Appendix V

SIGNAL-PROCESSING AND SENSOR FUSION METHODS (PAPER II)

Paul Gader, University of Florida

SUMMARY

Signal processing is a necessary, fundamental component of all detection systems and can result in orders of magnitude improvement in the probability of detection (PD) versus false alarm rate (FAR) of almost any sensor system. A recent Multidisciplinary University Research Initiative (MURI) program (now ended) specifically aimed at landmines achieved dramatic improvements by developing and applying appropriate signal-processing techniques to data acquired from advanced sensors. Several of these techniques were transferred, or are currently being transferred, to industries developing prototype hardware systems for landmine detection. Basic research and development (R&D) programs specifically focused on signal processing for humanitarian demining with strong technology transfer components will greatly enhance the likelihood of further improvements with existing and developing sensing methodologies. Focused programs of this nature do not currently exist in the United States.

BASIC PRINCIPLES

Signal processing for landmine detection seeks to exploit discrimination information in measured signals from a variety of sensors. This information can be of many types, including frequency, shape, and size. Signal processing is used to mitigate system effects, characterize and discard sensor responses to the environment, compute discriminating features for classification of sets of measurements as either mines or nonmines, and provide feedback to operators.

Signals used in mine detection are generally multidimensional. Some example systems that are very familiar to this author are described here:

Vehicle- or Cart-Mounted Ground-Penetrating Radar (GPR)—These systems generally consist of arrays of transmitters and receivers and produce volumes of data. Each volume element can correspond either to spatial position on the surface and time (which roughly corresponds to depth) or to spatial position and frequency. The Vehicle-Mounted Mine Detector (VMMD) and Ground Standoff Mine Detection System (GSTAMIDS) are of this type, as are the Geo-Centers Humanitarian GPR and the NIITEK-fielded Wichmann GPR. The relative and perhaps absolute position of the measurements on these systems is usually known with a fair amount of accuracy. They provide multiple looks at most small mines.

Handheld GPR—These systems generally consist of small numbers of transmitters and receivers and produce sequences of radar returns. Currently, the Handheld Standoff Mine Detection System (HSTAMIDS) is the primary example of such a system. Researchers at Ohio State University and researchers in Europe, for example, are investigating other designs. No positional information is contained in the HSTAMIDS radar. These systems have repeatedly demonstrated the capability for detecting small, low-metal antipersonnel mines. Reducing the FAR is the current challenge—and that requires signal processing.

Cart-Mounted Acoustic/Seismic Systems—In these systems, a force is applied to the ground (of course, not with sufficient magnitude to set the mine off) and the resulting vibration of the ground is measured. Methods for measuring vibration include laser, radar, and ultrasound. The signals from these systems will also be three-dimensional, but there is no depth information here—only surface spatial position and velocity. Systems such as these have demonstrated the capability to find and characterize antipersonnel mines in both limited outdoor experiments at the Joint Unexploded Ordnance Coordination Office (JUXOCO) site and in the laboratory.

Spectral Imaging—Images formed in various spectral bands, usually some region in the infrared, can provide useful information, not only for mines but also for tripwires associated with mines. These systems can provide sequences of two-dimensional images. Often the images have value in several spectral bands.

STATE OF DEVELOPMENT

Signal-processing methods are currently being developed and implemented in both academic and industrial settings and are under evaluation in both laboratory and field settings. Many techniques are developed in universities and small companies and undergo transfer in some form to private companies engaged in sensor system development.

A very productive mode of operation has been to perform basic research in the university and transfer the results to industry. The university research relies on data collected by systems under development by the government. When signal-processing methods show promise, they are presented to industry and government. The methods, often in a modified, refined form, are adopted by industries and incorporated into their systems.

A very generic view of a typical signal-processing algorithm for demining is shown in Figure V.1.

