sensor according to its potentially extracted features to define different targets (e.g. target radar cross section less than $10m^2$ and velocity greater than $10m/s$). What this in effect achieves, as a result, is a probability measure assignment to a specific target following its description as an algebraic definition of these feature events.

In our case, we will study two different datasets. The first one, combines a Radar sensor with an optical sensor. A radar used to explore the velocity of an object among other things, thus defining a sample space and a Sigma-Field with an associated probability measure, will be coupled to a telescopic sensor with also an associated probability space. This product space thus allows us to define a principled fusion framework where fusing often yields improved and robust performance. Similarly, the second dataset will involve a seismic sensor, coupled with an acoustic sensor.

II. BACKGROUND AND RELATED WORK

As noted earlier, sensor fusion has long been of interest, albeit with limited theoretical success particularly when heterogeneous data is present and for which a unified and systematic approach has remained elusive. An introduction and comprehensive survey to the area of fusion is provided in [1], [2]. Classical techniques in decision level fusion include Bayesian Inference [3] and Dempster-Shafer fusion [4]. Bayesian Inference has received a lot of attention and has shown success if prior knowledge about sensor reports is available and given. On the other hand, Dempster-Shafer Fusion was proposed to specifically lift such a restriction on information prior, at a cost of a substantial increase in computational complexity. Some applications of Dempster-Shafer Fusion can be found in [5], where, LIDAR data is combined with Multispectral imagery, and in [6], where, multi-sensor information like vibration, sound, pressure, and temperature is fused to detect engine faults. Furthermore, [7] provides a detailed comparison of Bayesian Inference and Dempster-Shafer Theory. In [8], a two-stage approach to fusion was proposed, involving knowledge-modeling, which learns from past behavior of the classifiers whose results are to be fused, and operation stage, that combines output of these classifiers based on knowledge learned in the first stage. More recent work proposes a sensor model based approach [9], where the sensor network is modeled as either that of similar or dissimilar sensors. Similar Sensor Fusion or Convex Quadratic

Approved for Public Release: 18-MDA-9560 (22 March 18)

Fusion [9]–[11] is used when all the sensors explore the same characteristics/features of the target (for example, a set of 5 radars, looking at the same target), while Dissimilar Sensor Fusion or Analytic Center Fusion [9], [11] is alternatively used when sensors explore different characteristics/features of the target (for example, a radar and optical sensor looking at the same target). These assumptions turn out to be too restrictive, in that some sensors, albeit dissimilar, may have some common features, while offering additional features to enrich an object/target characterization. Our goal is to explore such a case, and demonstrate that a systematic and principled approach maybe designed, and the overall solution is improved on account of this enhancement.

A. Mutual Information

Mutual Information (MI) has been an important and ubiquitous metric in many applications in data science, with a goal of revealing the synergy/differential between the dependence of two random variables and their lack thereof, by way of their associated probability density functions (PDF). Specifically, consider two random variables, X and Y , with a joint probability mass $p(x, y)$ and marginal probability mass functions $p(x)$ and $p(y)$. The Mutual Information, $I(X; Y)$, is the relative entropy between the joint distribution, $p(x, y)$, and the product distribution, $p(x)p(y)$ [12] as given by

$$I(X; Y) = \sum_{x \in X} \sum_{y \in Y} p(x, y) \log \frac{p(x, y)}{p(x)p(y)} \quad (1)$$

This relationship may alternatively and perhaps more revealingly, be expressed in terms of the variables' joint entropy, their marginal entropies and their mutual information [12],

$$H(X, Y) = H(X) + H(Y) - I(X; Y) \quad (2)$$

III. PROBLEM FORMULATION

We assume throughout a set of targets/objects, $O = \{o_1, o_2, \dots, o_I\}$, whose detection and/or classification are of interest. Let the k^{th} feature observed by the l^{th} sensor be F_k^l . Then, a set of mutually exclusive events, $\Omega_k^l = \{a_{kj}^l\}_{j=1, \dots, J_{kl}}$, may be defined for the feature F_k^l . Here, a_{kj}^l is the j^{th} event for F_k^l and is described as, $a_{kj}^l : F_k^l \in [u_j, v_j]$, $u_j \in \mathbb{R}^+$, $v_j \in \mathbb{R}^+$, and $v_j > u_j$. The probability report for the k^{th} feature from the l^{th} sensor is then defined as,

