Clutter Rejection by Clustering Likelihood-Based Similarities

Evan Hanusa
Dept. of Electrical Engineering
University of Washington
Seattle, WA, U.S.A.
Email: hanusaem@ee.washington.edu

David Krout
Applied Physics Lab
University of Washington
Seattle, WA, U.S.A.
Email: dkrout@apl.washington.edu

Maya R. Gupta
Dept. of Electrical Engineering
University of Washington
Seattle, WA, U.S.A.
Email: gupta@ee.washington.edu

Abstract— We implement and evaluate a likelihood-based method to cluster contacts in a multistatic active sonar setting. The underlying assumption is that a true contact will be detected by multiple receivers and any aspect-dependent feature must be consistent across all contacts in a cluster. Contacts which are contained in the same cluster can be appropriately fused and passed into a tracker. Clutter contacts detected are rejected if they are not in a cluster with any other possible objects. The use of the aspect dependent features Doppler and target strength allows for improved rejection of clutter. We show that clutter can be rejected with minimal false negatives.
Keywords: Clustering, Doppler, Multistatic Active Sonar, Clutter Rejection.

I. INTRODUCTION

In a multistatic active sonar array, a single transmitted ping reflects off of multiple objects and has the potential to be detected by any of the receivers. Each receiver estimates the location of the contacts which are within its detection range. Additional processing can result in the estimation of characteristics of the target (Doppler or amplitude, for example). If the receivers are not extremely far apart, it is reasonable to expect that if a target is present in the scenario area, it should be detected by multiple receivers. For the purposes of tracking and localization, it would be beneficial to combine each receiver's measurement of a single object's position. This is simple in a low-clutter, high probability-of-detection system. However, it becomes much more difficult in a multistatic active sonar environment, where probability of detection is low and there is a large amount of clutter.

In this paper, we propose a method for fusing contacts from multiple receivers by clustering the contacts, and we evaluate this approach on a standard multistatic tracking dataset. Figure 1 motivates the proposed clustering approach - one sees that contacts from the same object (marked by dots of the same color) cluster together, and thus clustering the contacts may help fuse contacts representing the same object.

We propose and implement clustering based on the similarity between contact position measurements and other aspect-dependent features of the contact. We investigate two clustering algorithms: a k-medoids clustering and a variant that incorporates assumptions consistent with the multistatic architecture. The different clustering algorithms are compared using several metrics motivated by the dataset. Section II
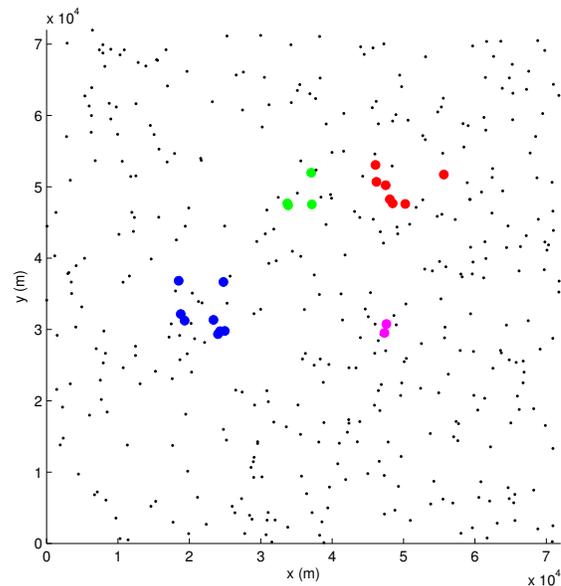
Fig. 1. Example contacts returned from one ping of a multistatic active sonar system. The small black points are the clutter contacts. The larger colored contacts originate from objects. Contacts which are colored the same are from the same object but were received by different receivers.

discusses related work. Section III introduces the similarity used in the clustering. Section IV describes the different clustering algorithms. Section V defines the metrics which will be used to evaluate the clustering performance and compares results across clustering algorithms. Section VI will discuss the results, suggest some methods for fusion within the clusters and pose some open questions.

