Appendix F

GROUND-PENETRATING RADAR (PAPER II)

James Ralston, Anne Andrews, Frank Rotondo, and Michael Tuley,
Institute for Defense Analyses[1]

BACKGROUND AND SUMMARY

It is understood that information on the performance of individual landmine sensors and groups of sensors is being considered in the context of a mine remediation process that includes site identification, the location of individual mines, and appropriate mine-by-mine remediation measures. This understanding must include several key points, detailed in the following paragraphs:

Study the remediation process. This must be done before attempting to model the process. It must also be recognized that remediation includes any measures applied to the mined area before sweeping with detection systems. For example, it is sometimes the practice to use heavy rollers in an attempt to detonate some mines before sweeping. This can be effective in preemptively disposing of many mines at relatively low cost and risk, but the aftermath can be detrimental to the function of detecting sensors if some unexploded mines are inadvertently buried to greater depths, misoriented, or tilted so that they become hazardous to probe manually, or if their soil context is changed in ways that may affect certain sensors, such as radar, acoustics, or infrared. Any attempt to draw conclusions about enhanced remediation on the basis of modeling must first ascertain that the model's assumptions are realistic.

---

[1] Originally published by the Institute for Defense Analyses, Alexandria, Va. Reprinted by permission.

Understand local differences. For a variety of political, cultural, and economic reasons the process of landmine remediation proceeds differently in different countries. This must be respected in assigning overall costs to various stages of the process. For example, operations that would be "prohibitively labor intensive" in one country may be acceptable in nations that have large populations of unemployed farmers. Elapsed time or man-hours is not generally a good proxy for cost. Properly administered, mine remediation in such countries may represent economic as well as humanitarian assistance.

Fusion is critical. The goal should be to find not the best sensor but the best combination of sensors. Fusing the two best individual sensors does not necessarily lead to the best fused performance. Data collections should emphasize establishing statistical correlations between detections and false alarms of all sensors, rather than simply measuring individual sensor performance.

Performance is situational. All sensors will work better in some environments and against some threats than in and against others. The best choice of sensors and fusion techniques will depend on the specifics of each environment.

Prospects for improvement. The most likely research and investment areas leading to possible significant reductions in mine remediation rate are as follows:

- Reducing false alarms by emphasizing fusion of multiple sensors or the use of multiple stages of detection and confirmation sensors.
- Identifying the specific classes of false alarms whose characteristic features can be recognized and confidently rejected. The classes may include particular man-made objects (e.g., soft-drink cans) as well as naturally occurring soil phenomena. The feature set investigated should not be limited to the output of a single sensor type.
- Seeking cost reductions in the most advantageous suite of sensors and remediation techniques, once it is identified.

Andrews, Ralston, and Tuley [1] document the results of an Institute for Defense Analyses (IDA) assessment of the state of current research on ground-penetrating radar (GPR) as applied to countermine and unexploded ordnance (UXO) clearance. Despite significant long-term investment in GPR for mine and UXO detection, it remains true that no GPR system that meets operational requirements has yet been fielded; however, recent advances in several mine detection radars under development have produced significant improvements in detection performance and false-alarm mitigation over what was achievable only a few years ago. The authors' report examines existing GPR research and development efforts with emphasis on missions where GPR has the potential to provide a unique capability and to achieve operationally meaningful performance. It identifies data collections and analyses that will be necessary both to make decisions about the suitability of GPR for particular missions and to achieve performance gains necessary for operational utility.

The Joint Unexploded Ordnance Coordination Office (JUXOCO) sponsored a GPR workshop held at IDA in June 1999. Investigators from all currently funded GPR efforts in countermine and UXO were invited to present their work. An independent panel representing government, federally funded research and development center, and university expertise in GPR was assembled to assist the government in the assessment role.[2] Panel discussions were held during the course of the three-day workshop and a one-day follow-on meeting. The report of Andrews, Ralston, and Tuley [1] does not an attempt to express a consensus of the panel, which likely does not exist; however, the comments and insights provided by panel members are reflected in the emphasis and conclusions of this report.