Preprocessing can be extremely important and involves tracking the stationary and nonstationary statistics of the background and subsequently removing the background and normalizing the data, decon-

RAND*MR1608-V.1*

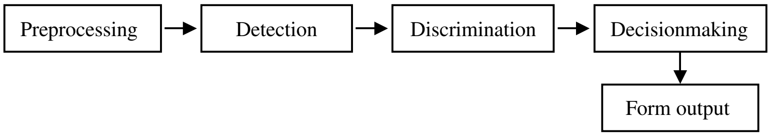

Figure V.1—Typical Signal-Processing Algorithm for Landmine Detection volution of the system response to the environment, removal of interferents such as radio frequency interference (RFI), and so on. In many systems, the response of the mines cannot be seen in visualizations of the data before preprocessing but are very apparent afterward. Preprocessing is by necessity adaptive, and system performance can degrade dramatically when the adaptive estimation methods fail.

Detection algorithms generally use preprocessed data and identify regions of interest that are then sometimes used for subsequent discrimination processing. Some detection algorithms are essentially anomaly detectors that look for differences from the background. That is, once the background has been removed, any remaining signal that is "sufficiently different" from the background is considered to be a mine. Constant false alarm rate algorithms are of this type. These algorithms must be adaptive. Other algorithms, such as hidden Markov models (HMMs), attempt to model the types of anomalies that constitute mines. These algorithms can be considered both detection and discrimination algorithms.

Discrimination algorithms try to model mines, and sometimes background and clutter, to characterize the mines. HMMs, matched filters, Generalized Likelihood Ratio Tests, Linear Discriminants, Support Vector Machines, and neural networks are all examples of these types of algorithms.

Both detection and discrimination algorithms often rely on the use of features. They can be computed mathematically, as is the case in Linear Discriminant Analysis, Principal and Independent Component Analysis, and wavelets, or they may be based on qualitative/ physical knowledge.

Decisionmaking involves post-processing—producing output responses for sensor fusion and for human operators. Post-processing can be used to reject responses based on such gross aggregate properties as size and shape. Responses may be produced for an operator, which is extremely important, or for the other sensors in multisensor fusion. If the system is not completely automated, i.e., responses are prepared for other algorithms or operators, then the algorithms must be more aggressive. Sensor fusion should generally require some quantitative knowledge.

CURRENT CAPABILITIES

Dramatic improvements have been achieved using signal processing on a variety of systems. The basic university research to developmental industry research has been a productive mode, resulting in tech transfer activities at various levels of completion in the VMMD, GSTAMIDS, University of Mississippi, HSTAMIDS Countermine and Humanitarian, Geo-Centers Humanitarian, and NIITEK Wichmann systems.

Early in the previous MURI, fuzzy clustering algorithms using the Defense Advanced Research Projects Agency (DARPA) backgrounds data reduced FARs from around 30 percent to around 4 percent as shown in Figure V.2.

That algorithm was refined and implemented in real time to the Geo-Centers VMMD system, which exceeded the exit criteria at the Advanced Technology Demonstrations in 1998, achieving over 90-percent detection at approximately 0.04 FAR in field testing on antitank mines. New algorithms based on HMMs improved on that performance in the lab, achieving performance of approximately 95-percent PD at approximately 0.02 FAR in the lab, again on antitank mines. This algorithm is currently undergoing tech transfer to the GSTAMIDS system and is currently achieving similar scores of

