A Multiple Migration and Stacking Algorithm Designed for Land Mine Detection

John Schofield, David Daniels, and Paul Hammerton

Abstract—This paper describes a modification to a standard migration algorithm for land mine detection with a ground-penetrating radar (GPR) system. High directivity from the antenna requires a significantly large aperture in relation to the operating wavelength, but at the frequencies of operation of GPR, this would result in a large and impractical antenna. For operator convenience, most GPR antennas are small and exhibit low directivity and a wide beamwidth. This causes the GPR image to bear little resemblance to the actual target scattering centers. Migration algorithms attempt to reduce this effect by focusing the scattered energy from the source reflector and consequentially improve the target detection rate. However, problems occur due to the varying operational conditions, which result in the migration algorithm requiring vastly different calibration parameters. In order to combat this effect, this migration scheme stacks multiple versions of the same migrated data with different velocity values, whereas some other migration schemes only use a single velocity value.

Index Terms—Land mine detection, reverse time migration.

I. INTRODUCTION

OUT of the many issues for countries recovering from conflict, one of the key problems is that of land mines. Land mines have been used in many conflicts over the years, and this has resulted in between 45 and 100 million unexploded land mines in the ground today spread across over 75 countries [5]. As a result of this, considerable research has gone into trying to find a cost-effective way of detecting land mines without excessive false alarms. Ground-penetrating radar (GPR) is an electromagnetic technique for finding the location of objects or interfaces buried beneath the Earth's surface [6], [8]. GPR works by sending electromagnetic waves into the ground, as shown in Fig. 1, and detecting the reflected signal caused by any impedance discontinuities in the ground, which can then be investigated if required.

The GPR image is highly dependent on the propagation characteristics of the ground as well as the antenna characteristics. Consider the idealized case where we are detecting a point source; here, we have the point spread function of the target being convolved with the antenna beam function, and this spreads the received signal in time and space, causing the
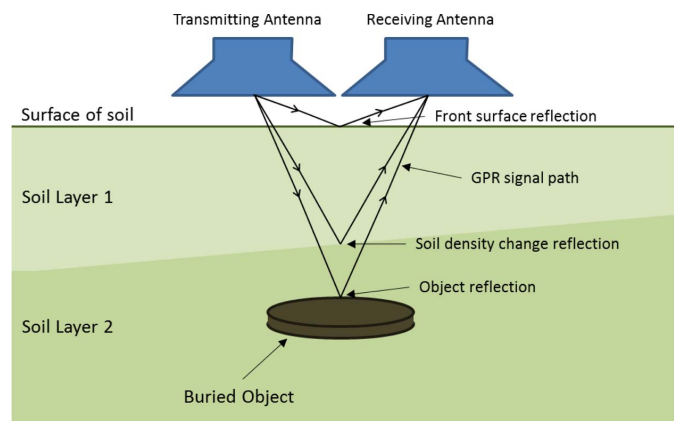
Fig. 1. Standard GPR arrangement demonstrating how to detect buried objects.

well-known diffraction hyperbola [10]. The problem with this idealized case is that the wavelengths used by GPR devices can be half the diameter of an antitank (AT) mine, and as such, the AT mines do not always act like point spread functions. However, the geometry of AT mines means they still form hyperbolic scattering because of multiple scattering centers which, of course, are not the same as a single point source. Throughout this paper, we will be making the assumption that we are dealing with perfect hyperbolas, which is a reasonable assumption considering the scattering patterns caused by the AT mines used in the investigation.

This paper discusses a new technique to improve the target detection rates by using a multiple migration and stacking scheme. Here, we will converge the diffraction hyperbola back into something more closely resembling a point source for a range of different ground conditions.

A standard technique for finding a certain class of shape in an image is the Hough transform (HT). The HT [11] and generalized HT (GHT) [2] are basic feature extraction techniques used to find imperfect instances of a certain class of shapes. Typically, the GHT is used for detecting basic shapes such as lines or circles [12], although more complicated shapes can also be considered [17], such as hyperbolas. As such, the HT has been used in various papers to help detect these hyperbolas [1], [3], and techniques have been developed to take account of variable soil density which then affects the propagation velocity [15]. Another approach to improve the clarity in GPR signals is to use migration. However, most migration literature deals with seismic data [4], [16]. A popular migration technique for improving the detection rates of land mine detectors has been reverse time migration, which various investigators have used to

reduce the scattering effects [9], [13], [14], [19]. Ideally, such an algorithm should be used in real time with the land mine detector and work on a variety of ground conditions without the need for recalibration. However, the migration algorithms tend to be relatively complicated, which leads to problems implementing them directly into a land mine detector while the HT encounters problems when dealing with hyperbolas of varying shapes.

