2012 14th International Conference on Ground Penetrating Radar (GPR)
June 4-8, 2012, Shanghai, China @ ISBN 978-1-4673-2663-6

# Mining Sea Turtle Nests

## An Amplitude Independent Feature Extraction Method for GPR Data

Vladimir Ermakov, Artur Dubrawski
The Robotics Institute, Auton Lab
Carnegie Mellon University
Pittsburgh, PA, USA
vermakov@cmu.edu, awd@cs.cmu.edu

Jessica Hodgins
Disney Research
Pittsburgh, PA, USA
jkh@disneyresearch.com

Tony Dohi
Walt Disney Imagineering
Research & Development
Tony.P.Dohi@disney.com

Anne Savage
Disney's Animal Kingdom
Education and Science Department
Lake Buena Vista, FL
Anne.Savage@disney.com

***Abstract*— We use a Ground Penetrating Radar (GPR) to localize eggs of sea turtles laid in sand. GPR technology has been developed to detect subsurface structures, and successfully applied in archeology, civil engineering, and demining. Typical uses rely on relatively strong signals due to high contrast in dielectric properties of the buried manmade objects and the soil. Signal to noise ratios in our task are substantially lower, as the variances in humidity and granularity of layers of salty sand, and the presence of nuisance artifacts such as rocks, clogs of seashells, air pockets, etc., contribute to making turtle nest detection a challenging task. We present a combination of signal processing, pattern recognition, and feature selection techniques that stand up to these challenges. Our approach is evaluated using ground truth data collected in the field. We believe that this method can be useful in a range of non-standard GPR applications, especially when the signals to noise ratios are low.**

***Keywords-ground penetrating radar; signal processing; pattern recognition***

## I. Introduction

All species of sea turtles that inhabit waters surrounding the United States are considered threatened or endangered. Depending on the species, most females may have multiple clutches in a season and can lay between 80-120 eggs in each nest. Nest success is influenced by many factors: temperature, natural disasters, predation, and poaching. Efforts to effectively and easily locate clutches will assist in the long-term monitoring of nesting success. Current techniques to locate clutches rely on the observer examining the sand for clues that indicate whether the sea turtle has nested or not. Depending on the species of turtle, locating eggs is complicated by the depth of the nest. Some turtles can dig more than 1 meter deep, thus significantly decreasing the probability of locating the clutch with ease. We are testing the utility of using Ground Penetrating Radar (GPR) technology combined with signal processing and statistical data mining, to enable robust and practical automation of this process.

### *A. Ground Penetrating Radar Processing Background*

Ground Penetrating Radars (GPRs) are sensors that use electromagnetic signal for detecting sub-surface objects. The principle of operation is that the signal travels through the medium undisturbed until there is a shift in propagation speed, caused by a change in electromagnetic properties of the materials. At these transition points, part of the signal is reflected back and captured by an antenna. GPRs are often used for detecting pipes, looking for underground water, determining road and pavement condition, archeological explorations, and for finding land mines [11]. The idea of applying these sensors for detecting biological materials is fairly novel. In particular, returns obtained in wet salty sand, which is known to be a difficult medium in itself, often do not provide substantially distinct signatures for turtle eggs given the similarity of their electromagnetic properties when compared to background.

The data is collected by scanning the ground underneath the radar. The GPR signal propagates as a cone, so reflections that are detected are not necessarily coming directly from under the radar antenna. Typically, the radar is dragged along the surface of the sand in a grid like pattern: e.g. first a series of scans parallel to the shore, and then a series of scans perpendicular and overlapping with the previous scans.

GPR data can be interpreted at the level of echoes from single impulses often called A-Scans [12]. A-Scans are wave forms with time being approximately proportional to depth, and amplitude being indicative of the reflection that is received at that depth. The single echoes can be appended together for the duration of the parallel or perpendicular swipe across the beach. The result of this may be interpreted as a 2-D grayscale image, often referred to as B-Scan [12], with depth on the Y axis, distance along the sensor trajectory on the X axis, and the level of intensity indicative of the amplitude of the reflection as seen in Fig. 1.

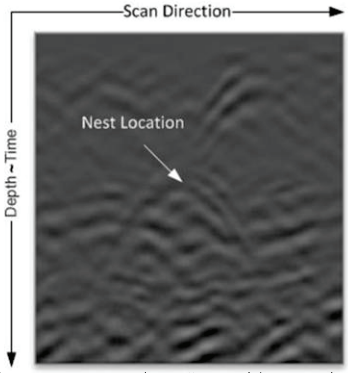

Figure 1. Example B-Scan with a nest in it.

