

LANDMINE DETECTION BY MEANS OF GROUND PENETRATING RADAR: A MODEL-BASED APPROACH

P.A. van Vuuren *

* School of Electrical, Electronic and Computer Engineering, North-West University, Private Bag X6001, Potchefstroom, 2520, South-Africa. E-mail: pieter.vanvuuren@nwu.ac.za

Abstract: The presence of landmines poses a worldwide humanitarian problem. Often, these mines are difficult to detect with metal detectors. Ground penetrating radar (GPR) is a promising technology for the detection of landmines with low metal content. Automatic landmine detection typically consists of two steps, namely preprocessing (or clutter removal) and classification. In this paper the clutter removal algorithm consists of a nonlinear frequency domain filter followed by principal component based filtering. Principal component analysis is performed in the frequency domain to build a background model for the clutter. The latter model is removed from the observed data in the log-frequency domain in order to preserve the phase component of the spectrum. Finally, the data is normalized and transformed to the time domain. The results presented in this paper show a marked improvement in the ability to remove general background clutter. Classification is performed on the basis of the prediction performance of neural network time-series models of the various classes of GPR responses. The classification system can correctly identify the position of metal anti-tank (AT) mines. It can also recognize specific examples of low metal AT and (anti-personnel) AP mines, but does have a low generalization ability for such mines.

Key words: Landmine detection, ground penetrating radar, neural networks, principal component analysis, clutter removal

1. INTRODUCTION

Landmines and other explosive remnants of war pose a debilitating threat to innocent civilians long after hostilities have ceased. World-wide, 4191 people (of which 75 % were civilians) were killed or injured during 2010 by landmines and other explosive remnants of war [1]. Each one of these casualties is one too many. Furthermore, their mere suspected presence in certain regions prevents agricultural activities, inhibit people's freedom of movement and hampers reconstruction and development of war-torn societies [1]. Clearance of contaminated areas is a costly exercise. The effectiveness of mine clearance operations obviously hinge on accurate detection of landmines.

Landmine detection is a vexing problem due to the low metal content of some anti-tank and anti-personnel mines. Furthermore, contaminated areas often are littered with other general metallic debris, causing a high false alarm rate in metal detectors [2]. Metal detectors are therefore unsuited to detect the broad range of landmines that can be encountered in practice.

Ground penetrating radar (GPR) is a proven technology for subsoil research [3]. The radar reflections received by a GPR system are a function of all subsoil dielectric changes. As such, GPR systems can in principle detect a wider range of mines compared to metal detectors. Whether this is true in practice remains to be seen.

GPR scans are typically performed from a moving platform such as an armoured vehicle. In order to cover as large an area of ground as possible, GPR

antennas are often arranged in arrays of identical antennas. GPR antenna arrays make it possible to reconstruct a three-dimensional image of the subsurface features. This 3D image is commonly known as a C-scan. The axis along which the GPR antennas are arranged is called the *cross-track* direction, since the entire array is moving in the *down-track* direction. The radar reflections obtained by a single GPR antenna at a specific (x,y) position form a 1-dimensional signal, known as an A-scan. A B-scan refers to a set of A-scans in either the down-track (x) direction or in the cross-track (y) direction.

In this paper a stepped-frequency, continuous-wave (SFCW) GPR antenna system is used to obtain a GPR C-scan. At each (x,y) position a sequence of GPR sinewaves is generated with a constant amplitude and stepwise increase in frequency. In this manner the A-scan at position (x,y) can (after some processing) be interpreted as the averaged frequency response of the ground to a wide band of radar frequencies (100 MHz to 1.8 GHz).

The term commonly used in the GPR literature for various phenomena that obscure the presence of valid targets (landmines) is *clutter*. Clutter increases the false alarm rate of a landmine detection system [4]. This is because clutter often dominates the observed data (especially in the frequency domain) [5].

Some researchers define clutter as all phenomena in an A-scan that remain constant over a large section of the corresponding down-track B-scan [6]. Such phenomena include antenna effects [7], certain instances of layers in the soil with differing electromagnetic properties [8], as well as the ubiquitous air-surface reflection (also referred

to as *ground bounce*) [4]. Other forms of distracting phenomena are however more transient in nature (e.g. rocks and roots [8] as well as soil roughness resulting in diffuse scattering [7]).

In landmine detection via GPR the preprocessing step is therefore primarily occupied with the removal of as much clutter as possible to facilitate successful landmine detection. Clutter removal is then followed by the next step which typically involves feature extraction and classification of feature vectors.

The data obtained by a SFCW GPR array can be interpreted as the averaged frequency response of the soil at a particular location. In the time domain this corresponds approximately to the impulse response of the soil. In this paper, the first in a two-part series, clutter removal is performed in the frequency domain. In the second paper (Landmine detection by means of ground penetrating radar: a rule-based approach) clutter removal is performed in the time domain. Classification is performed in the time domain in both papers. A model-based approach is followed to perform classification of GPR data in this paper, whereas a rule-based classifier is presented in the second paper.

This paper contributes to both the clutter removal stage and the final classification stage of the GPR landmine detection problem. A novel algorithm is developed in section 2 for the removal of clutter from frequency domain GPR data. After the resulting A-scans are transformed from the frequency to the time domain, classification is performed in a manner novel to the GPR literature by means of time-series models derived for various classes of data. Artificial neural networks are used for these time-series models. Section 3 discusses the design of the time-series models. The performance of the clutter removal algorithm is presented in section 4, while the ability of the time-series based classifier to detect a variety of mines is the topic of section 5. The paper is concluded in section 6.

2. CLUTTER REMOVAL ALGORITHM

2.1 Previous work performed on clutter removal

The fundamental objective of clutter removal is to accentuate the radar reflections made by landmines while at the same time diminishing the scattering due to soil features. Numerous clutter removal techniques exist, but they can all be classified into the following three categories.

- *Time or range gating techniques.* The simplest response to ground bounce is to merely discard the relevant sections of an A-scan [5]. This is known as time gating or range gating. The computational cost of time gating is very low. Unfortunately, this technique only removes a portion of the reflections due to ground bounce, since ground bounce reflections also occur deeper underground

[9]. Furthermore this technique is only applicable to time-domain data. Lastly, time gating can't be used with mines that are buried shallowly [9].