$$D_k^l = \{\Omega_k^l, \sigma_B(\Omega_k^l), P_k^l\} \quad (3)$$

Where, $\sigma_B(\Omega_k^l)$ is the Borel sigma algebra of Ω_k^l , and can be thought of as the set of all the possible events that can be described over the feature. P_k^l is the set of probabilities over the events in $\sigma_B(\Omega_k^l)$.

Since, we define objects as a combination of certain events occurring over different features, we will be working in the product space,

$$\Omega = \Omega_1^1 \times \Omega_2^1 \times \dots \times \Omega_{K_1}^1 \times \Omega_1^2 \times \Omega_2^2 \times \dots \times \Omega_{K_2}^2 \times \dots \times \Omega_1^L \times \Omega_2^L \times \dots \times \Omega_{K_L}^L \quad (4)$$

Where, K_l is the total number of features observed by the l^{th} sensor, and $l = 1, \dots, L$. Further, an object will be

defined as some combination of events in this product space, $o_i \in \sigma_B(\Omega)$. Given the object definitions and the probability distributions over various features, we then want to find the fused probability report over the objects, $D_f = \{O, P_f\}$.

IV. APPROACH

The sensor reports form a set $\{D_k^l\}_{k=1, \dots, K_L}^{l=1, \dots, L}$, that come from different sensors, and the target describing events may also be specified over different features. Further, the definitions of specific objects are the result of algebraic operations on the event space $\sigma_B(\Omega)$, a Sigma-algebra on the product space, Ω , with associated probability measures as noted in Section III. Thus, we must determine probability distribution on $\sigma_B(\Omega)$.

A. Determining the joint probability distribution

This discussion is in a 2-dimensional setting, in the interest of simplicity, but can be easily extended to an arbitrary number of reports. Let the reports $D_k^l = \{\Omega_k^l, \sigma_B(\Omega_k^l), P_k^l\}$ and $D_j^m = \{\Omega_j^m, \sigma_B(\Omega_j^m), P_j^m\}$ be defined over the k^{th} feature from the l^{th} sensor, and the j^{th} feature from the m^{th} sensor respectively. Then, we have the product space, $\Omega = \Omega_k^l \times \Omega_j^m$, and we wish to determine a probability distribution on $\sigma_B(\Omega)$.

When determining the joint probabilities, it is important to account for the extent of dependence between the features: Completely independent features yield minimal mutual information, and the joint distribution with minimum mutual information should be selected; a high correlation between features, on the other hand, yields maximal mutual information, and the joint distribution with maximal mutual information should be selected. These are clearly the extreme cases of dependence, and do not address the partial correlation case. A good approximation to determine the joint probability between partially correlated features is to consider a convex combination of the joint probabilities maximizing and minimizing the mutual information. The joint probability for the occurrence of events, $\alpha \in \sigma_B(\Omega_k^l)$ and $\beta \in \sigma_B(\Omega_j^m)$, can then be determined as,

$$P_\Omega(\alpha, \beta) = \rho \cdot P_{\Omega_{\text{MAX MI}}}(\alpha, \beta) + (1 - \rho) P_{\Omega_{\text{MIN MI}}}(\alpha, \beta), \quad (5)$$

where, $\rho \in [0, 1]$ is a pseudo-measure of extent of correlation between the features. $\rho \approx 1$ when the features are highly correlated, and $\rho = 0$ when features are independent of each other. ρ can be determined from the training data, by computing the correlation between the features using a measure like Pearson's correlation or distance correlation [13].

