

APPROACHING SEMANTICALLY-MEDIATED ACOUSTIC DATA FUSION

Baofeng Guo, Yi Wang, Paul Smart, Nigel Shadbolt and Mark S. Nixon

Department of Electronics and Computer Science,
University of Southampton, UK

T. Raju Damarla

U.S. Army Research Laboratory, 2800 Powdermill Road,
Adelphi, MD, USA

ABSTRACT

Our primary hypothesis is that it should be possible to enrich data fusion by semantic processing, with wide potential application. In order to achieve our aim we need to represent the semantic data and enable reasoning about it in a framework that can be aligned with data fusion. Ontologies are most suited to this task as they allow for rich representation of data structure; some approaches include probabilistic representation. These can be aligned with data fusion approaches, such as Bayesian, which can fuse by including estimates of uncertainty. We shall describe our initial approaches towards establishing our hypothesis, including a survey of the enabling technologies, a description of application data (acoustic sensors, military scenario), and our new method of feature selection for acoustic data fusion. We also explore the semantic attributes and the representations that can be deployed for enrichment purposes, showing how ontologies can be used in this context. In these respects we shall show how we can approach enrichment of data fusion by semantic technologies, how this can capitalise on the current stock of techniques, and illustrate the potential benefits associated with this new approach.

I. INTRODUCTION

The need for data enrichment is manifest in the plethora of approaches developed for data fusion. Essentially, data fusion approaches aim to fuse descriptions that capture different aspects of an artifact, so as to improve classification capability. Our motivation here concerns enriching the data, rather than the process; we seek to augment data for classification purposes, rather than to enhance decision making processes. Given that this approach is being developed within the Information Technology Alliance (ITA) our primary concern is military data. Here the military data is from acoustic sensors and this is used to classify vehicle type from data supplied by autonomous sensors.

These data are acoustic signals from multiple sensor sources. For fusion purposes, these signatures can be data

and represented in different ways by transformation of the source this essentially derives different invariance attributes in the feature description. However, to classify vehicles from such data still remains a difficult problem, such as lower SNR, and complex ambient interferences. The data can be enriched by semantic means, which concerns intelligence and human derived descriptions. In this way, the features can be better separated thereby improving classification capability. The mechanism we are using to effect this enrichment is by using ontologies, as they are suited to representation and reasoning within semantic data. There have been prior approaches to ontological enrichment of data fusion processes, but these have not been phrased in terms of enriching classification capability.

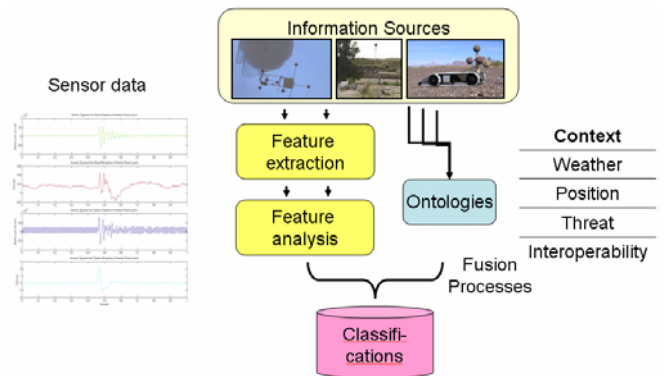


Fig. 1 Semantically-mediated data fusion

The framework of our approach is shown in Figure 1. Here, we seek to classify the vehicle that generated the sound recorded by microphones. In a conventional pattern recognition framework, features are extracted from this sensor data and these features are filtered according to perceived information content, prior to use in classification. We seek to enrich this process by semantic data, which we shall represent using ontologies. These will contribute to the data fusion processes which lead to the combined and enriched decision.

