CONF-971068--1

INDUCTIVE LEARNING AS A FUSION ENGINE FOR MINE DETECTION*

Stephen W. Kercel and William B. Dress
Oak Ridge National Laboratory
P.O. Box 2008
Oak Ridge, Tennessee 37831-6011
(423) 574-5278

RECEIVED

JUL 28 1997

OSTI

To be presented at the
1997 IEEE International Conference on Systems, Man, and Cybernetics
Orlando, Florida
October 12-15, 1997

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MASTER

*This research was performed at OAK RIDGE NATIONAL LABORATORY, managed by LOCKHEED MARTIN ENERGY RESEARCH CORP. for the U.S. DEPARTMENT OF ENERGY under contract DE-AC05-96OR22464.

DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

INDUCTIVE LEARNING AS A FUSION ENGINE FOR MINE DETECTION*

Stephen W. Kercel and William B. Dress
Instrumentation and Controls Division
Oak Ridge National Laboratory
PO Box 2008 MS 6011
Oak Ridge Tennessee 37831-6011, USA

ABSTRACT

Semiotics is defined by some researchers as "the study of the appearance (visual or otherwise) meaning, and use of symbols and symbol systems." Semiotic fusion of data from multiple sensory sources is a potential solution to the problem of landmine detection. This turns out to be significant, because notwithstanding the diversity of sensor technologies being used to attack the problem, there is no single effective landmine sensor technology.

The only practical, general-purpose mine detector presently available is the trained dog. Most research into mine-detection technology seeks to emulate the dog's seemingly uncanny abilities. An ideal data-fusion system would mimic animal reaction, with the brain's perceptive power melding multiple sensory cues into an awareness of the size and location of a mine. Furthermore, the fusion process should be adaptive, with the skill at combining cues into awareness improving with experience.

Electronic data-fusion systems reported in the countermining literature use conventional vector-based pattern recognition methods. Although neural nets are popular, they have never satisfactorily met the challenge. Despite years of investigation, nobody has ever found a vector space representation that reliably characterizes mine identity. This strongly suggests that the features have not been found because researchers have been looking for the wrong characteristics. It is worth considering that dogs probably do not represent data as mathematical number lists, but they almost certainly represent data via semiotic structures.

The most relevant, practical example of semiotic pattern recognition is Goldfarb's concept of the Inductive Learning Paradigm. He implements the concept with a mathematical structure called the Evolving Transformation System (ETS). The scheme is a semiotic-based, unifying theory of pattern recognition in which classical structural methods, vector-based systems, and even fuzzy logic, all fall out as special cases. In the past, structural methods

have had limited practical value, and vector-based methods attempt to estimate "decision surfaces" between classes in "feature vector space," but fail to reveal the structure of the data. In this paper, the authors explore the feasibility of applying Inductive Learning to the problem of mine sensor-fusion.

INTRODUCTION

Despite the profusion of sensor technologies being applied to the problem, the fact remains that countermining sensor technology has simply not progressed to a satisfactory level. The essence of the present state of countermining sensing is captured in a quote from Colonel Robert Greenwalt, "Today, highly trained, scared soldiers use all their senses, augmented with a coin detector and a stick." [1] The Defense Advanced Research Projects Agency (DARPA) offers the following advice to the developers of mine detectors. "When a soldier believes that a mine may be present - what does he reach for, a stick or your system?" [2].

Based on the comments of US Army countermining researchers in the literature and informal discussions between Army and Oak Ridge National Laboratory (ORNL) researchers, it appears that the Army is convinced that no single sensing technology is completely satisfactory for mine detection and sees the ultimate solution in the fusion of data from multiple sensor types. Hanshaw has described the specific fusion problem faced by the Army in some detail and has discussed selection criteria for what technologies should be fused, but does not give any details of specific data fusion algorithms. In describing how fusion operates, he considers how a soldier actually uses a metal detector. He argues that the audible output of the sensor is "fused" with the operator's visual view of the scene and that the brain's cognitive power melds these cues into an awareness of the size and location of the mine. Furthermore, the process is adaptive; the skill at combining the visual cues with audible data improves with experience [3].