Table I defines the key terms used in this paper.

II. RELATED WORK

The general topic of "tracking" spans several different domains, each with their own characteristics and challenges. Multistatic active sonar tracking tends to be characterized by large amounts of clutter and relatively low probability of detection. One approach to reject clutter in the hopes of improving later processing is to run two steps of multi-hypothesis tracking, with the first serving as a clutter rejection

κ	Contact
$N \in \mathcal{Z}^+$	Number of contacts
$c \in \mathcal{R}^2$	Contact position; $[x, y]$
$v \in \mathcal{R}^2$	Contact velocity; $[dx, dy]$
$r_i \in \mathcal{R}^2$	Receiver i location; $[x, y]$
$t \in \mathcal{R}^2$	Transmitter location; $[x, y]$
$\delta_j \in \mathcal{R}$	Doppler measurement for contact j
$d_j \in \mathcal{R}^+$	Range measurement for contact j
$b_j \in [0, 360)$	Bearing measurement for contact j
$z_j = [\delta_j, d_j, b_j]$	Measurements for contact j
$P(\cdot)$	Probability
$S(\cdot, \cdot) \in \mathcal{R}^+$	Similarity between two contacts
$k \in \mathcal{Z}^{++}$	Number of clusters
ψ	Cluster
m_j	Medoid of cluster j

TABLE I

TABLE OF KEY NOTATION.

step [1]. The work presented here differs from the *Maximum Likelihood - Multi-Hypothesis Tracker* (ML-MHT) approach in two key ways: this work uses a clustering approach, rather than the modified MHT, and also leverages Doppler, an aspect-dependent feature to improve the estimate of contact similarity.

Georgescu et al. propose a two step process for clutter rejection and contact fusion [2]. First, contacts are rejected using a ratio test which determines if a contact's features (SNR, Doppler) are distributed more similarly to a target or clutter. The remaining contacts are then combined in a pre-detection fusion step. The measurement error distribution is used to generate 100 particles for each contact. Then, all the particles for all the contacts are clustered by spatially binning them, where the spatial bins correspond to a regular grid of their $[x, y]$ coordinates. Any clusters which have a number of particles greater than a threshold are considered to be a contact, and these fused contacts are used in the tracker. The work presented here differs in a couple respects. First, the Doppler measurement is used in addition to the bearing and range measurements to create a pairwise similarity matrix between all contacts. This pairwise similarity matrix can then be used in a variety of different clustering algorithms. Additionally, the method presented here allows for the inclusion of any aspect-dependent feature in the similarity calculation, rather than only using the bearing and range measurements for clustering. Further, rather than the rough clustering produced by binning, we propose using a similarity-based clustering. We use a kernel k-medoids variant, but any clustering algorithm that operates on similarities could be used [3].

The use of clustering to combine region information is a common practice in video tracking. In general, clustering is used to combine multiple regions of interest that have been identified using segmentation or other techniques [4], [5]. These combined regions can then be tracked as one object, allowing for improved tracking performance. This is similar to the objective of the work presented here, however, the difference in the domain and type of measurement data necessitates a different approach.

Other approaches use the aspect-dependent target strength to only initiate tracks when a specular cue is observed, ensuring that clutter will not be used to initiate tracks [6]. This works well given that there is a large enough number of receivers that a specular cue can be observed, but it would not work as well in arrays with a smaller number of receivers or environments which result in a low probability of detection.

III. SIMILARITY BETWEEN CONTACT MEASUREMENTS

We desire a measure of similarity or distance between objects to determine what objects should be clustered together. In many situations, Euclidean distance is both the intuitive and correct choice, however this is not the case in multistatic sonar. To understand why this is the case and motivate a more appropriate choice, a brief recap of sonar measurements is included.