In this paper, a scheme is developed from observing reverse time migration algorithms and simplifying the method without losing any of its effectiveness, while also improving its robustness in dealing with varying ground conditions. This has been accomplished by creating a process which stacks multiple images with differing diffraction signatures. Typically, this has only been possible in either "common shot" or "common receiver" configurations [10], which requires either multiple transmitting antennas or multiple receiving antennas. These configurations are impractical for the purposes of demining; hence, why the following scheme has been developed for zero-offset data.

II. Multiple Migration and Stacking Algorithm

A. Introduction and Aims

In this section, we formally develop the algorithm used here to improve the target detection rates for land mine detection. Following this, we will test the algorithm on numerous data sets in different ground conditions and develop a target detection algorithm to quantify the improvement offered by the migration algorithm.

For each B-Scan [7], we expect to find hyperbolas with unknown curvature in a region which also contains a significant amount of noise and clutter. This technique will detect hyperbolas of any curvature within a reasonable range and highlight them for the user.

As was mentioned earlier, this is accomplished by stacking various images with differing diffraction signatures in order to give us a final image which will show minimal effects of diffraction, hence providing a clearer image for the purposes of demining. These images are created using a basic migration scheme which acts iteratively on each A-Scan. This approach looks at potential targets and enhances their probability by adding adjacent samples from regions where we would expect scattering to occur. This is then performed multiple times for the entire B-Scan with a range of calibration parameters to generate a range of images with differing diffraction signatures, which are then stacked together to give the final image.

The end result of the migration is similar to what could be achieved using an HT designed to detect hyperbolas of a specific curvature.

B. Mathematical Modeling

Due to the wide beamwidth, signals are received from a variety of locations other than directly above a target. We assume a beamwidth of θ (the GPR device which produced the scans used in this paper had a value of $\theta \approx 45°$). It can be seen that, by looking at a single A-Scan, there are a wide range

of possible locations for a signal response to originate from. Typically, multiple signal responses exist from the same target in a B-Scan. We can correlate these signal responses in order to find the target's actual location.

First of all, we need to develop a forward scattering model. Consider a sequence of A-Scans corresponding to the horizontal positions x_I, where $I = 1, \ldots, N$. For each of these A-Scans, a response of amplitude $R_{I,J}$ is recorded corresponding to a propagation distance z_J, where $J = 1, \ldots, M$. We consider a buried object at (x_a, z_b), and we expect to receive a signal response at (x_i, z_j), where z_j is the distance from the detector $(x_i, 0)$ to the buried object (x_a, z_b), provided the beamwidth is wide enough. From here, we wish to converge the set of points in a B-Scan, which belong to a single object, back into a pointlike object, which will allow us to see the target's location. In order to do this, we need to find the set of points $S_{I,J} = \{(x_i, z_j)\}$ satisfying

$$(x_i - x_I)^2 + z_j^2 = z_J^2 \pm \delta \qquad \left| \frac{(x_i - x_I)}{z_J} \right| \leq \beta \qquad (1)$$

where $\beta = \tan \theta$ and δ is the smallest correction term to ensure that $i \in \{1, \ldots, N\}$ and $j \in \{1, \ldots, M\}$. We need to find $S_{I,J}$ for all I and J (i.e., the whole B-Scan). Now, given a new point (x_p, z_q), we create the set $T_{p,q}$, which is defined as

$$T_{p,q} = \{(I, J) \text{ such that } (x_p, z_q) \in S_{I,J}\}. \qquad (2)$$

Finally, we define the migrated point $\hat{R}_{p,q}$ as

$$\hat{R}_{p,q} = \sum_{I,J \in T_{p,q}} R_{I,J} \qquad (3)$$

which is the sum of all points where a signal response could be generated from an object at (x_p, z_q). Now, the migration map $\hat{R}_{I,J}$ for $I = 1, \ldots, N$ and $J = 1, \ldots, M$ can be generated for the entire B-Scan, which will allow better object definition for the purpose of finding targets.