RAND*MR1608-V.2*

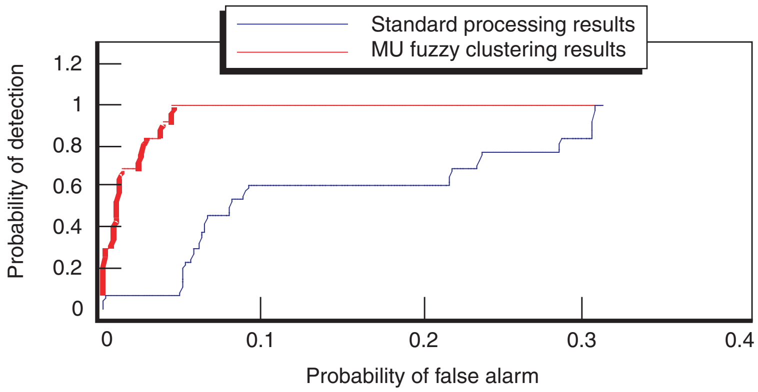

Figure V.2—Results of Signal Processing on DARPA Background Data around 95-percent PD at about 0.02 FAR in the laboratory. The HMM algorithms are currently being modified and implemented on the Geo-Centers Humanitarian system. HMMs show great promise, not only in this application but in several other applications in landmine detection.

Handheld systems, such as HSTAMIDS, are very important for the humanitarian demining mission. These systems have been tested, and the results are available in government reports. In addition, preliminary work with spatial algorithms has been implemented that decreases the FAR from 38 percent to 9 percent on the calibration area of the JUXOCO grid (see Figure V.3). When combined with an electromagnetic induction (EMI) processing algorithm, the performance of the automated spatial algorithm exceeded that of a human operator (see Appendix U). Thus, spatial signal processing has definite potential for reducing FARs and possibly reducing reliance on the expertise of the operator.

RAND*MR1608-V.3*

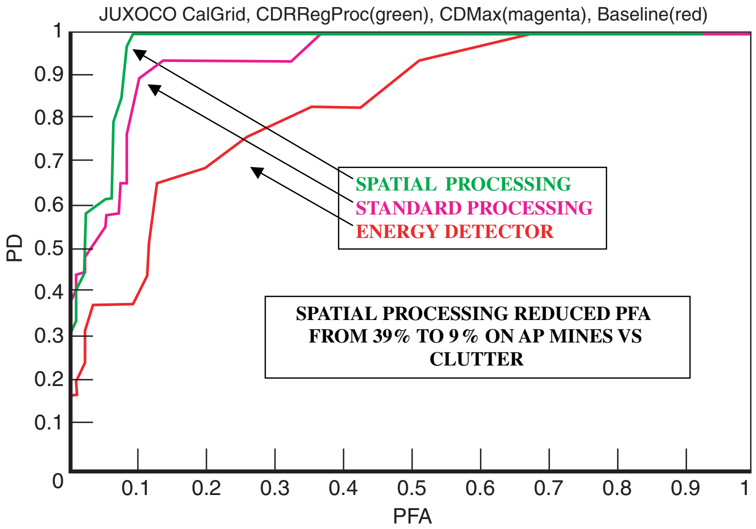

Figure V.3—Reduction in FAR Obtained from Spatial Processing of GPR Data Collected from a Handheld Unit

Image-processing algorithms have been developed for the acoustic-seismic system developed by the University of Mississippi. Data were acquired from some squares of the JUXOCO calibration grid and over grid squares containing low-metal antipersonnel mines and clutter objects, such as wood and plastic. The problem in this case was discriminating the antipersonnel mines from the clutter objects. Using a leave-one-out testing method, the algorithm outperformed the human experts, achieving the results shown in Figure V.4.

Signal-processing algorithms have consistently resulted in very significant drops in false alarm rates while maintaining probabilities of detection and can outperform expert human operators. Based on the experience of the past several years, this trend has not let up, and we expect these kinds of improvements to continue as new types of sensors and more advanced traditional sensors are developed.

RAND*MR1608-V.4*

(a) Using all

	FA	Mine
FA	27	0
Mine	1	11

(b) Leave-one-out

	FA	Mine
FA	27	0
Mine	1	11

(c) Human expert

	FA	Mine
FA	23	4
Mine	3	9

Figure V.4—Classification Results on the Acoustics Data

LIMITATIONS

Further R&D is required to achieve desired performance levels. Some limitations are discussed here. Although the limitations are grouped into discrete categories, they do overlap.