*This research was performed at OAK RIDGE NATIONAL LABORATORY, managed by LOCKHEED MARTIN ENERGY RESEARCH CORP. for the U.S. DEPARTMENT OF ENERGY under contract DE-AC05-96OR22464.

Hanshaw defines fusion as the merging of several different data streams of differing importance into a state of awareness of whether or not a mine is present. A good fusion system would mimic animal reaction. Multi-sensor fusion is needed to simultaneously achieve high probability of detection, P_d , and low probability of false alarm, P_{fa} , thus breaking the universal tradeoff that plagues single sensor mine detectors [3]. For combat minefield breaching, both low P_d and high P_{fa} are equally deadly. Low P_d is the failure to detect mines that are actually present. High P_{fa} causes unnecessary halts in moving military formations, increasing their vulnerability to enemy fire.

Some of the specific elements needed to solve the fusion problem are reasonably well known. The problem in mine detection is "identity data fusion," and conventional data fusion methods address it with classic vector-based pattern recognition algorithms [4]. Data-fusion systems reported in the countermining literature use conventional vector-based methods. Sometimes these are neural nets [3] and sometimes they are based on simple thresholding of vector quantities [5].

The need for some kind of generalized data fusion method is recognized in the countermining literature [3]. The approach offered by Goutsias has possibilities, but appears to be primarily concerned with image processing [6]. In fact, a considerable proportion of recent research has concentrated on image processing [7,8,9,10,11,12]. The need for some kind of structural-based image processing is beginning to be acknowledged in the literature [13]. However, most of the recent countermining pattern recognition research continues to use vector-based methods [14,15,16,17,18,19,20,21,22]. It is noteworthy that the countermining literature reports no exploration of symbolic pattern recognition as a fusion method.

There is reason to believe that the mine detector multi-sensor data fusion problem is tractable. Dogs can consistently detect mines, and this suggests that mines must have detectable features [2]. Whatever form they may take, characterizing features exist and are identifiable. Notwithstanding the argument that a consistent set of characterizing features exists, and notwithstanding the fact that many have tried to find it, nobody has ever found a set of feature vectors that reliably characterizes mine identity. This strongly suggests that the features have not been found because researchers have been looking for the wrong thing. It is worth considering that dogs probably do not represent data as mathematical number lists, but they almost certainly represent data via semiotic structures. This possibility is borne out by recent research that suggests that high-level mammalian cognition is inherently symbolic [23,24,25,26].

THEORETICAL BACKGROUND

The argument is commonly, and reasonably, put forward by countermining researchers that if dogs can detect mines, then mines are detectable. Appeal to this argument is the usual justification for the development of electronic systems that seek to emulate the dog's nose. However, there is at least one account of a successful mine detecting dog, who, to all practical effect, does not have a nose [27]. This suggests that what should be emulated by an electronic system is the process by which the dog (or any generic biological system) fuses multiple sensory cues to identify the presence of a mine.

Does the idea of computational emulation of a biological process actually provide a practical basis for a novel approach to countermining data fusion? Recent research sponsored by Landauer and Bellman suggests that, in principle, it does [26]. Biological systems process signs and symbols to gain awareness of their environment, and their processing skill improves with experience. They commonly use the data represented by these symbols to perform classification and grouping, and *they do not do so by identifying boundaries between classes*. The way that biological systems perform classification suggests that there exists a unifying principle of classification that is applicable to computational systems [28].

Landauer and Bellman define semiotics as "the study of the appearance (visual or otherwise), meaning, and use of symbols and symbol systems." From their examination of classification by biological systems, they conclude that it would require a radical shift in how symbols are represented in computers to emulate the biological classification process in hardware. However, they argue that semiotic theory should provide the theoretical basis for just such a radical shift. Landauer and Bellman do not claim to have discovered the unifying semiotic principle of pattern recognition, but they suggest that it must be inductive in character [29].

