Appendix U

# SIGNAL-PROCESSING AND SENSOR FUSION METHODS (PAPER I)

*Leslie Collins, Duke University*

This appendix focuses on the impact of signal-processing techniques on the landmine detection problem and suggests research investments that will allow continued performance improvement. The focus is primarily on processing of electromagnetic induction (EMI) data, although results for other sensors as well as sensor fusion will be discussed.

## BASIC PHYSICAL PRINCIPLES

Signal-processing algorithms for landmine detection must detect the presence of an object in the geological background and discriminate signals associated with landmines from signals associated with discrete clutter objects. In general, signal-processing algorithms perform best when the physics that define the problem are integrated within the mathematical constructs underlying the theory of signal processing and pattern recognition. Utilizing experimental data measured under realistic conditions to test the performance of algorithms, and using insight from the data to guide the algorithm development process, has proven to be crucial for reducing false alarm rates in the landmine detection problem. The utilization of computational models describing sensor phenomenology, physics-based feature selection, statistical models of mines and clutter, and spatial information have all led to dramatic reductions in the false alarm rates of landmine detection systems. Because the physics that governs each sensor modality differs, feature sets extracted from data collected by different sensors usually are not consistent across sensors. However, several common approaches to processing the raw signals or the extracted features have been applied across sensor modalities.

In the landmine detection scenario, a particular sensor, or set of sensors, is used to interrogate one or several spatial locations for which a mine/no-mine decision is to be made. The sensing process may be automated, as is the case for vehicular or autonomous systems, or may involve a human manually operating the sensor. A sensor may record all of the response defined by the phenomenology associated with that sensor, or it may only record a portion of the response. Signal-processing algorithms for landmine detection are necessarily constrained by the available sensor data. Algorithms are also constrained by the system configuration as well as their impact on operator training requirements. For example, in the Army's Handheld Standoff Mine Detection System (HSTAMIDS), a soldier operates collocated ground-penetrating radar (GPR) and metal detection sensors in two distinct modes. In a scanning mode, processing of both sensors is done via causal systems, and spatial information is not explicitly incorporated into the processing. Once one of the sensors signals the presence of a mine-like object, meaning a *detection* is made, the operator enters an investigation mode where the sensor is operated differently. In this mode, algorithms could potentially utilize spatial information and operate in a noncausal mode. The operator utilizes the information from the two sensors, as well as from visual and environmental cues, to effect *discrimination* via sensor and information fusion, whereby nuisance clutter items are potentially ignored.

## STATE OF DEVELOPMENT

In recent years, there have been substantial improvements in sensor technology, with resulting improvements in the quality of the signals available from landmine detection sensors. In the EMI regime, Johns Hopkins University has developed a high-quality time-domain system capable of recording the EMI signal very early in the response time, and Geophex Ltd. has developed a frequency-domain system that operates over a fairly broad band. Both of these sensors measure the entire sensor response and are transitioning from the laboratory to the field. In contrast, the fielded EMI sensor—the PSS-12—provides a single time sample of the EMI response curve at every spatial position sampled. Other sensor manufacturers have developed systems sensitive enough to detect the extremely low metal content of plastic mines but often do not record the entire time- or frequency-domain signature. In GPR, the recently fielded Wichmann/NIITEK radar is capable of collecting remarkably clean broadband time-domain data, and the radars that are components of HSTAMIDS (CyTerra Corporation) and the Ground Standoff Mine Detection System (GSTAMIDS) (EG&G Inc.) have demonstrated good performance in several test environments. Both seismic and quadrupole resonance sensors have also been developed and are transitioning to field tests.

In 1996, the Army Research Office funded three five-year Multidisciplinary University Research Initiatives (MURIs) to investigate phenomenological studies and signal-processing research for the humanitarian landmine detection problem. Previously, most research had been performed by government laboratories and by government contractors building systems and primarily was not basic (6.1-level) research. Prior to the MURIs, the majority of the signal processing performed in contractor systems was anomaly detection, and little if any discrimination of clutter from mines was performed. This was particularly true in EMI sensors, where energy detection was the primary mode of operation. The development of sensors that are providing better data, and the focus of the MURIs and other government-sponsored programs on advanced signal-processing and sensor fusion research, has resulted in the development and transition of algorithms that are beginning to effect discrimination, and thus positively impact the false alarm rate. The models developed under the MURIs have supported the signal-processing research. Because signal processing traditionally tends to lag sensor development, some of these algorithms are just beginning to be tested in blind tests. The MURI-based research has also resulted in an improved sense of the optimal feature set to use when processing data from the various sensor modalities, as well as performance bounds on some of the feature extraction techniques. However, there is much additional research that could be performed as additional high-quality sensor data become available, particularly in the areas of model-based signal processing and of sensor fusion. While single-sensor processing algorithms are beginning to become more sophisticated, phenomenological models have only begun to be incorporated into the processing, and most sensor fusion algorithms that have been tested in this application area are still fairly basic.