- *Classification-based techniques.* This is a loose collection of techniques that avoid explicit clutter removal in a variety of ways. One example entails extracting features that are insensitive to the presence of clutter [8]. A related approach involves performing statistical hypothesis testing on features based on linear prediction models of the data [10]. Such techniques highlight the blurry boundary between clutter removal and prescreening algorithms [11].
- *Background removal techniques.* Most clutter removal techniques essentially involve modelling the background clutter and subtracting it from the raw GPR data. This is most commonly accomplished one B-scan at a time (either in the cross-track or down-track direction). Models for the background clutter range from simplistic (e.g. the mean value at each depth in the down-track direction [6]) to first principles models based on the various physical processes responsible for clutter formation [7]. In fact, any technique that explicitly models the background clutter resorts under the banner of background removal techniques, whether via black-box system identification [9] or principal component analysis of time-series data [6].

Of the three above mentioned approaches, background removal techniques represent the best balance between accuracy and simplicity and is therefore further pursued in this paper.

Before delving into the details of the improved clutter removal algorithm we'll first pause a moment at a few fundamental assumptions concerning the process by which GPR reflections occur. GPR reflections occur when the electromagnetic properties (predominantly the permittivity) of the medium through which the radar signal is travelling change abruptly. These changes in permittivity occur at the boundaries of different objects (e.g. pebbles, roots, tunnels and mines) as well as at the boundaries of different soil layers. Some of these variations (specifically changes in the soil composition) can occur quite gradually. Most of the variations can only be regarded as random events and can't be predicted in advance. In fact, it is virtually impossible to construct an accurate and practically useful first principles model for the clutter in GPR data [12]. The only practical solution for building a model of the background scattering is therefore by means of adaptive numerical methods.

A moving average (or median) filter is a prime contender for a model of the background scattering, since it depends on minimal assumptions regarding the data and is inherently adaptive [6]. Typically, a background model is obtained by implementing a separate moving average filter at each depth index. The drawback of such an approach is however that the time axis of the A-scans is only an

inaccurate estimation of depth. This is due to the fact that the propagation velocity of the radar signal is dependent on the dielectric properties of the soil, which in turn is variable [13], [12].

The safest option is therefore to base a background model on an entire B-scan, either in the down-track or the cross-track direction. An entire B-scan can be regarded as an image (i.e. a vertical slice through the ground). Interpreting GPR data as images opens up the possibility of modelling the background clutter by means of principal components [14]. Principal component based clutter removal has been applied successfully in both the time- and frequency domain. The calculation of the principal components can also be performed recursively rendering an adaptive clutter removal algorithm [6].

2.2 Algorithm overview

The clutter removal process presented in this paper consists of a nonlinear frequency domain filter followed by principal component analysis also conducted in the frequency domain. The aim of the principal component analysis is to build a model of the background clutter. Clutter removal then entails that the background model is subtracted from the observed GPR spectra. Finally, the data is normalized prior to being transformed to the time domain for further processing by feature extraction algorithms. The entire procedure is as follows:

1. Removal of narrow-band disturbances via nonlinear filtering. During this step spurious peaks situated in the vicinity of 945 MHz are removed.
2. Form a background model by means of frequency domain principal component analysis and subtract it from the original spectrum in the logarithmic domain. Calculations are furthermore performed in the log-domain in order to prevent the phase component of the spectrum being amplified.
3. Normalize the resultant spectra.
4. Application of the inverse Fourier transform to obtain decluttered time-domain A-scans.

Each of the above mentioned steps will now be discussed in greater detail.

2.3 Spike removal via nonlinear filtering

The observed GPR spectra contain a number of distinctive peaks. These peaks have the following properties:

- They only occur at certain frequencies.
- The specific frequencies of these "spikes" exhibit a small variation from one A-scan to the next, but remain relatively constant throughout the entire C-scan.

- The specific "spike" frequencies are a function of the bandwidth of the stepped-frequency continuous wave signal of the antenna array.
- A large variation occurs in the amplitude of the peaks.
- A small variation occurs in the spectral width of the peaks.
- The peaks are visible in both the real and imaginary components of the spectrum.
- The peaks often (but not always) cause abrupt jumps in the phase component of the spectrum.

As an example, the magnitude and phase spectra of an A-scan are shown in figure 1. The peaks are clearly discernible in the magnitude spectrum and the corresponding jumps in the (unwrapped) phase spectrum are also quite obvious.




Figure 1: Frequency domain A-scan showing disturbance peaks

The semi-deterministic character of the spikes indicate that they are an artefact of the measurement system. Notch filters can be applied in the frequency domain to filter these peaks. This approach however unavoidably introduces its own artefacts to the spectrum.

An alternative approach is to employ a nonlinear filter to the observed spectra. The basic idea is to transform the data contained in a small window surrounding the spike in such a fashion to ensure that the resultant statistical distribution conforms to some desired distribution. More specifically, the "desired" distribution doesn't contain any outliers. This calls for a saturation-type nonlinearity to ensure that "reasonably-valued" data is left unhindered, while data points with large values are gradually limited to some upper (or lower) value. One option for this nonlinearity is the hyperbolic tangent function illustrated in figure 2.

The hyperbolic tangent filter is applied as follows:

1. Extract a small portion of the spectrum surrounding a spike. The width of this window obviously has




Figure 2: Hyperbolic tangent filter

an effect on the response of the filter. If it is too narrow, there won't be any outliers in the data in addition to the fact that the underlying distribution will be nonstationary due to the dominating effect of the spike. If, on the other hand, the window is too wide, then larger trends in the local spectrum will cause the data to also be nonstationary. In the end, the width of this window was chosen at 41 samples, since such a window only infrequently contained a nonstationary data sequence. (The stationarity of a data sequence can be tested by means of a run-test [15].)

2. Limit the amplitudes of the data in the window by means of the hyperbolic tangent function. This operation is performed separately on the real and imaginary components of the spectrum [6], since the data distributions of the latter spectra correspond more closely to the desired normal distribution than the magnitude and phase spectra.

In order to filter data at any general position and spread, it first has to be normalized to be commensurate with the input domain of the definition of the hyperbolic tangent function (which is approximately the unit domain). This normalization is done as follows:

$$y = \tanh\left(\frac{2x - 2\bar{x}}{\sigma}\right), \quad (1)$$

where x is the original data, $\bar{x}$ is the estimate of the position of the distribution (e.g. the median in the case of an outlier insensitive estimate) and σ is an estimate of the spread of the data (e.g. inter-quartile range).

To ensure that the filtered data has a distribution similar to the original data x , the filter output has to be modified as follows:

$$y = \frac{\sigma}{2} \tanh\left(\frac{2x - 2\bar{x}}{\sigma}\right) + \bar{x}. \quad (2)$$

As an example, the hyperbolic tangent filter was applied to the A-scan of figure 1. The spikes are completely removed by the hyperbolic tangent filter with no perceptible effect on the rest of the spectrum, as can be seen in figure 3.