C. Real Data Modeling

The formulation of the previous section is based upon horizontal and vertical distances, whereas the data consist of a sequence of A-Scans. We make the assumption that the propagation velocity through the ground is constant; hence, we can assume that there is a linear relation between time t_J of the signal response and the depth of the target z_J. This is not strictly true due to changes in soil type as a function of depth and lateral distance, as well as the varying water content of the ground, again as a function of depth and lateral distance, and objects in the ground. These all affect the propagation velocity; however, it is still an acceptable assumption for the local region associated with a hyperbola. This gives $z_J = \hat{c} t_J / 2$, where $\hat{c}$ is the propagation velocity of the radar waves through the ground. Also, GPR data are typically gathered using a handheld detector, which take a certain number of readings per second, and as such, we do not know the distance between adjacent A-Scans.

In order to use the formulation of migration from the previous section, we need to use expressions for the horizontal and

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

SCHOFIELD *et al.*: MULTIPLE MIGRATION AND STACKING ALGORITHM DESIGNED FOR LAND MINE DETECTION   3
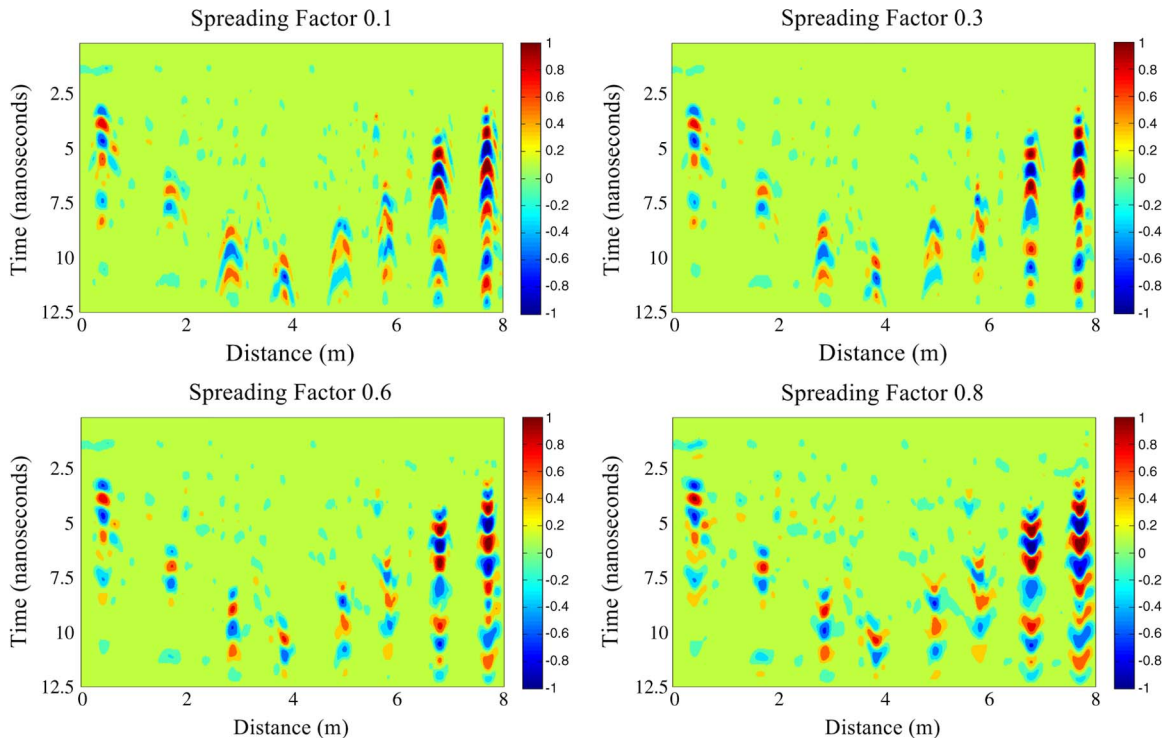

Fig. 2. Four images of the same migrated B-Scan showing the importance of using the correct spreading factor. Note that all images have been normalized between −1 and 1 and standard preprocessing techniques have been performed.