In the next Section we describe the current stock of approaches to (conventional) data fusion, including the JDL data fusion model for completeness. We then describe, in Section III, how ontologies can be used to

describe and reason within semantic data. Section IV concerns feature subset selection, which we use to explore the data space used within conventional data fusion. We describe our exemplar application data in Section V, together with the preliminary results we have obtained on it in Section VI. These show that data fusion and feature subset selection can be used to classify the data with success. Further, we have defined semantic attributes that can be associated with the data to explore the possibilities of semantic enrichment.

II. DATA FUSION

A. JDL Data Fusion Model

The JDL data fusion model is the most widely adopted functional model for data fusion. It was developed in 1985 by the U.S. Joint Directors of Laboratories (JDL) Data Fusion Group [1] with several recent revisions proposed [2, 3, 4]. The JDL distinction among fusion ‘levels’ (they may be heavily inter-related) provides an often useful distinction among data fusion processes that relate to the refinement of ‘object’, ‘situations’, ‘threats’ and ‘processes’ [1]:

- Level 0 - signal/feature assessment: semantic annotation and processing of high-dimensional data sets; guidance of information acquisition processes with respect of epistemic constraints; facilitation of feature extraction process using background knowledge, etc.
- Level 1 - Object assessment: improved certainty estimates with respect to fused data; facilitation of entity characterization in terms of implied features; support for identity inference, etc.
- Level 2 - Situation assessment: knowledge-filtered awareness, including support for contextual relevance reasoning, information triage, and representation of situation state.
- Level 3 - Impact assessment: support for rules-based processing of situation-relevant data in relation to decision support processes, e.g., automated mission planning, threat assessment, battlefield planning, etc.
- Level 4 - Performance assessment: identification of knowledge gaps/epistemic inadequacies and the implementation of appropriate remedial actions; provision of explanatory support to enable system evaluation and validation of knowledge system operation.

- Level 5 - User Refinement: adaptive determination of who queries information and who has access to information and adaptive data retrieval and display to support cognitive decision making and actions given social and political contexts.

Note that the levels were differentiated first on the basis of types of estimation process, which typically relates to the type of entity for which state is estimated. In general, the benefit of this scheme of partitioning fusion functions into these levels is due to the significant differences in the types of input data, models, outputs, and interferences applicable to problems at different levels.

B. Information flow across levels

Processing at each of these Data Fusion levels involves batching the available data for fusion into a network of fusion nodes where paradigmatically each fusion node accomplishes

- Data preparation (data mediation, common formatting, spatiotemporal alignment, and confidence normalization);
- Data Association (generation, evaluation and selection of association hypotheses; i.e. of hypotheses; i.e. of hypotheses as to the applicability of specific data to particular aspects of the state estimation problems);
- State Estimation and Prediction (estimating the presence, attributes, inter-relationships, utility and performance or effectiveness of entities of interest, as appropriate to the data fusion node).

In all the fusion levels, the accuracy of the fused state estimates tend to increase as large batches of data are fused; however the cost and complexity of the fusion process also increases. Thus, a knee-of-the-curve of performance versus cost fusion node network is sought in the system design and operation. As noted above, the data fusion levels are not necessarily processed in order and any one can be processed on its own or in combination given the corresponding inputs and there is feedback across Levels. The notion of inter-Level ‘informing’, controlling, and exploitation can in fact become quite complex in certain applications, and has similarities to the complexities of peer-to-peer internetworking processes at multiple levels of abstraction. In the course of one Level informing another, there should be some sense of added value, or utility, balancing the negative aspects of the additional processing complexity and time delay of enable such feedback. Moreover, the possibility of such feedback raises concerns for maintaining consistency in inference across levels.