## CURRENT CAPABILITIES AND OPERATING CHARACTERISTICS

An energy detector constitutes a signal-processing algorithm that is optimal for detecting a totally random signal in a totally random background. It assumes little if any a priori knowledge of the problem but is generally robust and is thus often used as an anomaly detector. It has been shown over the last several years that advanced signal-processing algorithms can reduce the false alarm rate substantially over such simple anomaly detection strategies. These more sophisticated algorithms have benefited from phenomenological models and understanding of the signal being sensed, advances in sensor capabilities, utilization of spatial data, and selection of physically based feature sets. Because the signals sensed from mines, background, and clutter are not deterministic quantities, statistical treatment of the various signatures has also had a positive impact on discrimination performance.

Metal detectors, for example, have advanced to the point where they can detect nearly all of the metal present in the environment down to tactical landmine burial depths. However, discriminating metal in a landmine from metallic clutter is a substantially more difficult problem. Similarly, discriminating a rock from a landmine in GPR data is more difficult than discriminating a landmine under the ground from the ground itself. Techniques that have been investigated to effect discrimination include Bayesian strategies, clustering techniques, hidden Markov models, inversion, support vector machines, and fuzzy processing. In a study performed at Duke University with the GEM-3 sensor, the false alarm rate was reduced by a factor of 10 when a statistical decision strategy was used in place of an anomaly detection (energy) strategy in a blind field trial. It is also possible to improve performance with systems that record only a portion of the received signal. For example, researchers at Auburn University demonstrated that the false alarm rate associated with the PSS-12 was reduced by a factor of 4 by processing the spatial pattern associated with the received signal measured over a suspect object. Various algorithms for EMI are currently being transitioned to fielded sensors to be tested, and algorithms are also being tested with some of the newer EMI technologies.

Algorithms for some of the newer technologies, such as acoustic and quadrupole resonance sensors, are less mature than those that have been developed for EMI and radar modalities. For example, the Quantum Magnetics quadrupole resonance sensor mitigates radio frequency interference in the demodulated sensor data via a least-mean-squares algorithm and then performs the detection using an energy detector in a band of frequencies around DC. Initial research has indicated that more advanced signal-processing techniques, such as Bayesian techniques or algorithms based on spectral estimation, may provide substantial reductions in the false alarm rate. However, such techniques must be tested on larger data sets to evaluate their robustness before definitive performance comparisons can be made.

Sensor fusion in systems currently being tested by the government is still in its research infancy. HSTAMIDS uses the operator to perform sensor fusion, although joint research between the University of Missouri–Rolla, University of Florida, Duke University, and CyTerra Corporation indicated that a fairly simple processing algorithm could meet or exceed the performance of the human operator *without* using any of the spatial information assumed to be used by the operator. Among other approaches, a voting scheme is being considered for GSTAMIDS, although more sophisticated techniques are being investigated. One reason for the lack of more sophisticated sensor fusion techniques is that collocated multisensor data have only recently become available in the community.

## LIMITATIONS AND RESTRICTIONS ON CAPABILITIES

Many of the discrimination techniques, as opposed to simple anomaly detection techniques, require training data. With the advent of accurate phenomenological models, training data that accurately mimic received sensor signals from mines under a variety of environmental and soil conditions are becoming available. Such training data will aid in the analysis of the robustness of discrimination algorithms. In addition to requiring training data for mines, most algorithms will require samples of background data local to the site under test to develop the statistics or features associated with the null hypothesis. Some mechanism for incorporating a priori knowledge of the class of targets likely to be present, environmental conditions, and other site-specific parameters into the processing algorithm by the sensor operators will also be necessary.

For most sensor modalities, there will always be discrete clutter items whose signature is similar enough to the signature of a mine that they will cause false alarms. Improved sensors should aid in this problem to some degree because more information can be extracted from the sensor and utilized in the signal-processing algorithm. Sensor fusion algorithm research, in concert with continued sensor development in alternative modalities (such as quadrupole resonance), should also help address this limitation.