vertical distances. We have that $z_j = j\hat{c}\Delta t/2$ and $x_i = i\Delta x$, where Δx and Δt are the step sizes in x and t, respectively, i is the ith A-Scan, and j is the jth signal response from an A-Scan. Note that Δx is taken as a constant, corresponding to a uniform velocity of the detector, but the algorithm allows for fluctuations, as is described later. Because of this change, we need to express (1) in terms of the new notation, and now, we need to find the set of points $S_{I,J}^{\alpha} = \{(x_i, z_j)\}$ satisfying

$$\frac{1}{\alpha^2}(i-I)^2 + j^2 = J^2 \pm \delta \qquad \left|\frac{(i-I)}{J}\right| \leq \alpha\beta \qquad (4)$$

where $\alpha = \hat{c}\Delta t/2\Delta x$ and will now be referred to as the spreading factor. Now, (2) remains the same, giving the set $T_{p,q}^{\alpha}$, defined as

$$T_{p,q}^{\alpha} = \left\{(I,J) \text{ such that } (x_p, z_q) \in S_{I,J}^{\alpha}\right\}. \qquad (5)$$

Again, we define the migrated point $\hat{R}_{p,q}^{\alpha}$ as

$$\hat{R}_{p,q}^{\alpha} = \sum_{I,J \in T_{p,q}^{\alpha}} R_{I,J}. \qquad (6)$$

This leads to the migration map $\hat{R}_{I,J}^{\alpha}$ for the whole B-Scan with spreading factor α.

In order for this method to work, we need to find the value for $\alpha = \hat{c}\Delta t/2\Delta x$. In this paper, the B-Scans have been generated using a MINEHOUND land mine detector developed by Cobham Technical Services with a bandwidth of approximately 500 MHz to 2500 MHz. With this GPR device, Δt (the sampling time) was 50 ps, and the propagation velocity $\hat{c}$ is given by $\hat{c} = c/\sqrt{\epsilon}$, where c is the speed of light in a vacuum and ϵ is the ground's relative dielectric constant. Consider the case where $\epsilon \approx 6.25$ as a test case to inform us of the approximate size

of this parameter; this results in $\hat{c} \approx 1.2 \times 10^8$ ms^{-1}. The step sizes in x, Δx, can be calculated by considering the data acquisition frequency (DAF). The MINEHOUND detector has a DAF of 62.5 Hz. Δx can be calculated as $\Delta x = $ Velocity DAF. If we assume the average velocity of the detector to be 0.5 ms^{-1}, this gives $\Delta x = 0.008$ m. Combining all these together gives $\alpha \approx 0.375$.

However, adopting a fixed value of α leads to problems due to the fact that both propagation velocity $\hat{c}$ and speed with which the GPR device is moved are subject to variations, which make any fixed estimate of α unreliable over the course of a B-Scan. The propagation velocity varies with changes in the ground type and the moisture content of the ground. The speed with which the GPR device is moved depends on the operator, and while you would expect an approximately constant velocity from the same operator, there will be some fluctuations. This means that any change to the curvature of the hyperbolas we are trying to resolve can be countered by a change to α. Typically, in a single B-Scan, multiple spreading factors will be needed to resolve each target back into something resembling a point source. This can be seen in Fig. 2, where we attempt to migrate eight AT mines, presented in the layout described in Fig. 4, back into pointlike sources using various spreading factors, and as can be seen, no single spreading factor migrates all targets back into pointlike sources. In Fig. 2, the four mines on the left are metal Tenkovska Mina Metalna-1 (TMM-1) AT land mines, whereas the four on the right are minimum metal Tenkovska Mina Antimagnetna-2 (TMA-2) AT land mines. The multiple reflections in the minimum metal mines can be explained due to the GPR waves entering the mine and reflecting internally inside the mine before returning to surface, hence causing the resonance effect. This is not possible with the metal mines as
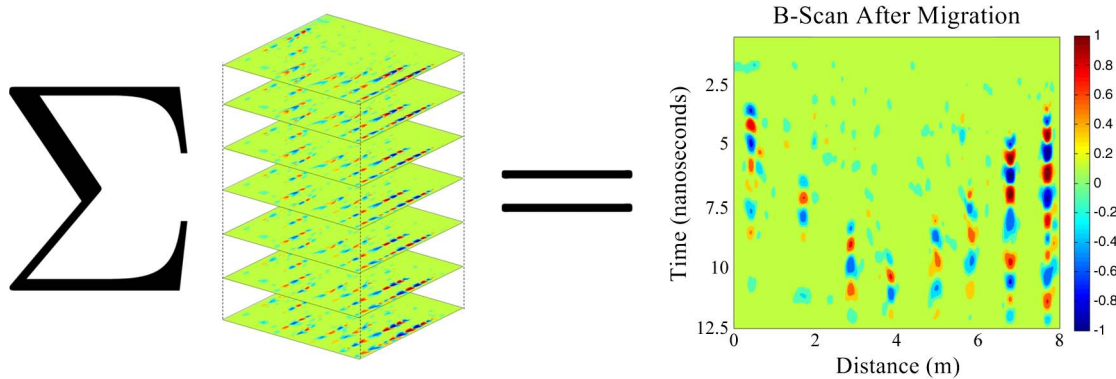
Fig. 3.   Example of how overlaying the images from Fig. 2 with various spreading factors works. Note how the diffraction effects visible in Fig. 2 can no longer be seen.
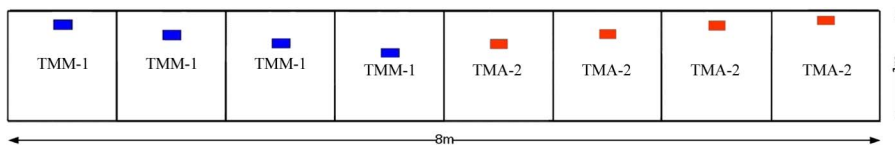