Indeed, the development of a unified inductive learning model is the key to artificial intelligence [30,31]. Induction is defined as a mode of reasoning that increases the information content of a given body of data. The application to pattern recognition in general (and to the mine sensor fusion problem in particular) is obvious. An inductive pattern recognizer would learn the common characterizing attributes of all (possibly infinitely many) members of a class from observation of a finite (preferably small) set of samples from the class, and a finite set of samples not from the class. The problem arises due to the fact that neither of the commonly used "learning" paradigms, symbolic and numeric, are actually capable of performing induction.

A symbolic learning paradigm is one that uses a fundamentally "discrete" mathematical model. For example, any finite or infinite language over a finite alphabet is a discrete model. Inductive learning inherently requires continuous shades of meaning, and these are impossible to adequately represent in a discrete model.

Goldfarb's inductive learning model emerged from the attempt to find a general principle of pattern recognition that would accommodate both vector-based and structure-based pattern recognition as special cases. The fundamental idea in both cases is to devise an appropriate measure of similarity. Regardless of the pattern recognition method used, members of the same class are measurably similar, and members of different classes are dissimilar. The inductive learning paradigm includes a means of measuring the (dis)similarity of the semiotic structures of sample data sets.

The following description is taken from Goldfarb [30]. Consider a *transformation system* as a 3-tuple,

$$T = (O, S, D)$$

where $O = \{o_1, o_2, o_3, \dots\}$ is a set of *structured objects*, $S = \{S_i\}_{i=1}^m$ is a finite set of reversible *operations* that can transform one object into another, and $D = \{\Delta_\omega\}_{\omega \in \Omega}$ is a *competing family of distance functions* defined on O whose parameter set Ω is an $(m-1)$ -dimensional unit simplex

$$\Omega = \{\omega = \{w^1, \dots, w^m\} \mid w^i \geq 0, \sum_{i=1}^m w^i = 1\}.$$

Each of the distance functions, Δ_ω , is defined as follows. Weight w^i is assigned to operation S_i and $\Delta_\omega(o_1, o_2) = \min \sum_{i=1}^k w_{s_i}^i$, where s_i is a sequence of operations $s_i = (S_{i_1}^{o_1}, \dots, S_{i_k}^{o_k})$ that transforms o_1 into o_2 . In other words, the minimum is taken over all possible sequences of operations that can transform structured object o_1 into structured object o_2 . To assure that Δ_ω is defined for every pair of objects, assume that the set S of operations is *complete*, i.e., there are enough operations to transform any object o_1 into any other object o_2 . To use this definition to compute distance between structures, the system must use its set of operations in a *cooperative and competitive* manner. All properties that result from this definition should be viewed as emergent properties.

Conventional classifiers fall out as special cases of the above transformation system. For example, when used in a string-edit transformation system, the above distance function leads to the Levenshtein distance. Similarly, for a fixed weight vector, and a fixed set of coupled operations, the transformation system becomes a conventional neural net.

The ETS results from the consideration that the set of operations, S , is not sufficient to distinguish members of a class from non-members. In this case, it is necessary to add new operations to the set S . This process of adding operations to the transformation system leads to a sequence

of transformation systems, T_i ,

$$T_i = \{O, S_i, D, CR\}$$

where $i = 0, 1, 2, \dots$, and where all operations S_i , except for S_0 , are obtained from S_{i-1} by adding to it one or more operations that are constructed from the operations in S_{i-1} with the help of a small set of *composition rules*, CR . Each rule, $R_i \in CR$, allows the construction of new operations from existing operations.

Concept learning in the transformation system reduces to an optimization problem, to maximize $f(\omega)$, where $\omega \in \Omega$, $f(\omega) = f_1(\omega)/(c+f_2(\omega))$, and $f_1(\omega)$ is the Δ_ω -distance within the class of interest, and c is a small constant to prevent overflow. If Ω_{max} is a subset of Ω consisting of all the global maxima of f on Ω , then for a given class C and a given set of operations S , every weighting scheme, $\omega^* \in \Omega_{max}$, generates the "best geometric configuration" of the training samples. Ideally (i.e. without noise), for $\omega^* \in \Omega_{max}$, $f_2(\omega^*) = 0$, and $f_1(\omega^*) \neq 0$, the set of S operations produces complete separation between members and non-members of a class.