It can readily be seen from Equation 1 that mutual information between two random variables is minimized when the joint probability distribution is selected as the multiplication of the marginal probabilities. Hence, we have,

$$P_{\Omega_{\text{MIN MI}}}(\alpha, \beta) = P_k^l(\alpha) * P_j^m(\beta). \quad (6)$$

The problem of maximizing mutual information given the marginal probabilities requires some steps. For some random variables X and Y , conditioning on the marginal probability distributions of X and Y yields constant $H(X)$ and $H(Y)$. As may be seen from Equation 2, the maximization of Mutual Information between two random variables then becomes

equivalent to minimizing their Joint Entropy, which is known to be a concave function.

$$\begin{aligned}
P_{\Omega \text{ MAX MI}} &= \min_{P_{\Omega}} \sum_{x \in \Omega_k^l} \sum_{y \in \Omega_j^m} -P_{\Omega}(x, y) \log P_{\Omega}(x, y) \\
\text{subject to: } &\sum_{x \in \Omega_k^l} P_{\Omega}(x, y) = P_j^m(y), \\
&\sum_{y \in \Omega_j^m} P_{\Omega}(x, y) = P_k^l(x), \quad P_{\Omega}(x, y) \geq 0.
\end{aligned} \quad (7)$$

A greedy approach suggested for minimizing joint entropy given the marginal probabilities in [14] is used in our implementation. The main idea here is to keep large probability masses intact and not break them down into smaller chunks. The contribution of a probability mass toward the joint entropy only increases if it is divided into smaller chunks. That is, for $p = a + b$, $-p \log(p) \leq -a \log(a) - b \log(b)$, when $0 < p < 1$ and $a, b > 0$. So, keeping the large probability masses from given marginal probabilities intact ensures that their contribution toward the joint entropy is minimized. As empirically demonstrated in [14], the minimal joint entropies are obtained to within 1 bit of the optimal values.

B. Determining Object Probabilities

Once the probability distribution over the product space, Ω , has been determined, the next step is to determine the object probabilities, $P_f(o_i)$, in the fusion report, $D_f = \{O, P_f\}$. Given the combination of events describing an object, it is pretty straight forward to determine the probability of the presence of an object using the general rules of probability. For an object, o , which is to be described in terms of a combination between $\alpha \in \Omega_k^l$ and $\beta \in \Omega_j^m$, we can have either $o : \{\alpha \wedge \beta\}$ or $o : \{\alpha \vee \beta\}$. Where, $\alpha \wedge \beta$ denotes the combination α 'and' β , while $\alpha \vee \beta$ denotes the combination α 'or' β . Then, for the first case we have, $P_f(o) = P_f(\alpha \wedge \beta) = P_{\Omega}(\alpha, \beta)$ and for the latter we have, $P_f(o) = P_f(\alpha \vee \beta) = P_k^l(\alpha) + P_j^m(\beta) - P_{\Omega}(\alpha, \beta)$. This can be extended to any combination of more than two events.

V. EXPERIMENTS AND RESULTS

To substantiate the above approach along with the various steps, we select two different datasets.

A. Dataset 1

For the first one, we select two sensors, namely a Radar sensor and a telescopic optical sensor, for one (latter) of which we have acquired real data. Due to technical difficulty, our Radar measurements were never co-measured with the optical data. Both the sensors are ideally synchronized when observing a given target, which in our case, can be any object in outer space, such as satellite or space debris. For radar simulations we use MATLAB Simulink, and for telescopic image simulation we use some existing telescopic images collected by The Czech Technical University in Prague, and model the object distributions from those images to generate more images with required object characteristics. Each generated

radar signal over one second is correlated with two telescopic images. Samples for objects with different velocities, cross-sections, ranges, and aspect-ratios are then generated. The radar signals are used to make decisions over velocity, range, and the cross-section, while the telescopic images are used to make decisions over the aspect-ratio, and displacement over time of an object in view.

1) *Implementation Details:* To proceed with evaluation, we generate the probability distribution over the features that we will use. Figure 2 shows a high level block diagram for implementation on the first dataset.