Figure 3: Frequency domain A-scan after spike removal

The hyperbolic tangent filter has two limitations. First, the data within the window should be stationary. Secondly, the data in the window should have an approximately normal distribution (at least it shouldn't be skewed). Both of these limitations can be addressed by a judicious choice of the width of the filter window.

2.4 Modelling background clutter by means of principal component analysis

It is well-known that image compression can be performed by principal component analysis [14]. In this technique the main features of an image can be reconstructed from its dominant eigenvalues (also known as the image's principal components).

The background clutter over a certain window of data in the down-track direction can therefore be modelled by means of the first few principal components of the B-scan contained in the window. These principal components are obtained by means of singular value decomposition (which is applicable to rectangular matrices). If the B-scan at position y_0 contained in a certain window of down-track data is represented by matrix $\mathbf{X}$, the singular value decomposition can be expressed as:

$$\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^T, \quad (3)$$

where $\mathbf{U}$ is an orthogonal matrix containing the left singular vectors, $\mathbf{V}$ is an orthogonal matrix containing the right singular vectors and $\mathbf{S}$ is a diagonal matrix whose off-diagonal inscriptions are zero and whose main diagonal elements comprise the singular values of $\mathbf{X}$ arranged in descending order.

By attempting to reconstruct the original matrix $\mathbf{X}$ with only the largest few singular values and corresponding

singular vectors, a matrix is obtained that contains the main characteristics of the original matrix, without some of the detail. The resultant reconstructed matrix is in other words a model of the background in the window of GPR data.

Clutter removal is then performed by subtracting the background model from the observed data. This is equivalent to a high-pass filter and typically results in the edges and other sharp features in an image being accentuated [14].

2.5 Background removal in the log-frequency domain

Unfortunately, the subtraction operator inherent in the above mentioned high-pass filter inevitably causes an enormous amplification in the phase spectrum of the data, as can be seen in the following short analysis.

Say that the original value of a frequency-domain A-scan at a specific frequency (ω) is given by: $z_{\text{original}}(j\omega) = a + jb$ and that the value of the background model at the same position and frequency is: $z_{\text{background}}(j\omega) = c + jd$. Then the value of the decluttered A-scan at the frequency of interest is given by:

$$\begin{aligned} z_{\text{decluttered}}(j\omega) &= (a - c) + j(b - d) \\ &= \sqrt{(a - c)^2 + (b - d)^2} \dots \\ &\dots \exp \left[j \tan^{-1} \left(\frac{b - d}{a - c} \right) \right]. \end{aligned} \quad (4)$$

By focussing on the phase component of (4), it is clear that its phase will become very large if there is little difference between the original A-scan and the background model (which is most often the case).

As the above analysis shows, subtraction of *any background model* causes an unacceptable increase in the phase component. Clutter removal by means of subtraction of a background model is furthermore flawed since it inherently assumes an additive model for the process that gives rise to the observed data. In this approach (which is rather similar to a time-series model) the observed GPR signals ($S_o(j\omega)$) can be modelled as the sum of surface reflection effects ($S_s(j\omega)$), antenna effects ($S_a(j\omega)$), ground scattering ($S_g(j\omega)$) and possible target scattering ($S_t(j\omega)$) as follows [8]:

$$S_o(j\omega) = S_s(j\omega) + S_a(j\omega) + S_g(j\omega) + S_t(j\omega). \quad (5)$$

Improved clutter removal however requires an improved model for the process that gives rise to the observed data. A more accurate model makes use of the fact that the input signal generally is known. Under the (admittedly simplistic) assumption that the process of generating GPR reflections is linear and time-invariant, a measured GPR A-scan ($S_o(j\omega)$) can be modelled as the convolution of the input signal ($S_i(j\omega)$) with the various above mentioned processes [7]. In the frequency-domain

convolution becomes multiplication, which means that the observed spectrum can be modelled as follows:

$$\begin{aligned} S_o(j\omega) &= S_i(j\omega) \times G_s(j\omega) \times G_a(j\omega) \times G_g(j\omega) \dots \\ &\dots \times G_t(j\omega) \\ &= S_i(j\omega) \times G_b(j\omega) \times G_t(j\omega), \end{aligned} \quad (6)$$

where all of the effects that can be regarded as background clutter have been condensed into a background model $G_b(j\omega)$. (Note that in (6) the various processes are modelled as transfer functions ($G(j\omega)$), rather than mere observed signals ($S(j\omega)$) as in (5).)

A time-domain version of the target response can be obtained from (6) by means of a deconvolution filter [16], [17]. Equivalently, the spectrum of the target response can be obtained from (6) by division:

$$\begin{aligned} G_t(j\omega) &= \frac{S_o(j\omega)}{S_i(j\omega) \times G_b(j\omega)} \\ &= \frac{G_o(j\omega)}{G_b(j\omega)}, \end{aligned} \quad (7)$$

where $G_o(j\omega)$ represents the measured frequency response, which is essentially the measured spectrum in a stepped-frequency continuous-wave GPR system.

Practical implementation of (7) is best performed in the log-domain to avoid division by zero problems. By taking the natural logarithm on both sides of (7) and processing the magnitude and phase components of the spectra separately, the background model can be removed from the observed frequency response as follows:

$$\begin{aligned} \ln |G_t(j\omega)| + j\theta_t(\omega) &= \ln |G_o(j\omega)| - \ln |G_b(j\omega)| \dots \\ &\dots + j\theta_r(\omega) - j\theta_b(\omega). \end{aligned} \quad (8)$$

The final frequency response of the target can then be obtained by taking the exponential function on both sides of (8) to obtain (7).

2.6 Frequency-to-time conversion

As we'll see in section 3, classification of GPR signals is performed in the time domain. The inverse fast Fourier transform (IFFT) can be used to perform the transformation from the frequency domain to the time domain.

Before the IFFT can be applied, the decluttered spectra have to be normalized. Normalization is typically employed in pattern recognition systems to ensure that the numerical ranges of the various features are commensurate [18]. The simplest form of normalization entails subtracting the mean value and dividing by the standard deviation, as in (9).

$$G_n(j\omega_k) = \frac{G_t(j\omega_k) - \overline{G_t(j\omega)}}{\sigma}, \quad (9)$$

where ω_k refers to a specific frequency index in the spectrum, $\overline{G_t(j\omega)}$ is the mean value of the entire A-scan and σ is the standard deviation of the same spectrum.

The most important aspect of (9) is that it removes the bias level from the spectrum. This is important, since the bias component in a spectrum causes an impulse (Dirac delta) to appear in the equivalent time-domain signal.