*B. Our Goals*

The goal of our research is to create an application that could confidently determine location of sea turtle nests in a new environment with the minimal amount of expert supervision. For this, we need a method for extracting features from the GPR data, a classifier to assess probability of a nest presence in data and to determine which features are informative of nest locations, and a mechanism for combining evidence from collected data in order to render a meaningful picture for the operator to interpret. To achieve ease of use, we want the interaction with the operators to take place at a relatively high level. We need them to tell the system about locations of the collected scans, and to supply any additional information, such as verified locations of actual nests, to enable feedback-loop based tuning of the system. We specifically intend to limit the user's exposure to the art of tuning filtering parameters, often required by commercially available GPR software.

## II. Related Work

One common approach for locating sub-surface objects is by processing individual B-Scans, and detecting hyperbolic reflections [8]. Reflections appear hyperbolic because of the geometric properties of the GPR echoes. The shape of a hyperbola is determined by the size, shape, and relative position of the sub-surface object. Algorithmic detection of such hyperbolic reflections is not a trivial task. The data is often very noisy and contains many artifacts, the most prominent of which is the reflection between air and the ground surface. Above ground objects are also known to create noise and distortion [10].

The general procedure for pre-processing GPR B-Scans usually consists of background removal, clutter removal, and thresholding. Background removal helps clear prominent air-to-ground and other constant reflections that disrupt the signal. This task can be accomplished by subtracting from the raw signal a windowed average of the amplitude for each depth [10]. It is worth noting that specific applications may dictate specific algorithms for background removal, e.g. when the features of interest are long and horizontal, such as utility pipes. Clutter removal is often accomplished by applying image filtering techniques, such as the Wiener filter, to the data [6]. Thresholding retains only strong reflections and its extent is often inversely proportional to the depth, as reflected signals typically become weaker with depth [10].

A common technique for detecting hyperbolic shapes in GPR data is a Randomized Hough Transform (RHT) [1,2,7]. It is often used in Computer Vision for detecting image primitives such as lines and circles, and has been adopted to detect hyperbolas. While being robust to noisy data, RHT suffers significantly from computational overhead. An alternative approach, Probabilistic Robust Hyperbola Mixture modeling has been proposed [3] which achieves similar results with much less computation. Both approaches rely on fairly robust inputs that can be provided via pre-processing, where pixels of interest are to a large extent isolated from noise. Those techniques have been applied with some success to land-mine and utility pipe detection, where the contrast of electric conductivity between the soil and objects of interests was significant enough that an effective threshold of returned amplitude could be picked to isolate the signal from noise.

## III. Amplitude-Independent Feature Extraction

Because of the relative weakness of the signal of interest compared to other reflections in our data, and the inability of previously introduced filtering techniques to separate the signal from noise sufficiently well to warrant a practical use, we had to develop a new feature extraction method. Despite the weakness of the amplitude of reflection, the signal coming back from buried turtle nests still forms hyperbolic shapes, which are often noticeable in B-Scans as local maxima and minima of the returned signal. Our extraction method identifies these local minima and maxima in every A-Scan, and groups them across subsequent A-Scans into what we call Weak Object Reflection Merges or WORMs, as can be seen in Fig. 2. The WORMs consist of groups of adjacent pixels, and properties of these groups of pixels can be used to identify which WORMs may be relevant to sub-surface objects of interest, and which can be discarded as noise.

To extract WORMs we implement the following algorithm:

*1)* For every A-Scan, vertically slide a window of a fixed size.

*2)* Mark pixels in an allocated B-Scan-sized buffer if the center of the window pixel has the strongest or weakest amplitude in the window.

*3)* Use a flood-fill algorithm to group all the adjacent marked pixels into WORMs.

For each of the WORMs, we compute the following features:

- Number of pixels in the group.
- Average depth of pixels in the group.
- Height to width ratio of the bounding box that fits the group.
- Sum of orthogonal deviations of all the intermediate pixels in the group from the line that connects the starting and ending pixels (left to right), normalized by the number of pixels.
- Maximum positive (upward) deviation from this line normalized by the number of pixels in the group.