Limited Training and Testing Data

Efficient gathering, storing, and indexing of data sets is crucial for signal-processing algorithm development and analysis for humanitarian demining. Orders of magnitude improvements have been achieved through statistical and other techniques that require sufficient training data. The size of training, validation, and testing data sets may be quite large. A typical two-to-three-day data collection may result in many gigabytes of data. The ground truth for these data is often only partially accurate, requiring significant effort in ascertaining the exact location of the mine signatures in the data. Many such data sets are required for robust algorithm development.

A few standards are in place, such as the Countermine Test Management System truth file format, but no standard exists for linking the truth file format to data files produced by sensor data collection systems. There have also been some attempts at creating data repositories, but they are limited.

Discrimination in High Dimensions

Many of the dramatic advances in the past few years have come from better anomaly detection. Many simple algorithms detect any anomalous behavior in measured data. Much of the anomalous behavior is due to system cross-talk and nonstationary effects of the ambient environment, such as the ground and temperature of the equipment. The result was that many algorithms declared mines present when there was no object present. Detection algorithms have become much better at eliminating these kinds of problems. In addition, new sensor systems, such as the Wichmann GPR, have much cleaner signals, reducing the probability of these anomalies. We are currently at the point that region-of-interest detectors can be used to discard most of the measurements that are not associated with mines. The remaining measurements must be processed more carefully to discriminate mines and nonmines.

Thus, the problem in some cases has shifted more to one of discriminating between objects and is similar in nature to pattern recognition problems. A standard pattern recognition problem is to assume that a set of features has been measured from a region of interest and that one seeks to categorize the features as belonging to a finite set of classes.

In the case of landmines, we wish to measure features on sensor data and categorize the features as mine or not mine. The distinction here is that the set of objects that are not mines is not a coherent class. Thus, one cannot estimate the distribution of these objects or really define them as a class. (This problem is not isolated to landmines but occurs in many other real-world pattern recognition problems. It is not a solved problem.)

This implies the need to characterize precisely the class of mines. This is a more difficult task than it might seem because most representations of mines via sensor data require high-dimensional representations. Mine signatures can be highly variable, so complicated regions in feature space are required. We often use low-dimensional intuition to guide the development of high-dimensional methods. However, regions in high dimensions can be counterintuitive. For example, in a unit cube in high dimensions, the distance to the side is constant as a function of dimension, but the distance to the corner goes to infinity as the dimension goes to infinity. Another example of counterintuitive behavior is given by the Busemann-Petty conjecture. It seems reasonable to assume that if every symmetric, convex central slice through object A is bigger than every symmetric, convex central slice through object B, then object A is bigger than object B. However, this is only true for objects with dimensions less than 5. The characterization of class regions in high-dimensional space is still not well understood, but it is important for discriminating mines from clutter.

As another example, it is known that multilayer perceptions (MLPs) are excellent discriminants if the input pattern is known to come from a fixed, finite set of classes. However, an MLP cannot create a closed region in feature space unless the number of hidden nodes is greater than the number of input features. Even if this property holds, it is not guaranteed to form a closed region, and it is an extremely hard problem to check if the region is closed. Thus, MLPs are probably not reasonable for landmine discrimination problems.

Some methods, such as robust or possibilistic clustering, self-organizing feature maps, and relevance vector machines, provide promise of helping to achieve better understanding and high-dimensional modeling and hold promise for improved discrimination capabilities for many sensors.

Integration of Physical Models and Algorithms

Physical models offer great promise in helping to identify discriminating features, to identify bounds on variability of signatures, and even to create synthetic training sets. The results of physical models can be qualitative understanding of the phenomenology associated with a sensor, parametric mathematical models that can model the range of responses from a sensor, or actually synthetic “raw” data. Some examples of qualitative features that are currently under investigation are the “double humped” form of the returns in forward-looking radar and the transmission zeros found in acoustic-seismic systems. The prototypical mathematical models are the decay curves for EMI. Some synthetic data have been generated for multifrequency EMI and the Wichmann GPR, but conclusive results on robust accuracy are not known by this author. These methods are very computationally intensive. Continued research in this area is important.