CONCLUSIONS

The principal conclusions of *Research on Ground Penetrating Radar for Detection of Mines and Unexploded Ordnance* [1] are as follows:

- Phenomenology controlling performance is not sufficiently well understood. Advanced development work must be preceded by concomitant research understanding.

---

[2] Presentations to the workshop are summarized in Andrews, Ralston, and Tuley [1].

- Too little analysis has been carried out on the data that have been obtained. A synergistic analysis effort that stretches across programs might provide real dividends.
- While there are exceptions, current system performance is typically limited by false alarms. That is, detection is clutter limited, not noise limited. Only when target and clutter characteristics are both well understood can signal processing be applied effectively.
- Much more effort has been spent studying target characteristics than has been spent on clutter. Efforts defining target signatures are necessary, and target-related research should continue; however, substantial efforts must be focused on clutter research and data collection.
- Predicting performance requires understanding sensitivities to the environment. Models will not provide the realistic data useful in algorithm development until the understanding of clutter is improved.
- Incorporation of diverse expertise in sensor hardware, algorithm development, modeling, and testing has been beneficial. The Multidisciplinary University Research Initiative (MURI) and the red team approach to the Handheld Standoff Mine Detection System (HSTAMIDS) program have resulted in a better understanding of the sensor functionality and performance improvements.
- There is a need for controlled, repeatable testing to evaluate sensor performance independent of operator skill and technique, and not subject to uncontrollable alterations in the environment. This capability is important for comparing different sensors and tracking changes in performance with sensor modifications.

RECOMMENDATIONS

1. GPR countermine performance is limited by clutter, and clutter is not well understood.[3] Thus, the focus of research should be on

---

[3] *Clutter* is defined to be returns identified by the sensor system as targets that do not correspond either to intended targets or to random system noise; that is, clutter repre-

defining, understanding, and measuring clutter. To that end, the following steps should be undertaken:

- Determine the range of clutter and target data needed to support system design decisions, algorithm development, and modeling research.
- Build a suite of research-quality data-collection instruments not constrained by operational requirements.
- Collect and analyze clutter and target data, with a focus on clutter. Data collection should be driven by three concerns: better understanding clutter characteristics, providing training and test data for signal-processing algorithm development, and providing both input and validation data for electromagnetic model development.
- Table F.1 provides our recommendations for the system design and parameter space to be covered by the instruments and the data collection. These recommendations are for reasonable, notional parameters for the instruments and the experiments, but they do not represent the results of a rigorous study of the trade space or practical engineering considerations. As such, final designs should be based on an extensive red team effort involving hardware engineers, signal processors, modelers, and test designers.
- Develop a research program to provide the necessary knowledge of clutter characteristics. Such a program should involve a careful physical and electromagnetic description of environments of interest, ranked in order of importance. These could be used to prioritize data collections. Clutter is highly variable, and that complicates its description. The focus of the research should be an attempt to group clutter into a limited number of classes relevant to system design. To that end, a careful evaluation of a combination of statistical and discrete approaches for clutter characterization is warranted.

---

sents real sensor responses to discrete items or environmental conditions that are not of interest.

Table F.1

Data Collection Matrix

Parameter	Value Range
Frequency range	200 MHz–6 GHz
Polarization	Full
Grazing angle	Full hemisphere
Aspect angle	Full hemisphere
Road/terrain/area	Increasingly complex media, small patches
Target type/configuration/ quantity	Individual target interrogation: —Buried mines —UXO —Discrete clutter objects —Standard metal and dielectric targets
Spatial resolution	Less than minimum target size, best attainable with radar, centimeters
Waveform	Stepped frequency
Azimuthal processing	Three-dimensional SAR
Antenna height	Close coupled to earth

- — Support research on the characterization of electromagnetic propagation and scattering in soils. Investigate a statistical paradigm similar to the atmospheric weak-scattering case. Bolster theoretical analysis with carefully calibrated measurements and computer modeling. Efforts should begin on simple, well-characterized media. As understanding is gained, more complex compositions should be tackled.