Let the received radar signal be, $x(n)$, and its corresponding Fourier transform, $X(k) = \sum_{n=0}^{N-1} x(n)e^{-i2\pi kn/N}$ for each object, whose labels are distinguished with the velocity, range, and cross-section values of that object, $[v, r, cs]$. Using the training data and labels, SVM classifiers using the Radial Basis Function kernel are trained over the events of interest defined over $[v, r, cs]$, and used to determine a classification score over these events when a test sample is fed in, as shown in Figure 2. The classification probabilities are then approximated from the classification score using Platt's technique [18]. These probabilities over the events of interest are the, P_k^l , described earlier in Section III.

We have two telescopic images associated with 1-sec of radar return for the same object. Due to space limitations, we forego the details of target detection which is not central to our paper. The object of interest is initially detected by using target detection as discussed in [19]. Upon its detection, the probability distribution over the object's aspect-ratio, and its displacement in the second image relative to its location in the first image is determined using the image flow technique discussed in [20].

2) *Event and Object Definitions:* For training and testing purposes, we define various events over the feature-sets from both the sensors. For the radar, as noted before, we use $[v, r, cs]$ and the events are defined as,

$$\begin{aligned}
a_1^v : 0 \leq v \leq 10 \text{ m/s}, \quad a_2^v : 15 \text{ m/s} \leq v \leq 35 \text{ m/s}, \\
a_1^r : 0 < r \leq 300 \text{ m}, \quad a_2^r : 300 \text{ m} < r, \\
a_1^{cs} : 0 < cs \leq 20 \text{ m}^2, \quad a_2^{cs} : 15 \text{ m}^2 \leq cs \leq 50 \text{ m}^2.
\end{aligned} \quad (8)$$

From the telescopic imaging sensor, the features, displacement and aspect ratio, $[d, AR]$ define the following events,

$$\begin{aligned}
a_1^d : 0 \leq d \leq 60 \text{ pixels}, \quad a_2^d : 90 \text{ pixels} \leq d \leq 210 \text{ pixels}, \\
a_1^{ar} : 0 < AR \leq 1.5, \quad a_2^{ar} : 1.5 < AR.
\end{aligned} \quad (9)$$

Furthermore, the objects for classification are defined in terms of these events as,

$$o_1 : \{a_1^r \wedge [(a_2^v \wedge a_2^d) \vee (a_2^{cs} \vee a_2^{ar})]\} \quad (10)$$

$$o_2 : \{a_1^v \wedge a_1^d \wedge a_2^r \wedge a_1^{cs} \wedge a_1^{ar}\} \quad (11)$$

Given these events and object definitions, we wish to determine the fused report, $D_f = \{P_f(o_1), P_f(o_2), P_f(\overline{o_1 \vee o_2})\}$. This can be thought of a classification problem with 3 classes, Class 1: Object 1, Class 2: Object 2, and Class 3: Neither Object 1 nor Object 2.




Fig. 1: ROC Curves for detection (Dataset 1) of (a): Class 1: Object 1, (b): Class 2: Object 2, and (c): Class 3: Neither Object 1 nor Object 2




Fig. 2: High level block diagram of the implementation, for the first dataset




Fig. 3: High level block diagram of the implementation, for the second dataset

B. Dataset 2

The second dataset we used is pre-collected data from a network of seismic sensors, and acoustic sensors deployed in a field, where people/vehicles were walking/driven around in specified patterns. Details about this sensor setup and experiments can be found in [15]. This dataset has been previously used for target detection in [16], [17], where, the authors focus on detection of humans. Here, we use this dataset to classify between humans, vehicular targets, and no targets.

1) *Implementation Details*: Figure 3 shows the high level block diagram for the implementation on the second dataset. Using the training data, SVM classifiers are trained over the corresponding events of interest, as discussed before for the first dataset in Section V-A1. The seismic sensor provides decisions over the features, target weight and target speed, $[w, s]$. True labels for weights of the targets are provided in the dataset, while that for target speed are obtained from the GPS data of the target. Similarly, the acoustic sensor provides decisions on the noise-level of the target, and the target speed, $[n, s]$. The two decisions over the target speed are combined into a single report by performing weighted averaging of the decisions of the two sensors. Here, the weights are selected as per the individual accuracies of the SVMs trained to detect events on target speed.