C. Data Fusion Algorithms

In the typical military setting, there are various platforms with sensors of different types. From such information, the data fusion system needs to produce reasonable hypotheses of the actual truth. Various mathematical techniques [5] have been developed to deal with this problem, which include:

- Bayesian methods, in which the degrees of belief are represented by a prior, conditional, and a posteriori probability. The typical techniques in this category are the Kalman filter, Multi-Hypothesis Estimation (MHE) filter, and the Joint Probabilistic Data Association (JPDA) filter.
- Evidential methods encompass several models such as the Dempster-Shafer theory [6] and transferable belief theory [9]. In the framework of the Dempster-Shafer theory information obtained from a source is represented by the Basic Probability Assignment (BPA). Fusion of independent and equally reliable basic probability assignments is performed by the Dempster rule of combination.
- Rough sets whose theory [14] deals with imprecision. The basic concept of the rough sets theory is to replace uncertain or imprecise information by two imprecise but certain information: the lower and upper approximations. The combination of imprecise information is realized by applying the set theory to the approximations.
- Possibility and fuzzy methods. In the framework of possibility theory, information obtained from a sensor is represented by possibility distribution. The combination rules can be based on t-norms and t-co-norms [5], the fuzzy translation of the intersection and union, etc.

As such, there is a selection of extant approaches to data fusion. We need means to represent the semantic data, and to enable reasoning in this space. Ontologies appear most suited to this task, with the further advantage that they can also be used to represent trust and uncertainty.

III. ONTOLOGIES

Ontologies are regarded as the basic building units and integral parts for the semantic representation, as they provide a reusable piece of knowledge about a specific domain. The use of ontologies for the explication of implicit and hidden knowledge is a possible approach to overcome the problem of semantic heterogeneity.

Generally speaking, the semantic heterogeneity deals with three types of concepts [7]: the semantically equivalent concepts, the semantically unrelated concepts and the semantically related concepts. In the first case, a model uses different terms to refer the same concept; in the

second case, the same term may be used by different systems to denote completely different concepts; and in the last case, different classifications may be performed.

As the ontology development process becomes more ubiquitous and collaborative, the difficulties in making ontologies inter-operable become a serious problem. In order to achieve effective semantic inter-operability in a heterogeneous information system, the meaning of the information that is interchanged has to be understood across the systems. The domains covered by ontologies have to be few, thus avoiding conflicts between useage of the context.

Semantic conflicts occur whenever two contexts do not use the same interpretation of the information. Goh identifies three main causes for semantic heterogeneity [8]:

- Confounding conflicts occur when information items seem to have the same meaning, but differ in reality, e.g. due to different temporal context.
- Scaling conflicts occur when different reference systems are used to measure a value.
- Naming conflicts occur when naming schemes of information differ significantly.

In nearly all ontology-based integration approaches ontologies used for the explicit description of the information source semantics. A question that arises from the use of ontologies for different purposes in the context of information integration is about the nature of the used ontologies.

The task of integrating heterogeneous information sources put ontologies in context. They cannot be perceived as standalone models of the world. They should rather be seen as the glue that puts together information of various kinds. Consequently, the relation of ontology to its environment plays an essential role in information integration. By the term mappings is understood the connection of an ontology to other parts if the application system. The two most important uses of mappings required for information integration are mappings between ontologies and the information they describe and mapping between different ontologies used in a system. There are further aspects to the use of ontologies, including the representations of trust and uncertainty, but we shall not explore these here.

We have yet to find any techniques developed to understand the feature content of the semantic data, known as feature subset classification for more conventional measured data. As we can apply feature subset selection to the data to be classified, we ought also to be able to understand the semantic space better by applying similar techniques there. As such, we shall move to techniques

which can be used to explore the data space in the next Section.

IV. FEATURE SUBSET SELECTION

A. Traditional Approaches

Feature selection is a processing to choose a subset of features or some combination of the input features that best represent the original data under a certain criterion. Data fusion can benefit from feature selection by discarding some of the redundant information and reducing dimensionality. Roughly speaking, the feature selection methods may be based on two approaches: (1) a pre-processing procedure, independently of the classification algorithm; and (2) a performance indicators, directly connected with the classification algorithms. Other dimensionality reduction techniques include principal component analysis, Fisher Discriminant Analysis, etc.