A. Multistatic Active Sonar Measurements

In a multistatic active sonar array, each transmitted waveform is received by many receivers. Each of these receivers processes the received waveform into a set of measurements. These measurements generally include a measurement of bearing b and bistatic range d for each contact. If the transmitted waveform is a *continuous wave* (CW) waveform, Doppler shift δ of the waveform due to the detected object's motion can be measured. It is important to note that these measurements are not comparable across receivers, because they are dependent on the location of the receiver. In addition, because the source and each receiver are located at different locations, there is a blanking region - a region in which no objects can be detected. Figure 2 shows the range and bearing measurements, as well as the blanking region.
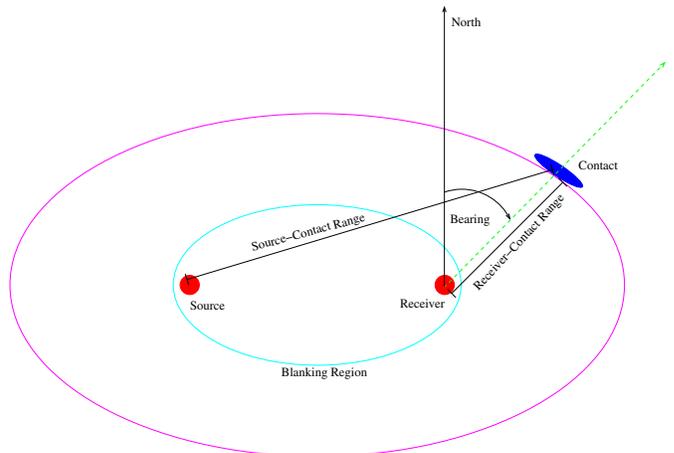
Fig. 2. A bistatic sonar system. The source and receivers are at different locations. The bistatic range measurement is the sum of the source-contact range and receiver-contact range. Bearing is measured clockwise from receiver north. The magenta equal-range line is an ellipse, with foci at the source and receiver. The blanking region is delineated by the smaller cyan ellipse with the same foci.

These measurements are also corrupted by noise due to algorithmic limitations or environmental effects. The distribution of the error is generally assumed to be known and additive

to the measurement. However, because the measurements are relative to the location of the receiver, it is non-trivial to create a simple mapping from the measurement error distribution on the bearing and range measurements of one receiver to bearing and range measurements of another receiver. Figure 3 is a plot of the posterior distribution $P(c|z)$ of an object being at a location c given measurements $z = [d, b]$ that have been corrupted by additive Gaussian noise in the measurement domains. Note that when mapped to the two-dimensional Cartesian coordinate space, the measurement error appears to be non-Gaussian. An extremely thorough discussion of multistatic sonar measurements is contained in [7].
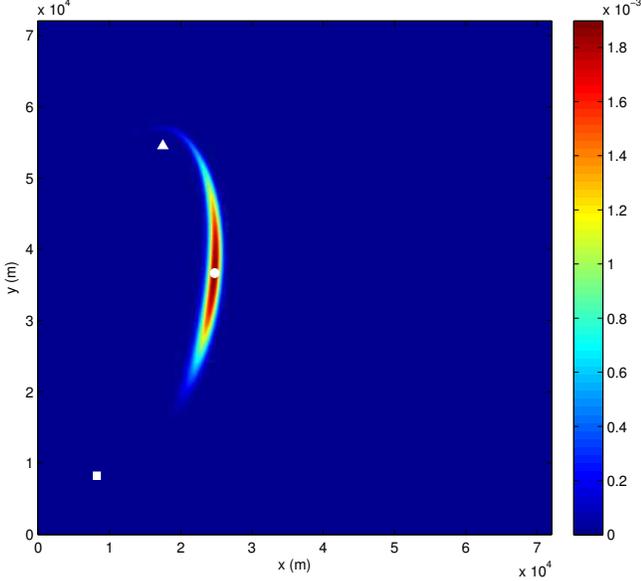
Fig. 3.   The posterior distribution over position given a range and bearing measurement with additive Gaussian noise. The true location is marked with a circle, the transmitter location is marked with a triangle, and the receiver location is marked with a square.

To overcome this non-linear measurement mapping between receivers, a likelihood-based similarity measure is proposed.