System-Level Optimization

Detection algorithms consist of a sequence of steps. Often each step is developed independently and optimized relative to short-term criteria. For example, preprocessing algorithms are generally evaluated based on reduction of “noise,” which is sometimes quantitative and sometimes qualitative but almost never related to the end goal of PD versus FAR. We have found that in aided target recognition and other pattern analysis problems, “noise” at times can be anomalies associated with objects of interest and the process of cleaning the noise so that the data “look” better can reduce performance. Algorithms that consist of multiple steps should be optimized as a system, using the end goal as the objective. Unfortunately, these objectives and algorithms are not well behaved in a mathematical sense. The algorithms are complex, highly nonlinear sequences of steps. The relationships of the input distributions to the output distributions are generally extremely difficult or impossible to compute. Gradient descent on the objective function can also be difficult because gradients tend to vanish in multistage, complex systems. Stochastic search algorithms offer promise in this area. System-level optimization is an important but difficult unsolved problem in the development of signal-processing algorithms for humanitarian demining.

Adaptive Processing

Adaptive processing is crucial to landmine detection. Humanitarian systems may need to operate for many hours at a time over a range of ground and atmospheric conditions. Data acquired by most sensors are highly nonstationary. The data may change rapidly or slowly, and the density of mines may be high or low. Adaptive methods have been used in all signal-processing algorithms that we know of that have been transferred from basic research to industry. These methods generally adapt the statistics of the background. Currently, if the mine density is high, they will fail; they can also fail when conditions change too rapidly. These methods have not been tested over long periods.

In addition to background removal, adaptive processing is required for removal of interferents, such as RFI for quadrupole resonance detection. This is a very important area of research.

Multisensor fusion should be adaptive. Different sensors are useful in different conditions, and a dramatic change in behavior of one sensor system may not be present in another sensor system, providing a sort of “check and balance” system. Continued research is needed that develops more sophisticated adaptive methods for demining.

Incorporation of Spatial Information in Handheld Systems

Handheld systems are very relevant to the humanitarian problem. Mines have spatial signatures; they are not one-dimensional. Currently, there are few algorithms for incorporating spatial information into the algorithms on these systems. Incorporation of spatial information is difficult because the sampling in handheld systems is time-based and there is no machine control of the position of the sensor when samples are collected. That is, samples are collected at fixed time intervals over an irregular set of points on the surface. Preliminary work with HSTAMIDS data demonstrates that spatial information can dramatically improve performance.

Incorporation of User Feedback

Currently, very few systems incorporate user feedback. HSTAMIDS allows the user to turn off adaptation when a potential mine has been encountered. This is a simple form of user feedback. Certainly, user feedback must be simple enough to provide for use by a wide variety of operators, but this may limit the types of feedback that can be used. Consider a case in which a deminer has just found a mine. The data for that mine could be used to update the signal-processing algorithm to improve the capabilities of the system in order to find mines in the current environment.

POTENTIAL FOR IMPROVEMENT

Signal processing is a necessary and fundamental component of all detection systems and can result in orders of magnitude improvement in the PD versus FAR of almost any sensor system. Continued focused research in this domain is necessary to reach higher levels of performance.

OVERCOMING CURRENT LIMITATIONS: SENSIBLE R&D PROGRAM

A basic research program specifically aimed at signal processing for landmine detection should be created. Excellent results were achieved in the past, but there is currently no focused basic research program.

Primary goals of R&D in signal processing should be the development of new techniques, education, and technology transfer. Subsidiary goals should be the establishment of demining specific, user-friendly software tools in a MATLAB/C environment and standardized databases useful for entire communities of researchers. These goals are not mutually disjoint. The content of the research should address the limitations mentioned above.