Fig. 4.   Layout of the mines from left to right as shown in Figs. 2, 3, and 5.

they act as a Faraday cage and, as such, allow no energy inside the mine.

As can be seen from Fig. 2, the most effective spreading factor seems to be $\alpha = 0.6$. However, for some targets, a spreading factor of $\alpha = 0.3$ seems more appropriate at resolving targets back to pointlike sources (possibly due to local variations in the water content of the ground). Rather than be faced with the difficult problem of finding the optimal spreading factor, the decision was made to superimpose all images together from Fig. 2. This results in an image where we have a strong response from the targets' locations due to the fact that there is a signal response at this location in all of the images. Conversely, any signal responses caused by the hyperbolic scattering or an incorrect spreading factor will only appear in one or two images and, as such, will not show up strongly in the overlaid image, as can be seen in Fig. 3. Mathematically, we can represent this as

$$\tilde{R}_{I,J} = \sum_{\gamma=1}^{r} \hat{R}_{I,J}^{\alpha} \Big|_{\alpha=\frac{\gamma}{10}} \tag{7}$$

where $r = 8$ has been used corresponding to spreading factors of $\alpha = 0.1, 0.2, \ldots, 0.8$. However, this may need to be changed according to the type of equipment, the velocity of the detector, and the ground conditions. However, a 33% change in the spreading factor (from changes of velocity or ground conditions) makes a minimal difference to the final image. It may be possible to implement an adaptive algorithm that alters r based on the data; however, this has not been considered in this paper.

III. STATISTICAL TESTING

A. Introduction

In order to impartially test how well the multiple migration and stacking algorithm performs, this section will show the results of testing this method on numerous data sets in order to obtain the probability of detection and the false alarm rate. The test site where these measurements were taken was prepared by

burying replica mines in various ground types. These are genuine TMM-1 and TMA-2 mines with the explosives removed and replaced with an inert material with similar dielectric properties as the explosive. These mines were then laid in ballast, a mixture of sand and shingle, and crushed concrete, with each ground type offering slightly different conditions such as subsurface anomalies and propagation characteristics. It is of worth noting that rocks and boulders do not pose much of a problem causing false alarms as they have much smaller radar scattering cross sections and vastly different internal structures. The results have been gathered with these mines laid in the configuration described in Fig. 4 in each of the three ground types to test how well the algorithm copes with varying conditions. These results will then be tested by a target detection algorithm based upon the total energy in each A-Scan.

B. Norm of Each A-Scan

A simple method of target detection is to consider the norm squared for each A-Scan, which is denoted by $\|\mathbf{R}_i\|^2$, where $\mathbf{R}_i = \{R_{i,1}, \ldots, R_{i,M}\}$ and is given by

$$\|\mathbf{R}_i\|^2 = \sum_{j=1}^{M} (R_{i,j})^2. \tag{8}$$

This can be quickly calculated for the entire B-Scan, and from there, we assume that, the higher the norm squared, the more likely it is that a target is present at that location. Upon inspecting the norm squared, it became clear that low-level background noise was providing mild interference. In order to reduce this background noise, thresholding was introduced to the B-Scan based upon a test section of ground with no mines present, which was performed on both original and migrated data sets. However, it is worth mentioning that this step is optional and has been mainly introduced in order to improve the clarity of the results. The results from applying the thresholding first and then calculating the energy in each A-Scan are shown