2. We should do a better job of exploiting data from current programs. There are two important facets of such an effort:

   - — Make data and specific analyses deliverable from contractors. Every effort should be made to ensure that data collections and analyses serve the broader goals of the countermine program.
   - — Set aside resources for independent analysis of data. Such efforts provide potentially valuable insights that are not likely to come out of program-driven analyses. An example is

the red team analysis of HSTAMIDS[4] data, which provided significant input to focus system improvements.

3. The HSTAMIDS red team is an example of how accessing a larger body of knowledge in the countermine area can pay dividends for a specific program. Research results coming out of MURI and applied to data from the BoomSAR, Wichmann, and Geo-Centers systems show significant performance improvements. Such interactions should be encouraged through a red team approach to system engineering decisions.

4. Measurement, modeling, and detection/discrimination algorithm development must be tightly integrated. As discrimination of mines from clutter is typically the problem faced by mine detection systems, discrimination algorithm development is the key to performance improvement. Algorithm success depends on the signals provided.

5. Sensors delivered to the government at the end of programs should be well documented and well calibrated.

6. Existing platforms from other Department of Defense programs should be leveraged to the extent possible. Specifically, the Defense Advanced Research Projects Agency ultra-high-frequency ultra-wide-band synthetic aperture radar (DARPA UHF UWB SAR) and the Army Communications–Electronics Command tactical unmanned aerial vehicle SAR should be tasked for data collection and baseline performance determination for countermine and UXO detection.

7. Develop protocols and equipment for standardized sensor testing.

8. Other specific recommendations are summarized in Table F.2.

---

[4] HSTAMIDS is under development for the U.S. Army. The system incorporates both GPR and electromagnetic induction (EMI) sensors.

Table F.2

Other Recommendations

Soil characterization	Develop a statistical description of soils, patterned on atmospheric physics Develop numerical modeling approaches that accurately represent realistic soils Initiate a measurement program to support above
Discrimination	Continue modeling efforts to identify discriminants Use above data collection for signal processing Investigate utility of polarization Investigate utility of spectral response Curtail complex natural resonance research Curtail 3rd harmonic research
Fusion	Require analysis and reporting of target and clutter statistics for current data Make raw and processed data deliverable. Initiate independent analysis Task collection of coregistered data sets for forward-looking radar with NQR and forward-looking and down-looking radar

HANDHELD DETECTOR PERFORMANCE

Below we survey results of recent tests of HSTAMIDS. This sensor system incorporates both a GPR and a metal detector. In the most recent testing [2], both sensors were used, and test results reflect a fusion of both sensor indications. In earlier tests [3], the performance of the individual sensors was separately tabulated. Here we use HSTAMIDS as a proxy for a state-of-the-art GPR sensor for small mine detection to quantify what performance improvements are necessary.

All of these results are influenced by the location and procedures of the particular test, as well as previous engineering design choices in the systems tested. Thus they are not easily extrapolated to general statements about capability. This points to important considerations about testing and evaluating detection systems. That is, there is a need to evaluate sensors in a standard environment on a common and reproducible target set where the influences of the operator are eliminated. This is particularly important for tracking system development and for evaluating competing sensor concepts. In field tests, true performance of sensors can be masked by contributions from footprint and coverage, as well as operator skill or fatigue, exploitation of visual cues, familiarity with the test site, or the means of presenting the data to the operator.

The HSTAMIDS Operational Requirements Document sets the requirement for detection of surface and buried antipersonnel and antitank mines at probability of detection (PD) = 0.90, with a false alarm rate not to exceed 0.6 per square meter [4]. Historically, handheld GPR sensors have been stressed by low-metal-content antipersonnel mines. (See, for example, Andrews et al. [5], where systems incorporating GPR and EMI sensors were tested in 1996.) Recent modifications in sensors under the HSTAMIDS program have resulted in performance improvements over what was achieved only a few years ago [2].