B. Information-based methods

Compared to the traditional methods, information-based methods directly measure the information content of each individual feature. If the measured information content is related to the level of discriminatory capability, the feature selection can be carried out by choosing those features with the higher information content.

According to Shannon's information theory, entropy measures information content in terms of uncertainty, and is defined by

$$H(A) = -\sum_{A \in a} p(A) \log p(A) \quad (1)$$

where A is a random variable taking values in the set a with probability distribution $p(A)$.

One may directly use the entropy as a criterion for feature selection. However, by examining the definition of entropy, it can be seen that entropy is calculated with respect to the single variable A , without reference to any objective. Thus, the amount of information measured by the entropy lacks a point of reference or benchmark. To improve the entropy-based methods, it is logical to extend the information measure to two variables: one for the measured feature itself and the other for the class label that is directly related to the classification objective. Mutual Information (MI) provides a framework to measure the similarity between two random variables, and was introduced recently for feature selection [15].

Given two random variables A and B , with marginal probability distributions $p(A)$ and $p(B)$, and joint probability distribution $p(A, B)$, MI is defined as:

$$I(A, B) = \sum_{A \in a, B \in b} p(A, B) \log \frac{p(A, B)}{P(A) \cdot P(B)}. \quad (2)$$

From (2), it is not difficult to find that MI is related to entropy by the following equations:

$$I(A, B) = H(A) + H(B) - H(A, B) \quad (3)$$

where $H(A)$ and $H(B)$ are the entropies, and $H(A, B)$ the joint entropy.

If we model each individual feature and the corresponding class label as random variables, MI can be used to estimate the dependency between the feature component and its class category. This can be used to investigate how much information a feature component contains about the class label, and use it for feature selection (see Fig. 2). In this way, we have approaches which can fuse data, and a new approach to explore the feature space to determine potency for classification capability.




Fig. 2 Illustration of mutual information

V. ACOUSTIC DATA

A. Data Description

We are testing our approaches based on an acoustic data set provided by the U.S. Army Research Laboratory (ARL). This data is most suited to our approach since it is an active and challenging problem, and one which is suited to semantic analysis. Essentially, the data reflects an interest in automatic sentinels that can automatically determine the presence of offensive vehicles. ARL has conducted several experiments in tracking a convoy of multiple targets using several spatially distributed sensor arrays. Fig. 3(a) shows an example of the configuration of a sensor array with six microphones placed at the vertices of a hexagon with a radius of 4 ft.

The data from each sensor array is collected at a sampling rate of 1024 samples per second. The raw data from each microphone in a sensor array is recorded for classification. After FFT, each second of acoustic signal is transformed into 351 dimensional spectral data. Fig. 3(b) shows an example of the location of six sensor arrays and a run test

around a prearranged track.




Fig. 3 (a) Sensor array configuration and (b) Test track and the positions of sensor arrays [16]

B. Semantic Representation

When considering how data is semantically represented, four fundamental questions to ask are: (1) How is the representation related to conceptual structures? (2) How is the meaning of each data represented? (3) How are the meanings of different data related to one another? With few exceptions, existing theories of semantic have made explicit claims concerning the representation of each meaning and the relations among different word meanings, while the relation between conceptual and semantic structures is often left implicit, and the issue of whether different principles are needed for representation of different content domains is often neglected.

In this paper we use a method called Focused ontology integration to represent the semantics within the data. This method consisting of following steps:

1. Identify a limited number (3-5) of (existing) ontologies close related to the problem
2. Find the places in the ontologies where they overlap;
3. Relate concepts that are semantically close via equivalence and subsumption relations (aligning);
4. Check the consistency, coherency and non-redundancy of the result.

Then we use this as the base to create more ontologies.

VI. ANALYSIS OF ACOUSTIC DATA

A. Semantic Analysis

The semantic enrichment of the data is done by identifying semantic concept and relations appearing in the data and from the data fusion. Of course this has a large coverage of different ontological domain, based on the discussion in section 3, we decided to focus on four type of ontologies

- a sensor ontology – semantic description of the sensors;

- a sequence ontology – semantic description of event where sensor was tested;
- a data ontology – semantic description of data received by the sensor; and
- a supporting ontology– semantic description of concepts that would effect all three mentioned ontology.