Development of new techniques involves not only new signal-processing algorithms but also new methodologies for applying signal-processing concepts to the humanitarian problem. The academic literature is rich with ideas but short on detailed examination of the applicability of these ideas to large-scale, real-world problems, such as humanitarian demining. Such basic research is best carried out by combining academic research with industrial R&D.

Education involves graduate-level courses at universities, short courses at industrial and government locations and conferences, small workshops involving multidisciplinary teams, and production of written materials that are tutorial or educational in nature. Education also overlaps significantly with technology transfer.

Technology transfer often begins with university or small R&D company researchers using data from prototype systems developed by industrial contractors. Techniques are developed and demonstrated in the lab using these data. They are then presented to the government and the industrial contractors. If techniques display success in the lab, then computer programs and descriptions can be provided to the industrial contractors as well as assistance in interpreting and implementing them within the real-time system environment. Industry developers often modify and refine techniques as a result of further evaluation and testing.

Research at universities and small R&D firms should be supported directly by the government to ensure broad applicability of new methodologies. Incentives for industry to work with basic researchers are also necessary.

BIBLIOGRAPHY

Gader, P. D., J. M. Keller, and B. N. Nelson, "Recognition Technology for the Detection of Buried Land Mines," *IEEE Transactions on Fuzzy Systems*, Vol. 9, No. 1, February 2001, pp. 31–43.

Gader, P. D., M. Mystkowski, and Y. Zhao, "Landmine Detection with Ground Penetrating Radar Using Hidden Markov Models," *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 39, No. 6, June 2001, pp. 1231–1244.

Gader, P. D., B. Nelson, H. Frigui, G. Vaillette, and J. Keller, "Landmine Detection in Ground Penetrating Radar Using Fuzzy Logic," *Signal Processing, Special Issue on Fuzzy Logic in Signal Processing*, Vol. 80, No. 6, June 2000, pp. 1069–1084.

Gori, M., and F. Scarselli, "Are Multiplayer Perceptrons Adequate for Pattern Recognition and Verification?" *IEEE Transactions on Pattern Analysis and Machine Intelligence,* Vol. 20, No. 11, November 1998, pp. 1121–1132.

Ho, K. C., and P. D. Gader, "A Linear Prediction Land Mine Detection Algorithm for Hand Held Ground Penetrating Radar," *IEEE Transactions on Geoscience and Remote Sensing,* Vol. 40, No. 6, June 2002, pp. 1374–1385.

Ho, K. C., P. D. Gader, and J. B. Devaney, "Locate Mode Processing for Hand-Held Landmine Detection Using GPR," in *Detection and Remediation Technologies for Mines and Minelike Targets VII*, J. T. Broach, R. S. Harmon, and G. J. Dobeck, eds., Seattle: International Society for Optical Engineering, April 2002.

Hocaoglu, A. K., P. D. Gader, J. M. Keller, and B. N. Nelson, "Anti-Personnel Land Mine Detection and Discrimination Using Acoustic Data," *Journal of Subsurface Sensing Technologies and Applications,* Vol. 3, No. 2, April 2002, pp. 75–93.

Keller, J. M., P. D. Gader, Z. Cheng, and A. K. Hocaoglu, "Fourier Descriptor Features for Acoustic Landmine Detection," in *Detection and Remediation Technologies for Mines and Minelike Targets VII*, J. T. Broach, R. S. Harmon, and G. J. Dobeck, eds., Seattle: International Society for Optical Engineering, April 2002.

Rotundo, F., T. Altshuler, E. Rosen, C. Dion-Schwarz, and E. Ayers, *Report on the Advanced Technology Demonstration of the Vehicular Mounted Mine Detection Systems,* IDA Document D-2203, October 1998.

Stanley, R. J., S. Somanchi, and P. D. Gader, “The Impact of Weighted Density Distribution Function Features on Landmine Detection Using Hand-Held Units,” in *Detection and Remediation Technologies for Mines and Minelike Targets VII*, J. T. Broach, R. S. Harmon, and G. J. Dobeck, eds., Seattle: International Society for Optical Engineering, April 2002.