We used stepwise regression [13], iteratively adding features one at a time keeping only those that maximized the R-square goodness-of-fit metric between our final estimate of nest location and true nest location. Other features we found to add little extra information and did not use in further experiments were maximum depth, minimum depth, distance between first and last pixel, number of pixels per unit length, average amplitude value of the pixels and average deviation from the line between first and last pixel.

Figure 2. Processed B-Scan with WORMs extracted and classified. Whiter lines are predicted to be potential nest WORMs, and darker lines correspond to predicted background WORMs.

In an average B-Scan, the method outlined in the previous section would extract close to 1,000 WORMs. Many of these would be very short, and omnipresent, while others would be very elongated, such as the air to surface reflection. Only a small subset of detected WORMs would have interesting characteristics, such as having hyperbolic appearance. We called WORMs that are omnipresent, and are not indicative of a subsurface obstacle as "background" WORMs, and those that had interesting characteristics that could be useful for discovering subsurface obstacles we called "Nest" WORMs. Our ultimate goal was to minimize involvement of the operator in the training process, so the only feedback data available to the system were the verified locations of discovered nests. Our classification approaches had to use proximity to these locations for learning the appropriate parameters for distinguishing between "Nest" and "Background" WORMs.

### B. Rule Based Classification

Our first approach was to use human expertise to identify a set of rules (fixed thresholds to compare feature values against) that best separated background and nest WORMs. We used a subset of B-Scans from the first 5 beach locations, also called "grids", to pick a reasonable set of rules for detecting "Nest" WORMs. We took into consideration the need to filter out air to surface reflection, which were a series of WORMs with a large number of pixels (more than 100) with low height to pixel count ratios (less than 0.1) and shallow depths (less than 110 pixels in our B-Scans). We also incorporated the knowledge about the range of expected depths of nests, ignoring features deeper than 2.5ft (250 pixels in our B-Scans). Additional constraints were determined in order to restrict the WORMs to have a reasonable minimal length (more than 10 pixels), to be concave, and to curve downwards (maximum deviation from line normalized by number of pixels is above 0.12).

The downsides of this method are that once the rules are determined for a set of training data, they are not adjusted as more data is gathered in the field. Surprisingly these simple rules showed to be remarkably robust and they performed slightly better than the other more sophisticated and flexible approaches that we explored for learning of the parameters.

### C. Determining Rules Automatically

In an effort to make the process of WORMs classification less dependent on expert knowledge and to avoid making specific assumptions about the data, we went on to explore other classification methods. In order to filter out signal from noise, we needed to identify WORMs that were more likely to be located in the vicinity of the object of interest. We separated our data into test and training subsets for leave-one-grid-out cross-validation (each cross-validation cycle used one grid for testing and the remaining grids of data for training predictors). We labeled WORMs that were within 1.5ft of the confirmed nest location as potential nest WORMs, and those that were outside of this radius as background WORMs. The task was to identify sub-regions of our parameter space in which WORMs were more likely to be part of the potential nest set, than in other regions of the parameter space.

To find these sub-regions we applied the PRIM bump hunting algorithm [4]. In a nutshell, the algorithm shrinks a bounding box in the multidimensional parameter space one dimension at a time, in order to optimize the objective function. The function we defined was the likelihood ratio:

$$L = \frac{P(Potential\ Nest\ |InBox)}{P(Non - nest|InBox)}$$

We adjusted this ratio with a generality score to make sure that the resulting restrictions were general enough to work across multiple B-Scans. The generality score was computed as follows:

$$G = \frac{\#\ BScans\ with\ potential\ nest\ WORMs\ in\ box}{\#\ BScans\ With\ potential\ Nest\ WORMs}$$

The final score was therefore:

$$S = L * G$$

Once the score function was optimized, WORMs inside the box were updated with an additional feature – their relevance score, were removed from the training dataset, and the algorithm was re-run on the remaining WORMs. The process was repeated until the optimized score fell below a predefined threshold. This threshold could be set as a percent improvement over the baseline score of the whole parameter space:

$$Baseline\ Score = \frac{P(Potential\ Nest)}{P(Non - Nest)}$$

The intuition behind this approach is that there are several clusters of features within our parameter space that may be relevant for sub-surface objects. Most notably there often are clusters of mini-hyperbola shapes, and clusters of almost straight diagonal lines (tails of these hyperbolas).