Inductive learning in the ETS occurs as follows. A sequence of operations S_i , is constructed in such a way that for the corresponding T_i , the minimum value of $f_2(\omega^*) =$ decreases, while $f_1(\omega^*) \neq 0$. Thus, the system learns new rules until it knows enough rules that the distance between all training samples in a class is zero (or close to zero in the presence of noise), and that the members of the class are a finite distance from all the non-members in the training set.

As a point of interest, it is reasonable to ask, is the ETS really a semiotic system? Meystel says that semiotic systems have the attribute of generalizing "the whole" by grouping specific "details." In other words, semiotic systems are "self-organizing," or equivalently, they exhibit "emergent behavior." In his introduction to Meystel's book, Albus puts it a little differently, saying that semiotic systems "are focused on increasing the knowledge about a situation which we seemingly perceive from the real world" [32]. Goldfarb believes that the inductive process is crucially important, and is the only process by which an agent increases semantic information about its environment. Furthermore, he says that the all the properties that emerge as consequences of his definition of the ETS are "emergent properties" [33]. Thus, ETS has several of the key attributes of a semiotic system, particularly since the very similarity measure is symbolic.

PROPOSED APPLICATION TO MINE DETECTION

The countermine literature suggests that reliable mine detection with high P_d , simultaneous with low P_{fa} , requires data from multiple sensors fused with a good algorithm [3].

Conventional vector-based systems, such as neural nets, have not produced the desired level of performance. The ETS is an alternative fusion scheme that is not vector-based, and may provide the breakthrough that enables robust, reliable mine detection with multiple sensor types.

Why might ETS succeed when vector-based methods have not been equal to the task? First, it is not a foregone conclusion that the mine sensor data output should be represented as a vector. Indeed, as already noted, countermine research has never discovered a reliable set of feature vectors for mine signatures. It is likely that some non-vector symbolic representation of the mine sensor output data remains to be discovered.

It should be noted that such a non-vector representation probably cannot be transformed into a vector representation and used with a neural net. Transforming a symbolic representation into a vector requires the assignment of numerical values to the symbols. This imposes an ordering on the symbols that is probably not indicated by the underlying physical reality. This is one reason that attempts to use vector-based methods to classify symbolic information have not been successful [33].

Even if the original features are represented as vectors, vector-based pattern recognition systems do not learn; they merely perform curve fitting. Thus, it is not reasonable to expect a vector-based data fuser to learn all the signatures for a certain class of mine by being trained on a small subset. The vector-based system merely fits a curve in a predetermined vector space, such that the in-class training samples are on one side of the curve, and the out-of-class training samples are on the other. The geometric structure of a finite-dimensional vector space is not modified (and inherently cannot be) by the training process.

On the other hand, the ETS does modify the geometric structure of the symbolic space in which it operates as it evolves. Essentially, as it evolves it learns the geometric structure of the mathematical space in which exemplars in-class are separable from exemplars out-of-class. Each step in the evolution produces a distance measure which is non-equivalent to the distance measure produced by the preceding step. The geometric properties of a space are characterized by the distance measure and the allowed set of operations.

In fact, vector-based systems can be seen to be a limited special case of ETS, in which characterization of data must be limited to vectors, and the properties of geometric space are not permitted to evolve. That is, the vector-based methods do not learn in the sense that ETS learns. Given these points, it is not surprising that ETS can learn class distinctions that are not noticed by a neural net. Also, if a dog can learn classes when the neural net cannot, then it is

fair to ask if the dog's behavior is not more like the ETS than like a vector-based system.