Finally, the IFFT can be applied. This is performed according to the procedure outlined in [16]. First, the spectrum is multiplied with a Kaiser window (with a coefficient of $1.32 \times \pi$). Next, the signal is zero-padded until the total length is 4096 samples. Finally, the signal is mirrored and the IFFT applied. As an example, the equivalent time-domain A-scan of the scan shown in figures 1, 3 and 7 is shown in figure 4. This A-scan clearly shows the response of a metal anti-tank (AT) mine.




Figure 4: Time-domain A-scan after clutter removal

3. CLASSIFIER DESIGN

Feature extraction is arguably the most important step in any classification system. If the features can easily discriminate between the different classes (i.e. the classes form clearly separable clusters in the feature space), then classification is a mere formality. If however the features result in severely overlapping clusters, classification by even the most advanced classifiers is all but impossible.

Most of the GPR landmine classifiers in the literature are based on features that describe some or other discriminating property in the observed data. In such an approach the classifier isn't concerned with the fundamental processes that produce the observed signals, but only how to summarize the observed data in such a manner that the different classes can be detected.

Within the above mentioned feature extraction paradigm there are many examples of quite successful classification systems. Some researchers prefer an image-analysis approach in which features are extracted from horizontal B-scans that mimic those features a human operator would focus on in order to detect mines [4]. These features

try to capture the essence of the concentric circles and hyperbolas in B-scans that betray the presence of buried objects. Others favour using features distilled from the energy density spectra of A-scans [5], [8].

There are however researchers that attempt to model the responses of buried objects to GPR signals in a much more principled manner. As an example Roth *et al.* derived expressions for the impulse response of a low metal content mine under idealized soil conditions [19]. A model-based approach is also quite useful to obtain a background model for the observed B-scans with the aim of clutter removal [7]. Inevitably these authors have to make numerous simplifying assumptions in order to obtain tractable models.

As previously mentioned, in practice a GPR landmine detector is faced with a staggering array of unknowns including: the exact location of different soil layers; the frequency of occurrence and size of rocks and pebbles; the water content of the soil (which may change as a function of location and weather); the position, depth and orientation of mines; and lastly, the dielectric properties of sediment and rocks (which is influenced by a huge number of factors that are difficult or impossible to measure beforehand [17]).

In the light of the above mentioned uncertainties it should be clear that first-principles modelling of all possible landmine responses is unfeasible. The only practical alternative is so-called black-box system identification. One such approach is to model the observed GPR data by means of parameterized LTI models. Accurate models for the various classes of buried objects imply that the parameter vectors of the models are also unique and therefore suitable candidates for classification features. Unfortunately from extensive simulation experiments, classical parameterized models (ranging from ARX models to Wiener-Hammerstein models) exhibit both poor generalization and discrimination ability on GPR data as compared to artificial neural networks. (The *generalization ability* of a model refers to its ability to model other instances of the same system that weren't included in the data set upon which the model's parameter were estimated. On the other hand, a classifier's *discrimination ability* is its ability to distinguish between different classes.)

One of the few remaining options to obtain a black-box model for the observed time-domain A-scans is by means of artificial neural networks.

3.1 Detail design

A full exposition of the vast field of artificial neural networks is beyond the scope of this paper (consult e.g. [18]). Certain classes of artificial neural networks (especially multilayer perceptrons and radial basis function networks) can be viewed as essentially parameterized nonlinear function approximators. Consequently artificial neural networks (henceforth referred to as neural net-

works) shouldn't be regarded as fundamentally different to any other parameterized nonlinear model structure.

Radial basis function networks (RBFNs) possess a number of advantages over multilayer perceptrons (MLPs), namely:

- Training of RBFNs typically entails clustering (for the hidden layer neuron centre-points) and least-squares approximation (for the output nodes). This process is much faster than the backpropagation and nonlinear minimization required for MLP training.
- Due to the limited extent of each basisfunction in the hidden layer, RBFNs don't commit huge extrapolation errors as is the case with MLPs. (MLPs approximate a function by means of hyperplanes, while RBFNs do the same job with Gaussian kernels.)
- The mapping performed by a RBFN can be interpreted in terms of a basisfunction expansion, which is slightly more insightful than the essentially unknown mapping done by a MLP.

In this paper, RBFNs are used as function approximators. More specifically, a model is estimated for each class of landmine (e.g. metal mines, low metal AT and anti-personnel (AP)). The neural network is trained to form as a nonlinear autoregressive (AR) time-series model for the time-domain A-scan of the particular mine in question. In an AR-model the current sample value of the signal is modelled by a linear combination of previous sample values of the signal. Mathematically, an AR-model can be expressed as follows:

$$y(t) = \sum_{i=1}^N a_i y(t-i). \quad (10)$$

An AR-model can be duplicated by a neural network by arranging the inputs of the network in a so-called tapped-delay line. This entails that the previous sample of the signal is used as the first input to the network, the sample prior to the previous sample is used as the second network input and so on. The output of this network is then taken as the current sample value of the signal.

The neural network architecture described above functions as a single step-ahead predictor. In other words, the current and previous sample values of the signal are used to predict the value of the forthcoming sample. This can obviously be extended to a k -step ahead predictor in which the network predicts the value of the signal k -sample values into the future on the basis of the current and previous samples. If a separate neural network is trained to perform as a k -step ahead predictor for each class of landmine, then these models can be used to classify a new (unknown) A-scan. This is done by simply allocating the A-scan to the class whose neural network has the best prediction performance. In this manner both feature extraction (represented by the

Gaussian basisfunctions of a RBFN) and classification can be performed by the same neural network.

The training set of each class model comprised a 3×3 neighbourhood of A-scans in the vicinity of a known mine position. Training was performed on these nine A-scans until either a maximum number of neurons were reached or a minimum error was reached. (Training of the standard Matlab[®] RBFN entails automatic positioning of the basisfunctions. These basisfunctions are added incrementally at each iteration. This is followed by a least squares solution for the linear mapping between the current basisfunctions and the future sample presented at the output of the network.)

4. RESULTS: CLUTTER REMOVAL

4.1 Clutter removal performance measures

Successful clutter removal algorithms are characterised by the following attributes:

- Low computational cost. This is important with real-time implementation in mind.
- Quality. High fidelity reproduction of target signals with only a minimal presence of undesired clutter. Quite often in the literature the quality of a clutter removal technique is "measured" in terms of a subjective assessment of the resulting A- and B-scans. An objective comparison of the quality various clutter removal algorithms however requires a quantitative measure of the performance of the respective algorithms.