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

SCHOFIELD *et al.*: MULTIPLE MIGRATION AND STACKING ALGORITHM DESIGNED FOR LAND MINE DETECTION 5
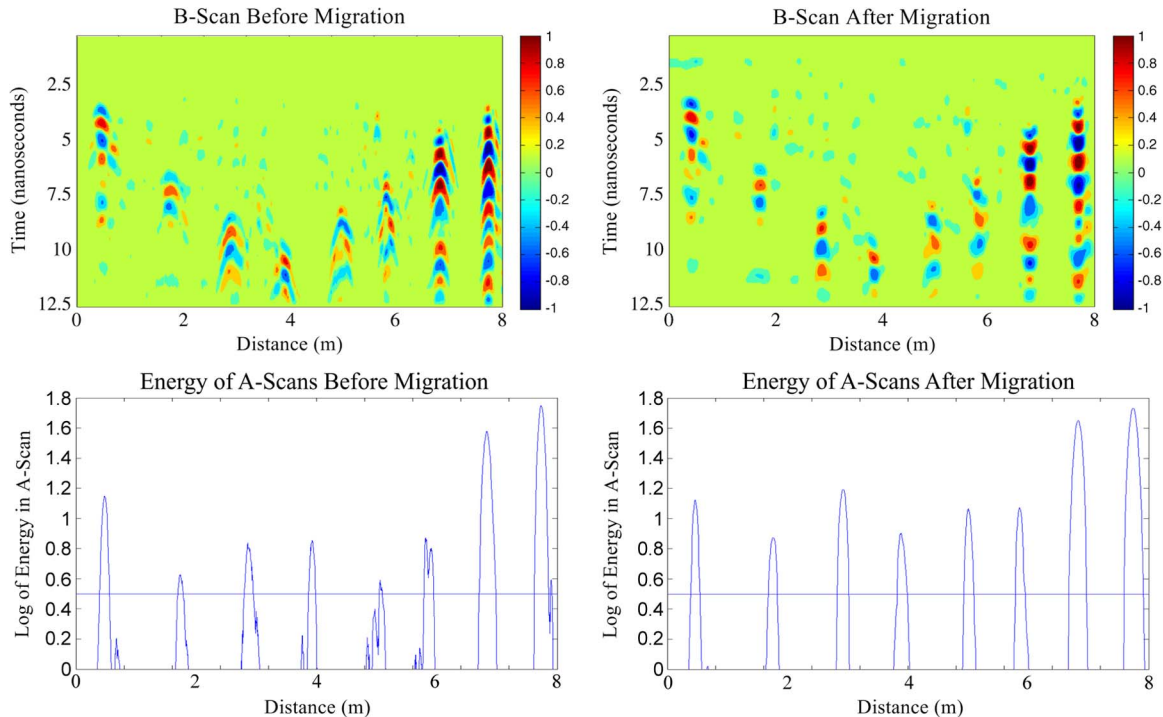

Fig. 5. Energy of each A-Scan in a B-Scan showing the effect of performing migration with an energy cutoff.

in Fig. 5. As can be seen from Fig. 5, the effects of migration are to increase the energy from a target and reduce the distance over which the energy is spread over, both consistent with what was expected from the algorithm. A further effect is that the migration algorithm will only pick out and highlight objects that are hyperbola shaped; any artifacts without this key shape are reduced below the level of the background noise, which leads to a reduction in the false alarm rate.

With the energy of each A-Scan extracted, we can now quantify how much of an improvement performing migration offers in terms of locating potential targets. This is accomplished by setting a cutoff value where any peaks over this value are classified as a target. The results of target detection are shown in Fig. 5, where a cutoff value has been included to better show the effects of the migration. By varying this cutoff value, we can generate receiver operator characteristic (ROC) curves to better see the effect of migration. Note that an arbitrary cutoff value has been selected as all data sets have been normalized. If the peak of energy is within a window predefined to be the location of a target from knowledge of the mine locations, then it is regarded as a successful detection. However, if the peak is outside this predefined window, then it is regarded as a false alarm.

C. Results

Having now developed a target detection algorithm, the multiple migration and stacking algorithm can now be tested on numerous AT mines. In order to gather the results, three test sites were prepared in ballast, a mixture of soil and shingle, and MOT. Each test site contained eight AT mines laid in the configuration described in Fig. 4. The data set consisted of 66 B-Scans following a predetermined path directly above the mines with 22 B-Scans from each ground type, resulting in a total of 528 AT mines. All results were obtained by the

same operator, and half of the data were collected a day after a substantial rainfall, while the other half was collected after a prolonged dry spell. The target detection algorithm was then used in order to measure the effect of the multiple migration and stacking algorithm (against not using the algorithm for the same data sets).