There are at present no results by any researcher using semiotic methods to indicate specific P_d or P_{fa} , or an operating curve indicating the tradeoffs between them for specific mine sensor suites. Based on the preceding discussion, the authors believe that ETS may lead to a significantly more favorable tradeoff between P_d and P_{fa} than the methods already tried by the countermine community. An extensive set of multiple-sensor cluttered data recently collected by DARPA would provide a good test of whether or not ETS can solve a pattern recognition problem that is widely recognized as unsolved by conventional methods.

CONCLUSIONS

The Army has a difficult problem in trying to identify landmine signatures from the output streams of multiple sensors, and to do so with simultaneous high P_d and low P_{fa} . Vector-based pattern recognition algorithms, which do not actually learn anything, but rather perform curve fitting in a predefined vector space, have so far proven unequal to the challenge. On the other hand dogs, which also rely on multi-sensory cues, do learn to identify mine signatures with high reliability. The ETS, which is not limited to vector representation or an *a priori* definition of geometry, is able to use a limited set of positive and negative exemplars to *actually learn* a set of transformation rules, and a consequent geometry. The geometry is specific to the training set, and simultaneously general to the entire class from which the positive training set is drawn. Dog learning appears to be much more like the ETS than vector-based pattern recognition. The ETS looks very promising as a data fusion engine for an electronic system to identify mine signatures from multiple types of mine sensors.

ACKNOWLEDGMENTS

The authors wish to thank Dr. Lev Goldfarb of the University of New Brunswick for taking the time to discuss the ETS in detail with the authors.

REFERENCES

- [1] Horowitz, P., *et. al.*, *New technological approaches to humanitarian demining*, Draft Report JSR-96-115, JASON, The MITRE Corporation, McLean VA, August 1996
- [2] Dugan, R., "Detection of Landmines and Unexploded Ordinance by Exploitation of the Chemical Signature," presented at DARPA UXO Detection by the Chemical Signature Conference, Arlington VA, 3 June 1996.

- [3] Hanshaw, T., "Multi-sensor fusion for the detection of mines and 'mine like' targets," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2496, pp. 152-158 (1995).
- [4] Hall, D.L., *Mathematical Techniques in Multisensor Data Fusion*, Artech House, Boston (1992).
- [5] Bargel, B., Bers, K., Stein, G. and Traeger, G., "Model-based sensor fusion for minefield detection," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2496, pp. 509-518 (1995).
- [6] Goutsias J.I., "Towards a Unifying Optimal Approach to Mine Detection Problems," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 2-13 (1996).
- [7] Abu-Tahnat, M. And Thompson, "Frequency-band Comparison Using Adaptive Filter for Multi-spectral Imagery," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 36-45 (1996).
- [8] Arigides, T., Libera, P., Fernandez, M.F. and Dobeck, G.J., "Adaptive-filter/feature-orthogonalization Processing String for Optimal LLRT Mine Classification in Side-scan Sonar Imagery," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 110-121 (1996).
- [9] Witherspoon, N.H., "Development and Operational Testing of COBRA System for Coastal Minefield Detection," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 14-14 (1996).
- [10] Muise, R.R., Wright, J.A. and Holmes, Q.A., "Coastal Mine Detection Using the Cobra Multispectral Sensor," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 15-24 (1996).
- [11] Earp, S.L., Marshall, J.W. and Anthony, E.R., "Prototype COBRA Near-Real-Time Processor," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 25-35 (1996).
- [12] Nevis, A.J., "Image Characterization and Target Recognition in the Surf Zone," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 46-57 (1996).
- [13] Banerji, A. and Goutsias, J.I., "Detection of Minelike Targets in Heavily Cluttered Environments Using the MNF Transform and Grayscale Morphological Image Reconstruction," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 58-68 (1996).
- [14] Schmalz, M.S., "Automated Detection and Recognition of Small Targets in Compressed Imagery," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 69-81 (1996).
- [15] Schmalz, M.S. and Ritter G.X., "Introduction to the Detection of Target Field Structure with Data Fusion of Density and Chord Information," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 397-408 (1996).
- [16] Bello, M.G., "Heirarchical Multilayer Perceptron Network-based Fusion Algorithms for Detection/Classification of Mines Using Multiple Acoustic Images and Magnetic Data," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 84-109 (1996).
- [17] Barnes C.F., "Sea Mine Detection with a Nearest Neighbor Classifier Based on Residual Vector Quantization," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 167-178 (1996).
- [18] Nelson, S.R. and Tuovila, S.M., "Automated Recognition of Acoustic Image Clutter," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 122-129 (1996).
- [19] Carter, P.H. and Dobeck, G.J., "Classification of Acoustic Backscatter Using the Generalized Target Description," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 190-200 (1996).
- [20] Gilmour, J.E. and Casasent, D.P., "Split-threshold Hit/miss Transform for Mines," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 149-156 (1996).
- [21] Goo, G. and Au W.W., "Detection and Identification of Buried Objects in Shallow Water," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 201-214 (1996).
- [22] Plett, G. and Doi, T., "Present and Future Methods of Mine Detection Using scattering Parameters and An artificial Neural Network," in *Detection Technologies for Mines and Minelike Targets*, Abinash C. Dubey, Editor, Proc SPIE 2765, pp. 385-396 (1996).