These allow us to have a much improved description than many of the previous research where semantic representation and data fusion were used individually.

Results from data fusion are used to represent the entities and properties contribute to the relationships between them. For example, some important relationships include “has a type of feature”, “is a kind of amplitude”, “has a function as”, etc.

Sensor ontologies support concepts such as “If A receives C from B”. In addition to the general properties used in any ontology, certain actions apply especially to sensors. The action, “detect,” is quite basic to sensors. Other verbs are “identify” and “classify”. Semantic relationships between actions play an important and often key role in sensor-related concepts.

Sequence ontology contains spatial and temporal information. For example, A may happen before B. X could be detected. In this case X enables B because B cannot occur without the existence X. It also implies that X must occur BEFORE A.

The data ontologies are the most complete description of the data fusion result. It totally depends on the algorithms selected and how wide the semantic description will cover.

The supporting ontology contains information and relations that is not clearly expressed in the last three ontologies. We focus mainly on the environmental and vehicle related classes and properties, such as weather, weight and size.

B. Feature level fusion for acoustic data

In our new assessment of MI-based feature selection for acoustic data, we propose a feature-level fusion method by combining two set of acoustic features. In details, harmonic features are used to characterize the fundamental frequency, and a group of key frequency components, selected by MI, are used to reflect other useful acoustic factors. Fusing these two set of features may provide a more complete description about the vehicles’ acoustic signatures, and improve the classification accuracy accordingly. To keep the same dimensionality of feature space, the fusion is devised by replacing the higher order harmonics components with the same number of key frequency components.

The harmonics are effective acoustic features. However, the acoustic model for vehicles may be more complicated, particularly in the area of non-harmonic features. Thus, an acoustic signature S regarding a working vehicle can be modelled as a combination of harmonic and other non-harmonic components:

$$S = H(f_0) + K(f_0) \quad (4)$$

In (4), the first term $H(f_0) = \sum_i h_i, i = 1, 2, \dots, M$ denotes a set of harmonic features. The second term $K(f_0) = \sum_j k(f_j), j = 1, 2, \dots, N, f_j \neq f_0$ is a group of features from the non-harmonic part. In this paper, we use $f_j, j = 1, 2, \dots, N$ to indicate the key frequency components that are not related with the fundamental frequency f_0 , but are considered also containing useful classification information.

Based on the vehicle-signatures' model in (4), the fusion can be implemented by combining information from both harmonics and other key frequency components. In details, we adopted a feature level fusion, where features from two sources are concatenated together to form a new feature vector, such as:

$$F = \{h_1, h_2, \dots, h_K, k(f_1), k(f_2), \dots, k(f_L)\} \quad (5)$$

Here we have $K + L = M$, and M is the dimensionality of the pre-specified harmonics feature space. Other fusion schemes, like those introduced in Section II, are also applicable and will be discussed in the future. The fused feature vectors now have the same dimensionality as the harmonic features', but with the L higher order harmonics replaced by the same number of key frequency components.

Experiments are carried out based on a multi-category vehicles data set from ARL, USA [10, 16-17]. In the experiments, half of data samples from each class were randomly chosen to estimate the statistical parameters, such as harmonic features' means vector, covariance matrix, and also mutual information. The remaining 50% samples forms the testing set on which performance was assessed.

Next, feature selection is carried out based on the methods introduced in Section IV. Following the results in [16], the harmonic number is chosen as 21. The dimensionality of the feature vectors extracted by mutual information is flexible, and the detailed number can be decided by the computing resources.