Discrimination algorithms also usually assume isolated anomalies, i.e., the signatures of individual items to be discriminated do not overlap. Although some preliminary studies at Duke University have indicated that there are techniques to separate overlapping signatures, this remains a difficult research problem. For highly cluttered sites, the development of robust algorithms to detect the presence of overlapping signatures and then separate the signatures prior to applying the discrimination algorithms is needed.

Other limitations include issues involving the necessity for real-time processing and training, and the requirement that algorithms are required to be simple for operators to execute. For example, Bayesian theory prescribes the optimal processor for the two-hypothesis testing problem, but this approach requires precise knowledge of the probability density functions describing the data under each hypothesis and sometimes requires a multidimensional integration. Both of these requirements could limit the applicability of the Bayesian approach in a field-deployed system. As another example, a discrimination algorithm could perform extremely well but require carefully controlled spatial data collections, which may be difficult to train an operator to perform but which might be possible with an automated system.

## ESTIMATED POTENTIAL FOR IMPROVEMENT

Dramatic performance improvements have been demonstrated over the past few years using improved signal processing that has been based on an improved understanding of the phenomenological underpinnings of the landmine detection problem. Additional performance improvements have been demonstrated via improved sensors. Because several new sensors are entering field tests, the potential for order of magnitude improvements in the speed at which landmines can be cleared is feasible.

There are several promising GPR technologies for which algorithm development is in its infancy. The Wichmann/NIITEK sensor operated in a simplistic anomaly detection (energy detection) mode is performing comparably to other radars with more advanced processing algorithms that are meeting the government exit criteria for handheld and vehicular detection systems. Preliminary tests with more advanced algorithms indicate that an order of magnitude improvement in false alarm rates at appropriate scanning rates will be achievable with this system in the two-to-seven-year time frame.

Most fielded EMI systems have not been optimized for the landmine detection problem per se and measure only a limited portion of the available signal. Initial evaluations of more advanced EMI systems utilizing statistical signal-processing algorithms in blind tests have suggested that an order of magnitude reduction in the false alarm rate is also possible with these systems. Additional tests in traditional government test sites are ongoing and appear to support the preliminary results. Quadrupole resonance is also a promising technology for use as a confirming sensor to further reduce the false alarm rate.

Several sensor technologies are near the point that they could be used to *individually* reduce the false alarm rate by an order of magnitude, particularly in realistically cluttered test sites. The combination of these technologies with appropriate sensor fusion algorithms, and confirmatory sensors, has the potential to dramatically reduce the false alarm rate and thus provide order of magnitude increases in landmine clearance rates. To achieve this goal, investments should be made in the individual sensor technologies, signal-processing research for each of the technologies, and basic research in sensor fusion. Additional care must be taken to ensure that the system-level operation is appropriate for soldiers or indigenous populations.

## OUTLINE OF A RESEARCH AND DEVELOPMENT PROGRAM

Several well-established research results should be used to guide the design of a research and development program. These include the following:

- Incorporation of the phenomenology associated with a particular sensing modality directly into the signal-processing algorithm or into the feature selection can result in substantial performance improvement.
- Advanced signal-processing techniques can improve discrimination performance over energy-based, differential-energy based, or anomaly detection techniques.
- Multisensor systems outperform single-sensor systems, even with fairly rudimentary sensor fusion algorithms.
- Multiple sensor designs within the same sensor modality can each perform well and may generate different false alarms.
- Utilization of spatial data improves performance.
- In EMI systems, the signatures from multiple objects combine approximately linearly, and can be separated under some conditions.

These results suggest parallel development of multiple sensors in concert with signal-processing algorithms, instead of isolating the research associated with each task. It also suggests considering a system-level optimization as part of the down-select for sensors and algorithms to be incorporated into the system. Substantial reductions in the false alarm rate could be achieved by leveraging and extending previous research findings, including those listed above, and by continuing a basic research program in humanitarian demining.

## REFERENCES

Most recent field tests of systems have been reported at the UXO/Countermine Forum and/or the International Society for Optical Engineering (SPIE) Detection and Remediation Technologies for Mines and Minelike Targets Conference. Research results from the three Humanitarian Demining MURIs and other Department of Defense (DoD)–funded projects are most commonly reported at SPIE. In addition, interim and final reports from the MURIs and other DoD-supported research programs are available from the sponsors. The Joint Unexploded Ordnance Coordination Office website (www.uxocoe.brtrc.com/lib.htm) has also archived several journal articles and technical reports.