- [23] Prueitt, P.S., "Is Computation Something New?," in *Intelligent Systems: A Semiotic Perspective*, J. Albus, A. Meystel, and R. Quintero, Organizers, pp. 333-337, National Institute of Standards and Technology, Gaithersburg MD (1996).
- [24] Yufik, Y.M., "Memory, Complexity and Control in Biological and Artificial Systems," in *Intelligent Systems: A Semiotic Perspective*, J. Albus, A. Meystel, and R. Quintero, Organizers, pp. 311-317, National Institute of Standards and Technology, Gaithersburg MD (1996).
- [25] Pribram, K.H., "A Brain Model for Remembering," in *Intelligent Systems: A Semiotic Perspective*, J. Albus, A. Meystel, and R. Quintero, Organizers, pp. 318-320, National Institute of Standards and Technology, Gaithersburg MD (1996).
- [26] Landauer, C. and Bellman, K.L., "Semiotics of Constructed Complex Systems," in *Intelligent Systems: A Semiotic Perspective*, J. Albus, A. Meystel, and R. Quintero, Organizers, pp. 35-40, National Institute of Standards and Technology, Gaithersburg MD (1996).
- [27] Informal discussion with General Carl Stiner, former Commanding Officer of the US Southern Command.
- [28] Landauer, C., "Constructing Autonomous Software Systems," *Workshop on Biologically Inspired Autonomous Systems: Computation, Cognition, and Action*, Durham NC (March 4-5, 1996).
- [29] Bellman, K.L., "When Intelligence is in Control," in *Intelligent Systems: A Semiotic Perspective*, J. Albus, A. Meystel, and R. Quintero, Organizers, pp. 10-11, National Institute of Standards and Technology, Gaithersburg MD (1996).
- [30] Goldfarb, L. and Nigam, S., "The Unified Learning Paradigm: A Foundation for AI," in *Artificial Intelligence and Neural Networks: Steps Toward Principled Integration*, V. Honovar and L. Uhr, Editors, Academic Press, Boston (1994).
- [31] Goldfarb, L., "Inductive class representation and its central role in pattern recognition," in *Intelligent Systems: A Semiotic Perspective*, J. Albus, A. Meystel, and R. Quintero, Organizers, pp. 53-58, National Institute of Standards and Technology, Gaithersburg MD (1996).
- [32] Meystel, A., *Semiotic Modeling and Situation Analysis: An Introduction*, AdRem Inc, Bala Cynwyd PA, pp. viii, 13-18, (1995).
- [33] Goldfarb, L., Abela, J., Bhavsar, V.C. and Kamat, V.N., "Can a vector space learning model discover inductive class generalization in a symbolic environment?," in *Pattern Recognition Letters*, Vol. 16, pp. 719-726, (1995).