Currently-popular support vector machines (SVMs) [18], were chosen as the classifiers in these experiments because

they are less sensitive to the higher dimensional data. Although SVMs are used here, other classification algorithms, such as the Multivariate Gaussian classifier [16], are also applicable. The kernel function used is an inhomogeneous polynomial. The penalty parameter C is tested between 10^{-3} and 10^5 and polynomial order is tested between 1 and 10 by a two fold validation procedure using only training data. A polynomial of order 3 and $C = 20$ were finally found as the best values for this SVM.

To avoid bias on random samplings, the testing was repeated 10 times to allow an estimate of the error inherent in this sampling process. The classification results based on different feature sets are shown in Table I. The first two rows of Table I list the classification accuracy before fusion, where the 21 dimensional harmonic features and key frequency features are applied individually. It is seen that the feature set represented by the key frequency components is relatively weak.

The next two rows show the accuracy results after fusion. In the third row, to keep the dimensionality as 21, 10

TABLE I
COMPARISON OF CLASSIFICATION RESULTS BASED ON
DIFFERENT FEATURE SETS

Feature set	Mean accuracy (%)	Standard deviation
Harmonics(21dim)	74.4	1.45
Key Freqs(21dim)	74.3	1.05
Harmonics(11dim)+ Key Freqs(10dim)	78.5	0.92
Harmonics(11dim)+ Key Freqs(21dim)	81.1	0.57

higher order harmonics are removed. The vacated space is then replaced with the 10 features based on the key frequency components. It can be seen that after this fusion, the classification accuracy is improved from 74.4% to 78.5%. In the fourth row, if we add more key frequency components (e.g., 21 dimensional key frequency components) to the 11-dimensional harmonic features, the classification accuracy can be further increased to 81.1%, but at the cost of higher dimensionality (i.e., 32 dimensional feature space).

The accuracy values shown in Table II are overall numbers, which can be seen as weighted results based on all five individual classification accuracies. In this multi-category vehicle data set, the number of testing set for each vehicle is different due to their different running speeds. A bias may occur if a method has preference to the classes which have more numbers than others. Therefore, we further list the classification accuracy for each individual vehicle, labelled by V1 to V5 in Table II

The same 10 times random samplings are carried out, and the (mean classification accuracy) \pm (standard deviations) for each of the five vehicles are listed. From Table II, it is

seen that for the case of the same dimensionality (21 features), the fusion method is consistently better than other un-fused methods, in all five classes.

Though Table I and II confirmed the improvement of classification accuracy by using the proposed feature-level

TABLE III
COMPARISON OF CLASSIFICATION RESULTS (MEAN (%) + STD)
FOR EACH INDIVIDUAL TYPE OF VEHICLES; FUSION METHOD 1:
HARMONIC (11DIM)+ KEY FREQS(10DIM); FUSION METHOD 2:
HARMONIC(11DIM)+ KEY FREQS(21DIM)

Vehicle types	HARMONICS (21DIM)	KEY FREQS (21DIM)	FUSION 1 (21DIM)	Fusion 2 (32dim)
V1	84.2±1.70	70.6±4.80	86.2±1.28	87.1±1.01
V2	78.7±1.06	83.4±0.87	82.8±0.80	85.7±0.88
V3	83.2±1.36	81.1±1.27	83.8±0.88	84.9±0.53
V4	50.7±5.69	54.2±3.70	59.4±3.08	65.2±1.69
V5	65.4±1.92	66.9±2.06	73.5±2.33	76.4±1.50

fusion, it is also realized that the change of the vehicles' velocity may affect the selected result by mutual information. In this sense, the features represented by the key frequency components are not as stable as the harmonics, and a further research is undergoing to address this problem.

VII. CONCLUSIONS AND FUTURE WORK

The case for enrichment of data fusion processes is quite compelling, especially in military scenarios, and we look forward to achieving this. Our target is to enrich classification capability by augmenting data fusion processes using semantic data. Our example application concerns the identification of vehicles from autonomously-sensed acoustic data. We have demonstrated that the component technologies needed to enable our overall aims are already sufficiently well developed for this task. In this, data fusion can improve classification capability from the acoustic data. Further, our new approach to feature set selection by information content can improve potency in the descriptions used, as here shown on the ARL data. We have described how we will deploy the semantic enrichment by using ontologies together with some of the semantic attributes we intend to explore. In this way we will enrich the results derived by data fusion and by feature subset selection and we look forward to improving the classification capability still further, by this new approach.