CLUTTER

A review of baseline data quickly leads to the conclusion that, in many cases, the fundamental problem of GPR performance is not the absence of a sufficient mine- or UXO-generated signal for the radar to detect.[5] Rather, the problem is the multitude of signals originating from surface and buried clutter. Here, the concern is separating target signals from clutter signals. Thus a thorough understanding of clutter becomes fundamental to understanding GPR performance and limitations.

We define *clutter* as returns identified by the sensor system as targets that do not correspond to intended targets or system noise, that is, real sensor responses to discrete items or environmental conditions that are not of interest. Clutter might be considered as a set of area- or volume-extensive attributes of the environment in which a GPR must work. Conversely, we might think of clutter as a collection of discrete, but undesired, targets to be separated from those we desire to detect. It is likely worthwhile to employ a mix of both views of

---

[5] An exception is likely to be detection of dielectric mines buried in soils with similar dielectric constants, where little or no contrast may exist.

clutter, and we do so here by defining two clutter study modalities: volume and discrete.

In either case, clutter statistics, which will quantify the ability of a feature or set of features to separate targets and clutter, must be measured for each potential discrimination feature of interest and will differ for each radar configuration. A library of target and clutter measurements taken in a variety of environments can be used to determine the robustness of clutter suppression approaches.

Features of discrete clutter objects (e.g., rocks, roots, cans, water-filled inclusions) may allow for discrimination and identification to reduce false alarms. It is possible but not currently known whether a significant fraction of false alarms arise from a small number of discrete types of clutter. This makes identification of the sources of false alarms an imperative part of any sensor improvement clutter study. If discrete objects or features of the ground can be identified and characterized, they can be screened out where they differ sufficiently from targets.

Studies of clutter are often neglected because of the desire to obtain data that are universally useful. Because the sensor must perform in a highly variable and continuously changing clutter environment, however, ways must be found to understand clutter. Framing experiments to study clutter is difficult because any study will be of the specific clutter at a specific site as seen by a specific instrument. There will be great variability in the clutter itself, based on uncontrollable variables such as geology, climate, and history of use. Clutter at the same site may have temporal variations depending on recent weather patterns. Further, the clutter will depend on features of the radar itself, such as grazing angle, spot size, resolution, frequency band, and polarization, as well as the processing. A clutter experiment therefore will require careful research planning, data collection, and analysis. Ongoing programs provide numerous data collection opportunities. There is a need to make sure the right data are taken, that they are analyzed, and that the analysis parameterizes the data in a meaningful way.

Recently, DARPA conducted a data-collection program that focused on understanding clutter for the buried mine and UXO problem [6]. Numerous sensors were tasked to survey four 1-hectare sites, with the goal of producing co-registered clutter maps from multiple sensor modalities for magnetometry, EMI, radar, and infrared sensors. All three radars in the study were sensors of opportunity, so in any event the data collected explored dimensions determined by previous design choices. The experiment also experienced navigation difficulties that made interpretation of the data difficult. Nevertheless, algorithm work done by Paul Gader on Geo-Centers radar data set resulted in a many-fold decrease in the density of false alarms at a comparable PD [7].

The necessary clutter study will require an effort such as this. This experiment should be conducted in a variety of clutter environments on well-characterized sites. Research-grade sensors with the flexibility to explore the widest accessible parameter space are required. Only through an effort such as this can we build the library of clutter data necessary for making engineering design choices; supporting modeling of real-world conditions, signal processing, and algorithm development; and for determining the robustness of sensor performance.

GPR SYSTEM CONSIDERATIONS

GPR design is complex and challenging because of the array of hardware and system choices and the coupling of many of those choices. This section provides a brief discussion of some of the choices that can be made and their implications.