As mentioned in the previous section, we vary the cutoff of how much energy is needed in an A-Scan for it to be termed "a target" to generate ROC curves. This allows us to see how effective the algorithm is in improving target detection rates. As can be seen in Fig. 6, performing migration offers a clear improvement in target detection. In Fig. 6, the false alarm rate is defined as the ratio of false alarms to genuine targets. Hence, a false alarm rate of one means that there is one false alarm for every successful detection. In practice, the false alarm rate will be much higher than that on the test site; however, the results in Fig. 6 show a clear improvement upon using the multiple migration and stacking algorithm.

IV. CONCLUSION

In this paper, a multiple migration and stacking algorithm has been developed to increase the probability of detection of AT land mines at a false alarm rate of 0.01. The algorithm provided a probability of detection of 0.85 compared with 0.6 before processing, hence demonstrating a substantial improvement. The aim of this technique was to eliminate the hyperbolic scattering present when detecting land mines in GPR images for a range of different ground conditions. The key feature of this new approach is the stacking technique which reduces the effect of fluctuations of the ground conditions and in the movement of the detector operator. This method was tested on real data to observe the effectiveness of this technique at improving the target detection rate.
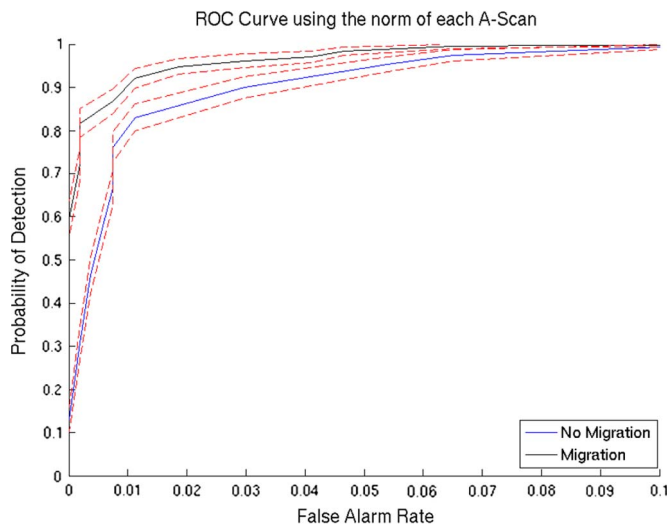
Fig. 6. (Dashed red lines) ROC curves based on the norm of each A-Scan highlighting the effects of migration with 95% error bounds generated using the Wilson score interval [18].

Further research could include testing this method on more soil types and types of mine across various sites to improve the reliability of the algorithm. Also, it would be of interest to introduce clutter in the ground to see how well the algorithm will distinguish between clutter and the mines. This method could also easily be extended to work in three dimensions to accommodate use for multiple wide swathe detectors which typically generate 3-D images.

Acknowledgment

References

[1] W. Al-Nuaimy, Y. Huang, M. Nakhkash, M. T. C. Fang, V. T. Nguyen, and A. Eriksen, "Automatic detection of buried utilities and solid objects with GPR using neural networks and pattern recognition," *J. Appl. Geophys.*, vol. 43, no. 2–4, pp. 157–165, Mar. 2000.

[2] D. H. Ballard, "Generalizing the Hough transform to detect arbitrary shapes," *Pattern Recognit.*, vol. 13, no. 2, pp. 111–122, 1981.

[3] L. Capineri, P. Grande, and J. A. G. Temple, "Advanced image-processing technique for real-time interpretation of ground-penetrating radar images," *Int. J. Imaging Syst. Technol.*, vol. 9, no. 1, pp. 51–59, 1998.

[4] J. F. Claerbout, *Earth Soundings Analysis: Processing Versus Inversion*, vol. 6. Cambridge, MA, USA: Blackwell Scientific Publications, 1992.

[5] M. Croll, *The History of Landmines*. Barnsley, U.K.: Leo Cooper, 1998.

[6] D. J. Daniels, *Surface-Penetrating Radar—IEE Radar, Sonar, Navigation and Avionics Series 6*. London, U.K.: The Institute of Electrical Engineers, 1996.

[7] D. J. Daniels, *Ground Penetrating Radar*. London, U.K.: Institution of Engineering and Technology, 2005, pp. 247–352.