One quantitative measure of the quality of the clutter removal process is the signal-to-clutter ratio (SCR) defined in [20]. The SCR is a measure of the relative spread of GPR data at a specific depth in the cross-track direction and assumes that all antennas in the array are equal, that the soil properties in the cross-track direction is constant and that landmines only occur infrequently. Although originally defined for cross-track B-scans, it can also be applied to down-track B-scans.

The SCR is quite useful to determine how much clutter has been removed by a particular algorithm. Unfortunately this measure can only be used to study the effect of a clutter removal algorithm one A-scan at a time. It would be much more convenient if a measure could be defined that allows a clutter removal algorithm to be evaluated on the basis of its performance on an entire C-scan. One such measure is based on correlation.

The presence of ground-bounce and gradually changing soil characteristics results in a high degree of cross-correlation between adjacent A-scans in all directions. Background removal would therefore ideally cause a large decrease in the general degree of cross-correlation between neighbouring A-scans, with the

only exceptions occurring at the locations of buried objects (e.g. landmines). (The latter conclusion is under the assumption that the occurrence of landmines can be viewed as statistical outliers as discussed in the appendix to this paper.) The quality of a clutter removal algorithm can therefore be assessed by calculating the average degree of peak cross-correlation between each A-scan and its eight immediate neighbouring A-scans.

Figure 5 shows the average degree of correlation of neighbouring A-scans after application of a time-domain principal component based clutter removal algorithm (which merely entails application of the algorithm described in section 2.4 on time-domain GPR data). This figure should be interpreted as a top-down view of an area of ground which has been scanned with a ground-penetrating radar. Areas showing high degrees of correlation may be indicative of interesting subsurface objects or soil layers. Superimposed on this figure are the positions of known mines (indicated with black squares). Clearly, there is room for improvement since large areas of ground are shown to have highly correlated A-scans.

Clutter removal performed via the procedure described in section 2 has a marked improvement on the average degree of correlation in the data. As figure 6 shows, the positions of metal mines are now clearly visible (a significant improvement on figure 5).

Individual A-scans are also much more informative after the log-frequency domain clutter removal process. This can be clearly seen by comparing the A-scan obtained after spike filtering (figure 3) with the A-scan at the same position after background removal (figure 7). The existence of potentially informative subbands can be clearly seen in both the magnitude and phase spectra.




Figure 7: Frequency domain A-scan after clutter removal

4.2 Descriptive power of the decluttered data

One of the vexing problems of data preprocessing is that the process occurs rather blindly. Educated guesses are made and algorithms applied in the hope that useful

information can be distilled from the raw data. Only after the final classification stage has been completed, can it be seen if the various preceding stages were successful. It would therefore be of great advantage if a quantitative measure could be developed with which the descriptive and discriminative power of a dataset can be assessed prior to feature extraction.

The traditional approach to solve the above mentioned problem is to cluster the data and measure the ratio of the inter-cluster distance to the cluster size [18]. In this manner a measure can be found of the class separability in the data. Such a clustering-based approach is however limited to data consisting of a limited number of features and is definitely not suited to raw data.

A viable option is to implement a primitive correlation-detector (also known as a matched filter) [21]. Here the objective is to determine whether there are other A-scans in a C-scan that are similar to a given A-scan. The difference between this correlation-detector and the correlation-based measure of the previous section is that the entire C-scan is analysed for potential similarities with a given A-scan and not only the A-scan's immediate neighbourhood. In this manner the data can be analyzed at different stages of the system to determine the separability of the classes.

As an example, the correlation-detector was applied to the decluttered time-domain data to determine whether there are any other A-scans similar to an A-scan at the position of a high metal AT mine. The results are shown in figure 8. Clearly, the A-scans of the two examples of high metal AT mines are quite similar to each other. Another encouraging result from figure 8 is that the high metal AT A-scans are quite different from the other A-scans in the data. Such a result predicts that good classification results can be obtained on the data.

Figure 9 however paints a different picture altogether. This figure shows the peak cross-correlation with an A-scan at the position of an AP mine. If the scale of the colour bar is taken into consideration, figure 9 tells us that the A-scans of AP mines aren't correlated with any other A-scans in the C-scan (despite the fact that there are two known examples of M14 mines in the dataset). Similar analyses on the other low metal mines indicate that plastic mines are all but invisible in the current decluttered time-domain data. Similar problems have been reported elsewhere in the GPR literature [22].

5. RESULTS: MINE DETECTION

All of the results in this paper were obtained on a small dataset consisting of seven landmines buried in an artificially constructed sandpit. This set of mines consists of the following specific types: metal AT (2 examples), M14 (2 examples), TMA-3 (1 example), VSHCT (1 example), No. 8 (1 example). Over time a meerkat family also made the sandpit their home, with their burrows quite conspicuous in the GPR reflections. A GPR antenna array




Figure 5: Degree of average correlation after time-domain principal component based filtering




Figure 6: Degree of average correlation of magnitude spectra after log-frequency domain processing




Figure 8: Peak cross-correlation with an A-scan at the position of a high metal AT mine




Figure 9: Peak cross-correlation with an A-scan at the position of an AP mine

(model V1821) from 3D-radar was used to obtain the GPR data. This GPR antenna array is a general-purpose instrument designed for applications ranging from airport and road inspection to archeology and military purposes. Most of the GPR studies in the literature are obtained with impulse type GPR, whereas the instrument used in this research is a SFCW radar.

5.1 RBFN parameter optimization

Optimizing a RBFN to serve as an accurate time-series model involves that optimal values for the following parameters have to be found:

- Prediction horizon. This refers to the value of k in a k -step ahead predictor.
- Length of the tapped delay line. The length of the tapped delay line determines how many previous samples are taken into consideration to form a prediction of a future sample. (This parameter is the same as the order of a classical AR-model.)
- Width of the basis functions.
- Maximum network size. This indicates the maximum number of basis functions that can be placed during training.

These parameters were optimized with the eventual objective of the neural networks models in mind, namely to serve as classifiers of GPR A-scans. Consequently, the RBFN parameters were chosen in such a fashion that the resultant network exhibits a good discrimination ability coupled with a good generalization ability. Both of these properties can be inferred from the peak prediction performance of a network on a 3×3 neighbourhood of A-scans. These A-scans were chosen from different locations in the available data representing known classes of data.

The discrimination ability of a network is revealed by the difference in the prediction performance between data from the training class and data from other classes. The generalization ability, on the other hand, corresponds to the degree of similarity in the prediction performance of the network on data from similar classes (e.g. two different instances of metal AT mines).