The most fundamental choice in GPR is the center frequency and bandwidth of the radar. Low frequencies provide improved soil penetration; the depth at which targets must be detected and the soil types within which they must be detected drive the choice for the lowest frequencies to be transmitted. For example, UXO detection would generally call for lower frequencies than mine detection because of the greater depths at which targets may be located. Practical limits on low-frequency performance are often determined by the maximum size of the antenna that can be deployed. Range resolution is governed by bandwidth, with the achievable resolution given as the speed of light in the medium divided by twice the bandwidth:

$$\Delta R = \frac{c}{2B\sqrt{\mu_r \varepsilon_r}},$$

where c is the speed of light, B is the bandwidth, μ_r is the relative permeability, and ε_r is the relative dielectric constant. Thus, if high resolution in range is desired, wide bandwidth is required, and the higher the center frequency, the narrower the percentage bandwidth for a given resolution and the more straightforward the radar design job. Because of the dispersive properties of soil, high frequencies will be attenuated more than low frequencies. Rather than considering the waveform that is transmitted, the GPR designer must plan his processing and detection strategies around the expected spectrum of the return after propagation to the target, reflection, and propagation back to the radar antenna. Thus, having low frequencies that penetrate well may be of little consequence if the detection algorithm depends on fine resolution and the higher frequencies that provide bandwidth are severely attenuated. The chosen frequency regime also controls less obvious radar characteristics, such as achievable cross-range resolution, in SAR systems and the level of radio frequency interference (RFI) with which the system must contend.

Most GPRs for mine detection are wideband devices because good range resolution is required to separate targets from clutter. Two general approaches to obtaining wideband performance are available to the system designer. Each has advantages and disadvantages. The first utilizes waveforms having time-bandwidth product that is near unity. These systems are represented by the family of impulse radars that have been developed for ground-penetration missions. The major advantages of an impulse radar are that lower dynamic range receivers are required to discriminate against clutter, the waveform generation time is short, and a high-range resolution display is available with little or no processing. The major disadvantages are the need to control radio frequency dispersion over a wide instantaneous bandwidth, susceptibility to RFI because of the wideband receiver front end, the need for very-high-speed analog-to-digital converters (or the inefficiency of a sampling oscilloscope approach) for waveform capture, and difficulty in controlling details of the transmitted spectrum.

The alternative to impulse is to employ a waveform with a time-bandwidth product much greater than one. Such systems have been implemented using stepped frequency, linear FM (frequency modulation) chirp, or phase codes. The major advantage of stepped frequency or linear FM chirp is that the frequency spectrum can easily be chosen to fit what the designer considers optimum. In fact, notches can even be placed in the transmitted spectrum to avoid interference with or by other systems. Stepped-frequency waveforms in particular allow narrow instantaneous receiver bandwidth, lower bandwidth analog-to-digital converters, and wider dynamic ranges. This last advantage is often offset by a need for the wider dynamic range because the large surface clutter return and target returns are not temporally separated as they are in an impulse system. Other advantages of high time-bandwidth product waveforms are higher average powers and an ability to tailor the frequency response on receive through processing. Phase and amplitude calibration and equalization are easily accomplished at each discrete frequency step. The major disadvantages are the required dynamic range mentioned above and the time required to generate one complete waveform.

A waveform and bandwidth having been chosen, the GPR designer must implement an antenna commensurate with the bandwidth. Antenna design becomes particularly critical in systems whose geometry provides little standoff from the surface of earth. A major obstacle in using wideband antennas is eliminating internal reflections over their entire frequency band. Such internal reflections result in antenna “ringing” that can hide target returns in systems that operate close to the surface. In down-looking systems, that problem is exacerbated by the very large surface clutter return that may also reverberate within a poorly matched antenna structure. Closely coupled antennas can reduce that problem, as can the use of cross-polarized antennas that tend to discriminate against the surface clutter return. While no designer would intentionally choose an antenna known to produce significant internal reflections, two design approaches are generally viable. In one, the designer makes heroic attempts to reduce antenna internal reflections, thereby simplifying the signal-processing problem. Such an approach is illustrated by the Wichmann radar, where the array antenna is designed to minimize both internal reflections and reverberation between the antenna face and the ground surface. The second option is to have a certain amount of antenna internal reflections and take those out in signal processing. This is most easily done with a stepped-frequency system, where the internal reflections can, in principle, be measured at each frequency and then coherently subtracted from the return. The flaw in such an approach is that the reflections will depend to some extent on the details of the surface clutter return, and as that changes, coherent subtraction may be less effective. Internal reflections are of less concern in standoff radars, but at the low frequencies often employed in GPR they may be a problem even in that case. In particular, at low frequencies the entire structure on which the antenna is mounted becomes part of the radiating structure, and reverberations may linger in time. That problem has been noted in several GPR implementations.