B. Likelihood Based Similarity

The similarity between two contacts is calculated using a likelihood-based approach [8]–[10]. Given the measurements of bearing, range and Doppler of two contacts, the similarity between contacts is calculated as follows:

First, given two measurements z_1, z_2 from two receivers r_1, r_2 the posterior distribution is

$$P(c, v|z_1, z_2, r_1, r_2, t),$$

where c is contact position, v is contact velocity and t is the transmitter location. By Bayes' rule,

$$P(c, v|\cdot) = \gamma^{-1} P(z_1, z_2|r_1, r_2, t, c, v) P(c, v|r_1, r_2, t),$$

where $P(z_1, z_2|r_1, r_2, t, c, v)$ is the *likelihood* of the measurements and γ is a normalizing constant,

$$\gamma = P(z_1, z_2|r_1, r_2, t).$$

Assuming the measurements from each receiver are independent,

$$P(c, v|\cdot) = \gamma^{-1} P(z_1|r_1, t, c, v) P(z_2|r_2, t, c, v) P(c, v|r_1, r_2, t).$$

In a bistatic sonar, there will be a blanking region for each receiver,

$$P(c, v|r_1, r_2, t) = 0 \; \forall c \; (D_1(c) < d_{blank} \vee D_2(c) < d_{blank})$$

where $D_i(c)$ is the mapping from location c to the bistatic range for receiver i. Assuming the measurements of bearing, range and Doppler are independent,

$$P(z_j|r_i, t, c, v) = P(\delta_j|r_i, t, c, v) P(b_j|r_i, t, c, v) P(d_j|r_i, t, c, v).$$

Assuming Gaussian measurement error,

$$P(\delta_j|r_i, t, c, v) = \mathcal{N}\left(\delta_j; \frac{v^T}{2}(\vec{u}_{tc} + \vec{u}_{r_i c}), \sigma_\delta^2\right),$$

where $\vec{u}_{tc}$ is the unit vector from the transmitter to location c and $\vec{u}_{r_i c}$ is the unit vector from receiver i to location c. This is derived from the bistatic Doppler equation [11]. For the bearing measurement,

$$P(b_j|r_i, t, c, v) = \mathcal{N}(b_j; B_i(c), \sigma_b^2),$$

where $B_i(c)$ is the mapping from location c to the bearing measured from receiver i. The range measurement is similar,

$$P(d_j|r_i, t, c, v) = \mathcal{N}(d_j; D_i(c), \sigma_d^2).$$

The similarity is defined here as the $\max$ over all possible contact positions and velocities,

$$S(z_1, z_2) = \max_{c,v} P(c, v|z_1, z_2, r_1, r_2, t). \tag{1}$$

In practice the posterior distribution is impossible to calculate analytically, and is approximated by sampling the position and velocity space. This is done for every pair of contacts that are not detected by the same receiver. Any two contacts that are detected by the same receiver are given a similarity of zero. This is done under the assumption that any object cannot generate more than one contact per receiver.

IV. CLUSTERING

Clustering is the process of assigning data points to groups such that points within each group are similar [12]. We propose clustering contacts, and in particular, propose using the likelihood-based similarity given in (1) to determine the clustering. This similarity is indefinite, that is, the $N \times N$ matrix of pairwise similarities between the N contacts is not necessarily a positive semidefinite matrix [12], and thus kernel clustering methods like kernel k-means are not applicable. Clustering methods that do act on indefinite similarities include linkage methods, spectral clustering, and k-medoids (see also [13] for a review of approaches to create kernels from indefinite similarities). Linkage methods can be very slow. For the experiments in this paper, we used a k-medoids algorithm as detailed in the next section.

A. K-Medoids Clustering

K-medoids is a k-means-like algorithm which uses only the similarity between points to cluster objects. K-means alternates between calculating the mean of the points assigned to each cluster, and then re-assigning points to the cluster with the nearest mean. K-medoids differs from k-means in that rather than calculating a mean of the points in the cluster, the point in the cluster to which all the other points are most similar (the *medoid*) is calculated.