Optimizing the prediction horizon: The peak prediction performance of the RBFN predictor is set out as a function of the prediction horizon in table 1. This table was obtained by fitting a nonlinear mapping between 15 previous samples of an A-scan and a sample k steps into the future. The networks for the different prediction horizons were all estimated on a 3×3 neighbourhood of A-scans belonging to one example of a metal AT mine. Furthermore, their basis functions have a spread of 0.5 times the amplitude range of the entire A-scan and a maximum network size limited to 200 hidden neurons. The prediction performance of the RBFN predictor on its estimation dataset is given in the first column of table 1. The other columns contain the peak prediction performance of the same predictor applied to data of other classes.

Table 1: Prediction performance [%] of a RBFN as a function of the prediction horizon

k	Metal AT	Low metal AT	Ground
1	95.49	80.99	82.70
2	96.30	57.80	63.84
3	93.94	31.83	47.43
4	90.76	-32.48	-0.29
5	87.95	-22.33	-11.26
6	86.15	-37.24	-37.46

From table 1 it is clear that the best combination of discrimination and generalization is attained by a RBFN with a prediction horizon of four.

Optimizing the rest of the RBFN parameters: The remaining three parameters of the RBFN (the length of the tapped delay line, the spread of the basisfunctions and the maximum network size) were optimized in a similar fashion as the prediction horizon. Consequently, we'll focus on the most important observations gained during the

optimization process without delving into the detail results.

With respect to the length of the delay line, it was found that the best combination between discrimination and generalization is obtained by using the previous 13 samples to form a prediction. Short delay lines generally result in poor discrimination ability, while the generalization ability reduces if the delay-line which is too long (i.e. the network is forced to perform an overfitting).

Concerning the spread of the basis functions, it was found that a scaling factor of 0.3 times the range of the A-scan amplitude values results in the best discrimination and generalization ability. If the spread is too small, then the network can't generalize at all. If, on the other hand, the spread is too large, the network loses the ability to discriminate between classes. Lastly, it was found that the best marriage of discrimination ability and generalization ability is obtained by a RBFN with a maximum of 200 basis functions.

In summary, the optimized parameter values for a RBFN model for GPR A-scans are given in table 2.

Table 2: Optimized parameter values for a RBFN

Parameter description	Value
Prediction horizon	4
Length of the tapped delay line	13
Spread of the basisfunctions	0.3
Maximum number of basisfunctions	200

5.2 Prediction performance of the RBFN on various classes of mines

The above mentioned RBFN was used to form time-series models for the following classes of GPR targets: metal AT mines, low metal AT mines (namely: VSHCT, TMA-3 and No. 8), AP mines (M14), meerkat tunnels and clean soil. Within the general class of clean soil two subclasses were modelled: one inside the sandpit in which the landmines were buried and the other (red ground) outside the sandpit.

The peak prediction performance of the various RBFN class-models are given in table 3. The prediction performances in this table were measured in 3×3 neighbourhoods at the locations indicated in the columns of the table. Each row represents a different model (estimated on a particular class). Ideally, this table should resemble a diagonal matrix, with a few exceptions. One would for example expect that a model estimated on one instance of a metal AT mine should perform well on both examples of metal AT mines in the dataset.

The results in table 3 generally indicate that the individual models can discriminate well between their A-scans of their class and other classes. Unfortunately only the models fitted to the high metal AT mines exhibited useful generalization ability. The lack of generalization ability of

Table 3: Peak prediction performance [%] of RBFN models trained on different classes

	Metal AT(1)	Metal AT(2)	VSHCT	No.8	TMA-3	M14(2)	M14(1)	Tunnels	Sandpit
Metal AT(1)	89.77	44.48	13.61	12.31	17.73	12.37	24.22	5.72	-4.44
Metal AT(2)	38.68	81.69	-31.19	-19.15	-24.13	10.95	15.55	-47.54	-10.67
VSCHT	12.77	11.65	74.13	16.13	20.88	22.36	35.03	15.72	37.72
No. 8	8.51	14.97	31.61	81.94	30.00	32.46	27.67	26.76	21.24
TMA-3	10.87	12.69	26.64	33.82	80.07	27.42	33.54	31.16	26.95
M14(2)	10.96	9.26	32.19	17.44	35.55	72.54	37.31	25.83	25.49
M14(1)	12.18	11.91	28.71	7.96	33.70	24.61	74.83	23.72	27.25
Tunnels	18.01	17.88	27.39	21.57	30.42	26.11	44.28	79.32	30.54
Sandpit	14.64	8.58	34.05	25.74	-2.35	28.01	37.06	2.32	81.39

the models estimated on the AP mines and examples of clear ground will be investigated in greater depth shortly.

By applying any of the models mentioned in table 3 to an entire C-scan, it is possible to form a plot of the prediction performance of the particular model as a function of the cross-track and down-track indices. Figure 10 is an example of the prediction performance of a RBFN model estimated on one example of metal AT mines. From this figure it is clear that the model only performs well on two areas. The latter areas coincide with the positions of metal AT mines. Figure 10 therefore shows that the RBFN model estimated on metal AT mines can both discriminate and generalize well.

The generalization ability of the RBFN models estimated for high metal AT mines, however, isn't shared by the models for the various low metal mines. As an example, figure 11 shows the prediction performance of the model for TMA-3 mines. Clearly the model can be used to recognise A-scans from its estimation dataset. The lack of duplicate versions of this specific type of mine in the dataset makes it difficult to come to general conclusions, but it doesn't seem as if the model has the capability to generalize well to other types of low metal ATs.

Some of the other low metal AT and AP mines (e.g. VSHCT and M14) aren't clearly distinguishable from the soil in which they are buried. This can be clearly seen from the prediction performances of the respective models trained on these classes of mines. As an example, figure 12 shows the prediction performance of the RBFN model estimated on one of the M14 mines. The first observation that can be made from this figure is that the model can't generalize well to include the other example of the M14 mine. More interesting however is the fact that a relatively large number of A-scans in the immediate vicinity of the training set are also modelled quite well by the RBFN. This suggests that the "M14-model" is actually to a large extent a model for the specific soil in which the mine was buried.

5.3 Classification performance of the RBFN-based classifier

The complete procedure for the entire classification system is as follows:

1. Train the individual RBFN models on examples of their respective classes. This training phase occurs offline before the system has to sweep an area for mines. Consequently, the time constraints for training are less stringent than during the final classification stage.
2. When an area has to be swept, the following steps have to be followed:
 - (a) Perform clutter removal. At present, this algorithm is performed one down-track B-scan at a time.
 - (b) Determine the prediction performance of each RBFN model in the classifier's database. This step is performed for each A-scan in the cross-track B-scan.
 - (c) Determine the model with the best prediction performance. This will form the system's hypothesis of the identity of the A-scan.
 - (d) Only allocate the identity of the best model to a particular A-scan if the prediction performance of the model is better than a minimum threshold (e.g. 50 %). In this manner extrapolation errors can be avoided and control exercised over the false alarm rate of the system.
 - (e) Repeat the entire process for the next cross-track B-scan.