2) *Event and Object Definitions*: For training and testing purposes, we define various events over the feature-sets from both the sensors. For the seismic sensor, we use $[w, s]$, while for the acoustic sensor we use $[n, s]$.

$$\begin{aligned} a_1^w : 96.08 \text{ pounds} \leq w \leq 230.61 \text{ pounds}, \\ a_2^w : 1311.61 \text{ pounds} \leq w, \\ a_1^s : 0.37 \text{ m/s} < s \leq 2.12 \text{ m/s}, a_2^s : 1.7 \text{ m/s} \leq s, \\ a_1^n : n \leq -30 \text{ db}, a_2^n : -10.6658 \text{ db} \leq n \leq 7.84 \text{ db}. \end{aligned} \quad (12)$$

The range of an event can be determined from the training data. The mean of the feature in question over the samples of the same class is computed, and a range of two times the standard deviation is taken on either side of the mean. Furthermore, the targets are defined as,

$$o_1 (\text{human target}) : \{a_1^s \wedge (a_1^w \vee a_1^n)\} \quad (13)$$

$$o_2 (\text{vehicular target}) : \{a_2^s \wedge a_2^w \wedge a_2^n\} \quad (14)$$

Given these events and object definitions, we wish to determine the fused report, $D_f = \{P_f(o_1), P_f(o_2), P_f(o_1 \vee o_2)\}$, where, $\{o_1 \vee o_2\}$ represents the no target case.

C. Performance Analysis

Table I, and II show the classification performance of different techniques (averaged over 10 runs of the technique) when implemented on dataset 1 and 2 respectively.

TABLE I: Performance Comparison for the First Dataset

Method	Average Accuracy
Radar	86.47%
Telescopic Imaging	81.31%
Feature Concatenation	85.93%
Convex Quadratic Fusion (Similar Sensor Fusion)	86.07%
Analytic Center Fusion (Dissimilar Sensor Fusion)	88.61%
Dempster-Shafer Fusion	87.18%
Proposed Technique	90.36%

Classification accuracy is often not the best measure to quantify performance, particularly in cases where different classes have different numbers of samples, which is the case here. A better way to compare performance is to look at the Receiver Operating Characteristic (ROC) curves. Fig. 1 and 4 show the ROC curves for classification in case of datasets 1 and 2 respectively. It can be seen from the ROC curves




Fig. 4: ROC Curves for detection (second dataset) of (a): Class 1: Human Target, (b): Class 2: Vehicular Target, and (c): Class 3: No Target

TABLE II: Performance Comparison for the Second Dataset

Method	Average Accuracy
Seismic Sensor	85.41%
Acoustic Sensor	67.62%
Feature Concatenation	81.63%
Convex Quadratic Fusion (Similar Sensor Fusion)	86.69%
Analytic Center Fusion (Dissimilar Sensor Fusion)	89.96%
Dempster-Shafer Fusion	87.93%
Proposed Technique	92.04%

(for dataset 1) in Fig. 1 that other techniques show limited performance in correct classification of objects from class 2 due to the low number of samples for class 2 in comparison to those in class 1 and class 3. This causes the classifier to bias toward selecting class 1 or class 3, even when the sample is from class 2 in order to achieve high classification accuracy. But, our technique trains over occurrence of events rather than the object itself, hence does not face this issue. Improvement in performance is also seen for Dataset 2 (Fig. 4). In particular, detection of human targets is significantly improved, by taking 'or' between noise level event and weight event, which reduces misclassification due to noise of winds.

VI. CONCLUSION

We propose a novel sensor fusion technique at the decision level that looks at targets as combinations of events over the features that describe it, while also considering the extent of correlation between different features. Experiments on various datasets showed that the proposed technique can outperform existing fusion techniques on the decision level.