There are several variants of k-medoids, the variant used here has the following steps [14]:

1) Randomly assign k contacts as medoids $\{m_1, m_2, \ldots, m_k\}$
2) Assign all non-medoid contacts to the nearest medoid
3) Calculate score $\mathcal{S}$ of this assignment (described below)
4) Swap medoid j with a non-medoid contact
5) Repeat Steps 2-4 for all non-medoid contacts
6) Repeat Steps 2-5 for all medoids $j = 1 \ldots k$
7) Select the set of medoids from steps 2-3 with the highest score $\mathcal{S}$
8) Repeat Steps 2-7 until convergence

The k-medoids algorithm gives a candidate clustering the following score:

$$\mathcal{S} = \sum_{j=1}^{k} \sum_{\kappa_\ell \in \psi_j} S(m_j, \kappa_\ell), \qquad (2)$$

where κ_ℓ is the ℓth non-medoid contact, ψ_j is the set of points in the jth cluster, and m_j is the medoid of the jth cluster.

B. Constrained K-Medoids

We also investigated a constrained k-medoids clustering. In this clustering variant, a constraint is added to the k-medoids algorithm that does not allow multiple contacts from the same receiver to be placed in the same cluster. This constraint incorporates the assumption that a single target cannot generate multiple contacts at a receiver. This modification is done by greedily assigning contacts to clusters: the un-clustered contact with the largest similarity to any medoid is assigned to that cluster if and only if the cluster does not already contain a contact from that receiver. If the cluster already contains a contact from that receiver, the similarity is treated as zero.

1) Randomly assign k contacts as medoids
2) Greedily assign contacts to medoids, restricted s.t. only one contact per receiver is in a cluster
3) Calculate score $\mathcal{S}$ of the assignment
4) Swap medoid j with a non-medoid contact
5) Repeat Steps 2-4 for all non-medoid contacts
6) Repeat Steps 2-5 for all medoids $j = 1 \ldots k$
7) Select medoids that result in highest score
8) Repeat Steps 2-7 until convergence.

In the next section, the metrics used to evaluate the performance of these clustering algorithms are described and the clustering algorithms are compared on a standard tracking dataset.

V. METRICS AND CLUSTERING PERFORMANCE

The ideal clustering of contacts would be to have one cluster which corresponds to each object present, and to have one or many clusters which contain only clutter contacts. This is a slightly different objective than most clustering problems due to the lack of a single true "clutter" cluster. To evaluate the clustering performance, two metrics are defined, *clusters per target* (CpT) and *object cluster impurity* (OCI). Clusters per target is defined here as the number of clusters which contain a true object. The optimal value for CpT is one. This metric becomes artificially small when clusters grow extremely large. For example, if all the contacts were clustered into a single cluster, CpT would be one but the clustering would not be very useful. To balance this, the "impurity" of the clusters which contain true objects is defined here. The impurity of a cluster is the percentage of false contacts in the cluster, and is ideally zero for any cluster that contains at least one true object. To measure the OCI of a clustering result, the average of the impurities for all clusters that contain at least one true contact is calculated.

A. Evaluation Dataset

The first scenario of the Metron dataset is used to evaluate the clustering performance of the different algorithms [15]. This simulated dataset has 4 transmitters (one active per ping) and 25 receivers aligned in a grid. This dataset is characterized by a low per-receiver probability of detection ($P_d \approx 0.13$ for each receiver) and large amounts of clutter ($N_C \approx 390$ per ping). In addition, bearing and range measurements are corrupted by a large error ($\sigma_b = 8°$, $\sigma_d = 600m$). Figure 1 is a plot of all the contacts from a single transmitted ping. An ideal clustering of these contacts would have one cluster per color, and each of these clusters would not contain more than one color nor any clutter (black dots). Clustering performance is evaluated for 30 pings from the first scenario which contain Doppler measurements.