Signal-processing and display options are strongly driven by the waveform choice and the antenna implementation. For example, with a single antenna that is manually scanned, it is very difficult to generate a display output more sophisticated than a simple one-dimensional range profile. Such a display is available with little or no processing from an impulse system and with simple pulse-compression processing from a large time-bandwidth product system. Linear array antennas or those antennas that scan across mine lanes provide more flexibility in display. A linear array can be used to provide a waterfall plot of time (range) images closely spaced in the cross-track direction, as in the Wichmann radar; a real aperture two-dimensional image, as in the Geo-Centers display; or a form of synthetic aperture image. A scanned antenna can also be used to produce a synthetic aperture image. Finally, scanning in both cross-track and down-track allows formulation of three-dimensional images. Doing so, however, requires careful attention to knowledge of antenna position and correction of propagation effects within the soil.

There should be a clear connection between an operational concept for GPR in countermine operations and a focused plan for conducting the necessary supporting research and obtaining needed engineering design data. The focus of the JUXOCO report [1] is on defining the research necessary to design an optimal operational system, not to design the operational system forthwith. As such, it is important to remember that the recommendations herein are for instrumentation and experiments to collect necessary data rather than to support operational requirements. We will generally want data-collection experiments to cover a wider range of operating parameters than would be practical for a fielded system so that we can be confident that the limits of the operating system are optimally chosen.

REFERENCES

1. A. Andrews, J. Ralston, and M. Tuley, *Research on Ground Penetrating Radar for Detection of Mines and Unexploded Ordnance: Current Status and Research Strategy*, Alexandria, Va.: Institute for Defense Analyses, D-2416, December 1999.
2. F. Rotondo, E. Ayers, A. Calhoun, E. Rosen, and L. Zheng, *Engineering Development Test Results for the Handheld Standoff Mine Detection System: May 2000*, Alexandria, Va.: Institute for Defense Analyses, D-2510, March 2000.
3. F. Rotondo, E. Rosen, and E. Ayers, *Test Methodology and Results for the Handheld Standoff Mine Detection System in Check Tests 1 and 2: June–October 1999*, Alexandria, Va.: Institute for Defense Analyses, D-2443, March 2000.
4. Operational Requirements Document for the Handheld Standoff Minefield Detection System, U.S. Army Training and Doctrine Command, August 19, 1995.
5. A. M. Andrews, T. W. Altshuler, E. M. Rosen, and L. J. Porter, *Performance in December 1996 Hand-Held Landmine Detection Tests at APG, Coleman Research Corp. (CRC), GDE Systems, Inc. (GDE), and AN/PSS-12*, Alexandria, Va.: Institute for Defense Analyses, D-2126, March 1998.
6. V. George and T. W. Altshuler, “Summary of the DARPA Background Clutter Experiment,” Proceedings EUROEM Conference 1998, Tel Aviv, Israel; George, V. and Altshuler, T. W., “Summary of the DARPA Background Clutter Experiment,” Proceedings FUZZ-IEEE 1998.
7. P. Gader, J. Keller, and H. Liu, “Landmine Detection Using Fuzzy Clustering in DARPA Backgrounds Data Collected with the Geo-Centers Ground-Penetrating Radar,” in *Detection and Remediation Technologies for Mines and Minelike Targets III*, A. C. Dubey,

J. F. Harvey, and J. Broach, eds., Seattle: International Society for Optical Engineering, 1998, p. 1139.