The final classifier only consisted of five models. One trained on one example of high metal AT mines. The low metal AT class was represented by a model estimated on a No. 8 mine, while only one of the two M14 mines was used to obtain a model for the AP mine class. Lastly, meerkat tunnels and clear ground were also represented by one model each.




Figure 10: Prediction performance [%] of the RBFN trained on one example of metal AT mines




Figure 11: Prediction performance [%] of the RBFN trained on the TMA-3 low metal AT mine




Figure 12: Prediction performance [%] of the RBFN trained on one of the M14 mines

The classification performance of the above mentioned system is reported in figure 13. This figure shows a horizontal slice of the time-domain GPR data at a "depth" of 6.1 nanoseconds. Classification was performed with a minimum threshold of 50 % and the classification decisions are superimposed on the time-domain GPR image. The symbols used in figure 13 convey the following meaning:

- Red dots indicate high metal AT mines (as identified by the system).
- Blue stars indicate low metal AT mines (as identified by the system).
- Green stars represent AP mines (as identified by the system).
- Magenta dots represent meerkat tunnels (as identified by the system).
- Yellow squares indicate the true position of high metal AT mines.
- Cyan squares show the true location of low metal AT mines.

- Black squares represent the true position of AP mines.

The results in figure 13 show that the system can recognize both instances of high metal AT mines, but has problems generalizing between the various low metal mines in addition to a high false alarm rate for low metal AT mines. A positive feature of the system is that it can correctly distinguish between high metal mines and meerkat tunnels, which is important since meerkat tunnels often resemble high metal AT mines quite closely.

The small size of the dataset and its relative uniqueness (due to the use of a commercial SFCW GPR antenna array) coupled with the realism of the test setup makes comparison with the existing literature quite difficult. It does however seem as if the detection system presented in this paper is quite capable of accurate detection of high metal AT mines, but does struggle with low metal mines.

6. CONCLUSION

In this paper a new clutter removal algorithm for GPR data is presented as well as a novel application of neural network-based time-series models for landmine detection. The system does seem promising for the detection of metal




Figure 13: Classification results for a decision threshold of 50 %

AT mines. The algorithm's inability to reliably detect low metal mines is a source of concern, since this defect does limit its practical applicability for humanitarian demining operations. Similar problems have however been recently reported elsewhere in the GPR literature [22].

The large variation in the dielectric properties of both the landmines (ranging from AT to AP mines) as well as the soil in which they are buried, leads us to the conclusion that it seems unlikely that one single sensor and classifier will be able to detect all mines in every practical situation. Combining (or fusing) the recommendations of different classifiers operating on the outputs of different sensors (e.g. GPR and electromagnetic induction sensors) has recently been reported on in the literature [23] and [22]. The results obtained by this approach seem promising enough to warrant additional work, but significant improvements will still have to be made to solve the mine detection problem in practice. The next paper describes alternative approaches to both clutter removal and classification.

7. APPENDIX: MINES REPRESENT OUTLIERS IN GPR DATA

The clutter removal task can be rephrased as removing everything from the data that isn't due to the presence of a landmine. The choice of clutter removal techniques is heavily influenced by the frequency of occurrence of a landmine in the field. With the exception of the border between North and South Korea, the "Cordon Sanitaire" in Zimbabwe represents one of the minefields in the world with the highest density of landmines [24]. This 25 meter wide corridor has an average density of 5,500 mines per kilometer, which can for the sake of argument be approximated as a density of 0.22 mines per square meter (under the assumption of an uniform distribution of landmines). A model V1821 GPR antenna array from 3D Radar has an effective scan width of 1.575 meter [25], which would result in the occurrence of one landmine every 2.886 meter of down-track movement. This area corresponds to 606 A-scans (if down-track samples are separated by 10 cm). Consequently, it is safe to assume that landmine signatures represent outliers (at most one in 606 A-scans) in field-measured data.

REFERENCES

- [1] *Landmine monitor 2011*. Canada: Mines Action Canada, October 2011.
- [2] K. Schreiner, "Landmine detection research pushes forward, despite challenges," *IEEE Intelligent Systems*, vol. 17, pp. 4–7, March/April 2002.
- [3] L. Robledo, M. Carrasco, and D. Mery, "A survey of land mine detection technology," *International Journal of Remote Sensing*, vol. 30, no. 9, pp. 2399–2410, May 2009.
- [4] P. Gader, W.-H. Lee, and J. N. Wilson, "Detecting landmines with ground-penetrating radar using feature-based rules, order statistics, and adaptive whitening," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, no. 11, pp. 2522 – 2534, 2004. [Online]. Available: <http://dx.doi.org/10.1109/TGRS.2004.837333>
- [5] J. N. Wilson, P. Gader, W.-H. Lee, H. Frigui, and K. Ho, "A large-scale systematic evaluation of algorithms using ground-penetrating radar for landmine detection and discrimination," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 8, pp. 2560 – 2572, August 2007. [Online]. Available: <http://dx.doi.org/10.1109/TGRS.2007.900993>
- [6] S. Tjora and E. Eide, "Evaluation of methods for ground bounce removal in GPR utility mapping," in *Tenth International Conference on Ground Penetrating Radar*, 2004.
- [7] O. Lopera, E. C. Slob, N. Milisavljevic, and S. Lambot, "Filtering soil surface and antenna effects from GPR data to enhance landmine detection," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 3, pp. 707 – 717, March 2007. [Online]. Available: <http://dx.doi.org/10.1109/TGRS.2006.888136>
- [8] K. Ho, L. Carin, P. D. Gader, and J. N. Wilson, "An investigation of using the spectral characteristics from ground penetrating radar for landmine/clutter discrimination," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 46, no. 4, pp. 1177 – 1191, 2008. [Online]. Available: <http://dx.doi.org/10.1109/TGRS.2008.915747>