B. Clustering Performance

An important step in any clustering problem is choosing the number of clusters. The number of clusters chosen can have a non-intuitive effect on the metrics chosen, so the two clustering methods described in Section IV are all compared for the same value of k. As a rough guideline, we suggest that

$$k \approx \frac{\bar{N}_C}{P_d N_{Rx}}, \qquad (3)$$

where $\bar{N}_C$ is the average number of contacts per ping, P_d is the per-receiver probability of detection, and N_{Rx} is the number of receivers. This results in the average cluster size being equal to the expected number of contacts that originate from an object. Initial results suggest that using a higher value for k can result in additional rejection of clutter, however the optimal value will likely be application-dependent.

A plot of the clustering that results from the constrained k-medoids algorithm (k=160) is shown in Figure 4. The clusters are displayed by connecting the medoid contact to all other

contacts in the cluster. This figure illustrates the result of using a non-Euclidean distance in the clustering: many contacts that are "close" in Euclidean distance are not clustered together, while other contacts that are quite far in Euclidean distance are in the same cluster.
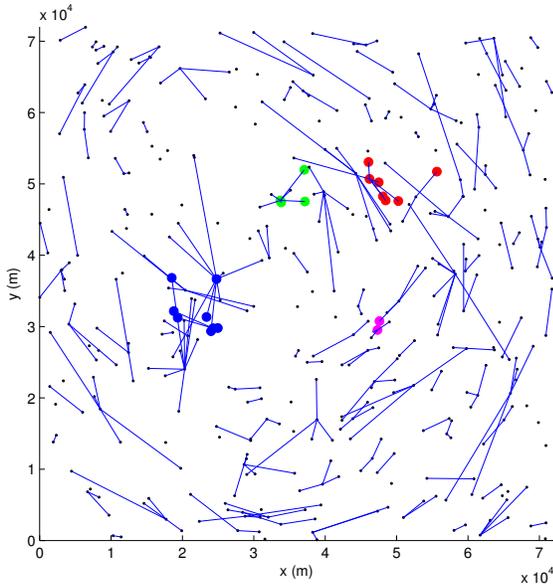
Fig. 4. Clustering results using the constrained k-medoids algorithm, with k=160. Blue lines indicate contacts that are clustered together.

The two clustering methods described in Section IV are compared in Table II. The metrics are calculated for each target on each ping, and the mean and variance of the CpT and OCI are calculated over all pings. These scores suggest that the standard k-medoids algorithm is a better choice for this dataset and choice of k. The metric of CpT is calculated for unclustered data to establish a baseline.

Method	Metric	Mean	Variance
Unconstrained	Object Cluster Impurity	0.555	0.057
Constrained	Object Cluster Impurity	0.519	0.913
Unconstrained	Clusters per Target	1.81	0.898
Constrained	Clusters per Target	1.99	1.07
Standard (Unclustered)	Clusters per Target	3.36	4.13

TABLE II

STATISTICS OF THE CLUSTERING PERFORMANCE METRICS FOR 30 DIFFERENT PINGS OF THE METRON DATASET, WITH K=160.

C. Clutter Rejection Using Cluster Size

A simple way of using the clustering results is to discard contacts that are in clusters that are smaller than a certain size. This is based on the assumption that a target will generate contacts on multiple receivers, and that these contacts will be clustered together due to a high similarity. Table III shows the rejection rates when clusters of sizes one and two are discarded. These rejection rates are calculated by dividing the

number of clutter or true contacts that would be rejected by the total number of clutter or true contacts in the pings. Using the unconstrained clustering and discarding clusters of size one results in rejection of 19.5% of clutter contacts and only rejecting 1.99% of target contacts.

Method	Rejected Cluster Size	Target Contact Rejection Rate	Clutter Rejection Rate
Unconstrained	1	0.0199	0.195
Constrained	1	0.0347	0.151
Unconstrained	2	0.0596	0.311
Constrained	2	0.137	0.347

TABLE III

TARGET AND CLUTTER REJECTION RATES WITH DIFFERENT CLUSTER SIZE THRESHOLDS.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we defined a likelihood-based similarity between contacts and compared the results of two clustering algorithms on a standard tracking dataset. The likelihood-based similarity uses position and Doppler measurements and is easily expandable to other aspect-dependent features.