[8] D. J. Daniels, D. J. Gunton, and H. F. Scott, "Introduction to subsurface radar," *Proc. Inst. Elect. Eng. F—Radar Signal Process.*, vol. 135, pp. 278–320, Aug. 1988.

[9] E. Fisher, "Examples of reverse-time migration of single-channel, ground-penetrating radar profiles," *Geophysics*, vol. 57, no. 4, pp. 577–586, Apr. 1992.

[10] J. Gazdag and P. Sguazzero, "Migration of seismic data," *Proc. IEEE*, vol. 72, no. 10, pp. 1302–1315, Oct. 1984.

[11] P. V. C. Hough, "Method and means for recognizing complex patterns," U.S. Patent 3 069 654, Dec. 18, 1962.

[12] D. Ioannou, W. Huda, and A. F. Laine, "Circle recognition through a 2D Hough transform and radius histogramming," *Image Vis. Comput.*, vol. 17, no. 1, pp. 15–26, Jan. 1999.

[13] C. J. Leuschen and R. G. Plumb, "A matched filter based reverse time migration algorithm for ground penetrating radar data," *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 5, pp. 929–936, May 2001.

[14] M. Neyrat, C. Guiffaut, and A. Reineix, "Reverse time migration algorithm for detection of buried objects in time domain," in *Proc. IEEE Antennas Propag. Soc. Int. Symp.*, 2008, pp. 1–4.

[15] G. Pochanin, "Some advances in UWB GPR," in *Unexploded Ordnance Detection and Mitigation*. New York, NY, USA: Springer-Verlag, 2009, pp. 223–233.

[16] R. H. Stolt and A. K. Benson, *Seismic Migration: Theory and Practice*, vol. 5. London, U.K.: Geophysical Press, 1986.

[17] H. Wechsler and J. Sklansky, "Automatic detection of rib contours in chest radiographs," Ph.D. dissertation, Univ. California, Irvine, CA, USA, 1975.

[18] E. B. Wilson, "Probable inference, the law of succession, statistical inference," *J. Amer. Stat. Assoc.*, vol. 22, no. 158, pp. 209–212, Jun. 1927.

[19] Z. L. Han, B. Pei, J. X. Zhou, and H. Wang, "Terrain correction and migration of GPR profile fulfilled simultaneously using reverse-time migration," *J. Jilin Univ.*, vol. 34, no. 3, pp. 459–463, Jul. 2004.

John Schofield was born in Yorkshire, U.K., in 1988. He received the MMath degree from the University of East Anglia, Norwich, U.K., in 2010, where he is currently working toward the Ph.D. degree in mathematics.

His research interests include nonlinear acoustics, ground-penetrating radar, and airline network and timetable optimization.

David Daniels received the B.Sc. and M.Sc. degrees from the University of Manchester, Manchester, U.K., in 1970, and the University of Northumbria, Newcaste upon Tyne, U.K., in 1984, respectively.

He is currently a Managing Director with Short Range Radar Systems Limited, West Sussex, U.K. He was previously with Cobham Technical Services, Leatherhead, U.K. He is an internationally recognized expert on ground penetrating and ultrawide-band radar and has published over 100 technical papers and award winning books on ground-penetrating radar as well as contributed to a number of seminal reference books in the field, including a contribution to *Skolnik's Radar Handbook*, third edition. He also authored *EM Detection of Concealed Targets* (Wiley, 2009), which describes techniques ranging from nuclear quadrupole resonance to submillimeter wave radar. He was appointed a Commander of the Order of the British Empire in the 2012 Queen's Birthday Honors List.

Paul Hammerton was born in Derby, U.K., in 1965. He received the B.Sc. and Ph.D. degrees from the University of Cambridge, Cambridge, U.K. in 1986 and 1990, respectively. His Ph.D. concerned the propagation of nonlinear acoustic waves through the atmosphere, under the supervision of Prof. David Crighton, Fellow of the Royal Society.

From 1990 to 1993, he worked on boundary-layer receptivity with Prof. Ed Kerschen in the Department of Aerospace and Mechanical Engineering, The University of Arizona, Tucson, AZ, USA, before returning to the University of Cambridge as a Fellow of Christ's College. Since 1995, he has been a Faculty Member with the School of Mathematics, University of East Anglia, Norwich, U.K. His recent research interests have included ultrasound scattering by colloids, boundary-layer transition, and solitary waves on free surfaces.