ACKNOWLEDGEMENTS

We gratefully acknowledge support by ARL/ MoD through IBM on the Information Technology Alliance research programme.

REFERENCES

- [1] White, F.E., "A Model for Data Fusion", Proc. 1st National Symposium on Sensor Fusion, 1988
- [2] A. N. Steinberg and C. L. Bowman. chapter 2. Handbook of Multisensor Data Fusion, CRC Press June 20, 2001, ISBN: 0849323797.
- [3] E.P Blasch., and Plano, S., "Level 5: User Refinement to aid the Fusion Process", in Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications 2003, B. Dasarathy (Ed.), *Proc of the SPIE*, **5099**, 2003
- [4] J. Llinas, C. Bowman, G. Rogova, A. Steinberg, E.Waltz, and F. White. Revisions and extensions to the JDL data fusion model II. In P. Svensson and J. Schubert, editors, *Proc. of the 7th Int. Conf. on Information Fusion*, 2004.
- [5] G. V. Veres, Trung Dong Huynh, Mark S. Nixon, Paul S. Smart, and Nigel R. Shadbolt, "Military knowledge information fusion via Semantic Web technologies", *ITA Report*, 2006
- [6] G. Shafer, "A Mathematical Theory of Evidence". *Princeton*, 1976.
- [7] Z. Cui and P. O'Brien, "Domain ontology management environment," in *Proc. of the 33rd Hawaii Intl Conf on System Sciences*, 2000.
- [8] C. Goh, "Representing and reasoning about semantic conflicts in heterogeneous information sources," Ph.D. dissertation, MIT, 1997.
- [9] P. Smets and R. Kennes, "The transferable belief model," *Artificial Intelligence*, **66**, pp. 191–243, 1994.
- [10] T. R. Damarla, T. Pham, and D. Lake, "An algorithm for classifying multiple targets using acoustic signatures", *Proc. of SPIE Signal Processing, Sensor Fusion and Target Recognition*, pp. 421-427, 2004.
- [11] B. V. Veen, "Minimum Variance Beamforming", *Adaptive Radar Detection and Estimation*, Edited by S. Hakin and A. Steinhardt, John Wiley & Sons, Inc, 1992.
- [12] T. Pham and B. Sadler, "Adaptive wideband aero acoustic wideband array," *8th IEEE SP Workshop on Statistical Signal and Array Processing*, pp. 295-298, June 1996.
- [13] T. Pham and B. Sadler, "Wideband acoustic array processing to detect and track ground vehicles," *1st Annual ARL Sensors and Electron Devices Symposium*, College Park MD, January 1997.
- [14] Z. Pawlak, "Rough Sets - Theoretical Aspects of Reasoning about Data", Klumer Academic Publisher, 1991.
- [15] B. Guo, S.R. Gunn, R.I. Damper, and J.D.B. Nelson. "Band selection for hyperspectral image classification using mutual information". *IEEE Geoscience and Remote Sensing Letters*, **4(3)**:522– 526, 2007.
- [16] T. Raju Damarla and Gene Whipps. "Multiple target tracking and classification improvement using data fusion at node level using acoustic signals". Technical Report, ARL.
- [17] D. Lake. "Harmonic phase coupling for battlefield acoustic target identification". *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 2049–2052, 1998.
- [18] B. E. Boser, I. M. Guyon, and V. N. Vapnik. "A training algorithm for optimal margin classifiers". In *Proceedings of the 5th Annual Workshop on Computational Learning Theory*, pages 144–152, Pittsburgh, Pennsylvania, United States, 1992.