The unmodified k-medoids algorithm performed the best, resulting in less than two clusters per target. The addition of the constraint that clusters may only contain one contact per receiver decreased the performance of clustering algorithm. This is due to two factors: the greedy nature of the algorithm used and the difficulty of the evaluation dataset. The constrained algorithm greedily assigns contacts to clusters, which results in sub-optimal clusterings. In tandem with the greediness of the algorithm, the large amount of clutter present in each ping makes it likely that a clutter contact will be more similar to a medoid than a target contact.

The scope of this work was only to cluster the contacts and compare two clustering algorithms. A potential research question is the how the clustered contacts should be used in a tracker. One option is to simply choose the medoids of the clusters and use them as inputs to a tracker instead of the full set of contacts. As described in Section V-C, clusters smaller than a threshold could be discarded under the assumption that a target should result in at least a certain number of contacts. Additionally, the measurements of the contacts in the cluster could be fused using a likelihood-based approach and then used in any contact-based tracker [9].

One open question is whether a different clustering method would offer a performance or computational advantage. We would also like to further investigate how the chosen number of clusters k affects the performance, and whether there is a better way to set k than what we have proposed here.

REFERENCES

[1] S. Coraluppi and C. Carthel, "An ML-MHT Approach to Tracking Dim Targets in Large Sensor Networks," in *Information Fusion, 2010. 13th International Conference on*, July 2010.

[2] R. Georgescu, P. Willett, and S. Schoenecker, "GM-CPHD and ML-PDA Applied to the Metron Multi-Static Sonar Dataset," in *Information Fusion, 2010. 13th International Conference on*, July 2010.

[3] R. Arora, M. R. Gupta, A. Kapila, and M. Fazel, "Clustering by Left-Stochastic Matrix Factorization," in *Proc. International Conference on Machine Learning (ICML)*, 2011.

[4] B. Gunsel, A. Ferman, and A. Tekalp, "Temporal video segmentation using unsupervised clustering and semantic object tracking," *Journal of Electronic Imaging*, vol. 7, no. 3, pp. 592–604, 1998.

[5] A. Lipton, H. Fujiyoshi, and R. Patil, "Moving target classification and tracking from real-time video," in *Applications of Computer Vision, 1998. WACV'98. Proceedings., Fourth IEEE Workshop on*. IEEE, 2002, pp. 8–14.

[6] D. Grimmett, "Reduction of false alarm rate in distributed multistatic sonar systems through detection cueing," in *OCEANS 2007 - Europe*, 2007, pp. 1 –6.

[7] S. Coraluppi, "Multistatic sonar localization," *Oceanic Engineering, IEEE Journal of*, vol. 31, no. 4, pp. 964 –974, 2006.

[8] L. Stone, T. Corwin, and C. Barlow, *Bayesian multiple target tracking*. Artech House, Inc. Norwood, MA, USA, 1999.

[9] E. Hanusa, D. Krout, and M. R. Gupta, "Estimation of Position from Multistatic Doppler Measurements," in *Information Fusion, 2010. 13th International Conference on*, July 2010.

[10] D. Krout and E. Hanusa, "Likelihood Surface Preprocessing with the JPDA algorithm: Metron Data Set," in *Information Fusion, 2010. 13th International Conference on*, July 2010.

[11] C. Clay and H. Medwin, *Acoustical oceanography: Principles and applications*. Wiley, 1977.

[12] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction*. Springer Verlag, 2009.

[13] Y. Chen, E. K. Garcia, M. R. Gupta, A. Rahimi, and L. Cazzanti, "Similarity-based Classification Concepts and Algorithms," pp. 747–776, 2009.

[14] L. Kaufman and P. Rousseeuw, "Clustering by means of medoids," Technische Hogeschool, Delft(Netherlands). Dept. of Mathematics and Informatics., Tech. Rep., 1987.

[15] K. Orlov, "Description of the Metron Simulation data set for the MSTWG," *Released to the MSTWG*, June 2009.