- [9] A. Van der Merwe and I. Gupta, "A novel signal processing technique for clutter reduction in GPR measurements of small, shallow land mines," *IEEE transactions on geoscience and remote sensing*, vol. 38, no. 6, pp. 2627–2637, November 2000.
- [10] K. Ho and P. 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, pp. 1374–1384, June 2002.
- [11] V. Kovalenko, A. G. Yarovoy, and L. P. Ligthart, "A novel clutter suppression algorithm for landmine detection with GPR," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 11, pp. 3740 – 3750, 2007. [Online]. Available: <http://dx.doi.org/10.1109/TGRS.2007.903694>
- [12] D. Daniels, *Surface-penetrating radar*. London: The Institution of Electrical Engineers, 1996.
- [13] P.-A. Sandness, *3DR, 3d-Radar data format: version 1.3*, 3d-radar, December 2009.
- [14] R. Gonzalez and R. Woods, *Digital image processing*. Reading, Massachusetts: Addison-Wesley, 1992.
- [15] J. Bendat and A. Piersol, *Random data: analysis and measurement procedures*. New York: Wiley-Interscience, 1971.
- [16] B. Karlsen, J. Larsen, K. Jakobsen, H. Sørensen, and S. Abrahamson, "Antenna characteristics and air-ground interface deembedding methods for stepped-frequency ground penetrating radar measurements," in *In Proc. of SPIE, AeroSense 2000*, vol. 4038, April 2000.
- [17] J. Lester and L. Bernold, "Innovative process to characterize buried utilities using ground penetrating radar," *Automation in Construction*, vol. 16, p. 546555, 2007.
- [18] C. Bishop, *Neural networks for pattern recognition*. USA: Oxford university press, 1996.
- [19] F. Roth, P. Van Genderen, and M. Verhaegen, "Convolutional models for buried target characterization with ground penetrating radar," *IEEE transactions on antennas and propagation*, vol. 53, no. 11, pp. 3799–3810, November 2005.
- [20] X. Feng, M. Sato, Y. Zhang, C. Liu, F. Shi, and Y. Zhao, "CMP antenna array GPR and signal-to-clutter ratio improvement," *IEEE Geoscience and Remote Sensing Letters*, vol. 6, no. 1, pp. 23 – 27, January 2009. [Online]. Available: <http://dx.doi.org/10.1109/LGRS.2008.2006634>
- [21] E. Ifeakor and B. Jervis, *Digital signal processing: a practical approach*. Harlow, England: Addison-Wesley, 1993.
- [22] A. C. B. Abdallah, H. Frigui, and P. Gader, "Adaptive local fusion with fuzzy integrals," *IEEE Transactions on Fuzzy Systems*, vol. 20, no. 5, pp. 849 – 864, 2012. [Online]. Available: <http://dx.doi.org/10.1109/TFUZZ.2012.2187062>
- [23] H. Frigui, L. Zhang, P. Gader, J. N. Wilson, K. Ho, and A. Mendez-Vazquez, "An evaluation of several fusion algorithms for anti-tank landmine detection and discrimination," *Information Fusion*, vol. 13, no. 2, pp. 161 – 174, 2012, bayesian fusion;Borda Count;Context dependent;Decision template;Dempster-shafer;Fuzzy integral;Land mine detection;. [Online]. Available: <http://dx.doi.org/10.1016/j.inffus.2009.10.001>
- [24] *Landmine monitor report 2009: toward a mine-free world. Special ten-year review of the mine ban treaty*. Canada: Mines Action Canada, October 2009. [Online]. Available: www.lm.icbl.org/lm/2009
- [25] "V-series multi-channel antenna arrays for GPR: Mkt-ds-v-series-040810v1," 2010.

THE INFLUENCE OF MAJOR EXTERNAL AND INTERNAL EVENTS ON THE CULTURE OF AN ENGINEERING ORGANISATION

W. Theron*, L. Pretorius** and K.-Y. Chan***

* Master's student, Department of Engineering and Technology Management, Graduate School of Technology Management, University of Pretoria, Lynnwood Road, Pretoria, 0002, South Africa
E-mail: willie@wazat.co.za

** Professor, Department of Engineering and Technology Management, Graduate School of Technology Management, University of Pretoria, Lynnwood Road, Pretoria, 0002, South Africa
E-mail: leon.pretorius@up.ac.za

*** Senior Lecturer Department of Engineering and Technology Management, Graduate School of Technology Management, University of Pretoria, Lynnwood Road, Pretoria, 0002, South Africa
E-mail: alice.chan@up.ac.za

Abstract: A Case study company that was set up as a project where the technical focus, activities and behaviour set the initial culture is considered in this research. Over a period of 11 years the Case study engineering organisation was exposed to many influences in the electrical utility industry that now give lead to questions such as: How did events influence the engineering culture and how did the culture change over time? Engineering organisations are subjected to external and internal events which are not always within their control. These include technological changes, economical changes or new competition, change in ownership, business focus or technical leadership. The ability to absorb such events is not only a function of the organisation's technology infrastructure, availability of funding or skills, but also of the organisational culture prevailing at the time. The objective of the research is to determine how eight events impacted the culture of an engineering organisation over a period of six years. The results show that the culture is indeed influenced by events, with an indication that the different work areas within the organisation experienced the cultural changes differently. The employees that worked for the organisation six years or longer also experienced the changes differently from those that were only employed for the last five years of the organisation's life. These results may assist the understanding of the impact that events may have on an organisation and allow early risk mitigation to counter undesirable culture forming.

Keywords: organisational culture, change management, event impact, engineering organisation.

1. INTRODUCTION

All engineering organisations are subjected to external and internal events, which are not always within their control. External events can include technological changes, global or local financial impacts, market shifts or new competition entering the market. Some of these events may be initiated by a black swan event as defined by Taleb [1]. On the other hand internal events can include critical skills shortages, change in ownership, and change in business focus or change in technical leadership. These events will all test the resilience of the organisation. The ability to absorb such events is not only a function of the organisation's technology infrastructure, knowledge management or availability of funding or skills, but is also driven by the culture prevailing in the organisation at the time.

Understanding how a major event may impact the culture of an engineering organisation, may make it possible to prepare the organisation in advance of the event with the intent to stay on course or to adjust course to be able to absorb the event.

1.1 Theory and research review

The research topic implies that organisational culture is investigated as part of the research. Schein [2] offers the following definition for organisational culture: "A pattern of shared basic assumptions that the group learned as it solved its problems of external adaptation and internal integration, that has worked well enough to be considered valid and, therefore, to be taught to new members as the correct way to perceive, think, and feel in relation to those problems".

Brown as quoted by Manetje et al [3] defines organisational culture as: "the pattern of beliefs, values, and learned ways of coping with experience that have developed during the course of an organisation's history, and which tend to be manifested in its managerial arrangements and in the behaviours of its members".

Both these definitions suggest that culture is developed over time, but do not imply how an organisation may behave should an event be imposed on the organisation. A large number of change management models have been developed, but when approaching change management