REFERENCES

- [1] D. L. Hall and J. Llinas, An introduction to multisensor data fusion, *Proceedings of the IEEE*, vol. 85, no. 1, pp. 623, Jan 1997.
- [2] Bahador Khaleghi, Alaa Khamis, Fakhreddine O. Karray, and Saiedeh N. Razavi, Multisensor data fusion: A review of the state-of-the-art, *Information Fusion*, vol. 14, no. 1, pp. 28–44, 2013.
- [3] Arthur P Dempster, A generalization of bayesian inference., *Classic works of the dempster-shafer theory of belief functions*, vol. 219, pp. 73104, 2008.
- [4] Glenn Shafer et al., *A mathematical theory of evidence*, vol. 1, Princeton university press Princeton, 1976.
- [5] Franz Rottensteiner, John Trinder, Simon Clode, Kurt Kubik, and Brian Lovell, Building detection by dempster-shafer fusion of lidar data and multispectral aerial imagery, in *Pattern Recognition, 2004. ICPR 2004. Proceedings of the 17th International Conference on. IEEE*, 2004, vol. 2, pp. 339342.
- [6] Otman Basir and Xiaohong Yuan, Engine fault diagnosis based on multi-sensor information fusion using dempstershafer evidence theory, *Information Fusion*, vol. 8, no. 4, pp. 379386, 2007.
- [7] Dennis M Buade and Paul Girardi, A target identification comparison of bayesian and dempster-shafer multisensor fusion, *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 27, no. 5, pp. 569577, 1997.
- [8] Yea S Huang and Ching Y Suen, The behaviorknowledge space method for combination of multiple classifiers, in *IEEE computer society conference on computer vision and pattern recognition. Institute of Electrical Engineers Inc (IEEE)*, 1993, pp. 347347.
- [9] Lingjie Li, *Data fusion and filtering for target tracking and identification*, Ph.D. thesis, 2003.
- [10] Lingjie Li, Zhi-Quan Luo, Kon Max Wong, and Eloi Bosse, Convex optimization approach to identify fusion for multisensor target tracking, *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 31, no. 3, pp. 172178, 2001.
- [11] Mihai Cristian Florea and Eloi Bosse, Critiques on some combination rules for probability theory based on optimization techniques, in *Information Fusion, 2007 10th International Conference on. IEEE*, 2007.
- [12] Thomas M Cover and Joy A Thomas, *Elements of information theory*, John Wiley & Sons, 2012.
- [13] Gbor J. Székely and Maria L. Rizzo, Brownian distance covariance, *Ann. Appl. Stat.*, vol. 3, no. 4, pp. 12361265, 12 2009.
- [14] Murat Kocaoglu, Alexandros G. Dimakis, Sriram Vishwanath, and Babak Hassibi, Entropic causal inference, *CoRR*, vol. abs/1611.04035, 2016.
- [15] Sylvester M Nabritt, Thyagaraju Damarla, and Gary Chatters, Personnel and vehicle data collection at aberdeen proving ground (apg) and its distribution for research, Tech. Rep., ARMY RESEARCH LAB ADELPHI MD SENSORS AND ELECTRON DEVICES DIRECTORATE, 2015.
- [16] Kyunghun Lee, Benjamin S Riggan, and Shuvra S Bhattacharyya, An accumulative fusion architecture for discriminating people and vehicles using acoustic and seismic signals, in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on IEEE*, 2017, pp. 29762980.
- [17] Siddharth Roheda, Benjamin S Riggan, Hamid Krim, and Liyi Dai, Cross-Modality Distillation: A case for Conditional Generative Adversarial Networks in Acoustics, Speech and Signal Processing (ICASSP), 2018 IEEE International Conference, *In Press*.
- [18] John Platt et al., Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods, *Advances in large margin classifiers*, vol. 10, no. 3, pp. 6174, 1999.
- [19] X Bian, B Gibson, and H Krim, A probabilistic approach of space objects detection from non-resolved optical observation, in *Advanced Maui Optical and Space Surveillance Technologies Conference*, 2013.
- [20] Ajit Singh and Peter Allen, Image-flow computation: An estimation-theoretic framework and a unified perspective, *CVGIP: Image understanding*, vol. 56, no. 2, pp. 152177, 1992.