*D. Using Spectral Scores for WORMs Classification*

In an alternative method for WORMs selection, we applied Principal Component Analysis (PCA) on the WORMs belonging to the background in the training data (split using leave-one-grid-out cross validation approach just like in previous section). We then used the Mahalanobis distance metric along the top principal component vectors to measure spectral distance. The idea behind this approach is that WORMs that are similar to the background have low spectral distance, and WORMs that are less similar to the background, presumably indicative of subsurface obstacles, are further away. We then used Linear Discriminant Analysis (LDA) to find a decision boundary that maximized the distance between the mean "nest" and mean "background" WORMs.

The main advantage of this approach is that the features are completely abstracted out, and the determination of relationships between them is left to the algorithm. This property is desirable since it should generalize directly to other types of GPR data and other types of objects to be detected, as long as the relevant training data is available. The model parameters are learned from training data, and the method itself is ignorant to whether one is looking for sea turtle nests or for land mines.

*E. Aggregation of WORMs*

Previous sections discussed methods of selecting WORMs that are likely reflections of a buried nest. To determine the actual nest location from the filtered subset of WORMs, a meaningful aggregation of evidence needs to take place. Our approach was to use kernel density estimation with a hyperbolic Gaussian kernel on B-Scans. By hyperbolic Gaussian kernel we simply mean the aggregation of scores obtained by sampling orthogonal distances of selected pixels to a fixed hyperbola from a Gaussian. This process implements the following algorithm:

1. Subdivide the B-Scan into a grid.
2. For each intersection of grid lines:
   a. Determine parameters for a hyperbola centered at that grid location.
   b. For each pixel belonging to "nest" WORMs (determined using one of the classification methods) :
      i. Compute the orthogonal distance to the hyperbola by applying the method described in [3] and [9] and determine the score by sampling into a normal distribution function.
      ii. Compute the weight by summing up the scores for all pixels.
      iii. Save the weight to the location in the grid.

This process is similar to the migration technique [8] often used as a filtering step in GPR data processing.

After the hyperbolic kernel is applied to accumulate evidence from WORMs in individual B-Scans, a 3-dimensional Gaussian kernel is applied to combine evidence from multiple B-Scans. The most likely location of a nest is then assumed to be the cell in three dimensions where aggregated values are maximized. In practice of sea turtle nest detection, depth information is of secondary importance to the horizontal placement of the nest. For that reason, we project the 3-dimensional likelihood grid onto a plane spanning the top view of the grid area. Each 2-dimensional grid element corresponds to the maximum value for each vertical column of the 3-dimensional density map. Hot-spots in the resulting 2-dimensional heat map are considered as predictions of the most likely nest locations.

*F. Determining Parameters for the Kernels*

The previous section mentioned a few fixed parameters, without suggesting a method for determining the best values for them. These parameters are

- Parameters for defining the shape of the hyperbolic kernel.
- Standard deviation of the hyperbolic Gaussian.
- Standard deviation of the three-dimensional Gaussian.

These parameters can be determined by gradient descent optimization on training data. For each grid where a nest was found we measured the error estimated as distance between the actual nest location (determined by digging of the soil confirm presence of the nest), and the suggested location that is the output of the algorithm described in the previous section. The cost function we optimized was mean absolute error across all grids in our leave-one-grid-out cross validation experiments.

## IV. Field Application

As we mentioned, our goal was to develop an easy to use application to help biologists in the field discover sea turtle nests effectively while reducing labor. We combined the techniques discussed in previous sections into a prototype implementation in MATLAB. This application allows the user to upload scans from the ground penetrating radar as they are being collected, and it generates a heat-map that shows the likelihood of having a nest in any particular location of the scanned area. The application implements two classification methods for WORMs – the hand-tuned parameters (Section III-B) and the spectral distance method (Section III-D), and allows the user to select their method of preference. The application allows users to input locations where they dug the ground and mark whether they discovered a sea turtle nest or not. This information is immediately used to update the spectral distance model for classifying WORMs, and the heat-map is recomputed using the updated model. Fig. 3 shows a screenshot of the application being used, rendering a likelihood map of a section of the beach with the confirmed nest location marked with a green square.

## V. Empirical Evaluation

*A. Data-Set Description*

The data for our experiments was gathered using a GSSI 800 MHz antenna. 26 nesting sites were surveyed rigorously, recording location and direction of each scan. In most cases B-scans were taken 1ft apart, on an 8ft by 8ft square area. An expert biologist then dug the sand at suspicious locations to check whether a nest was there or not and the location was

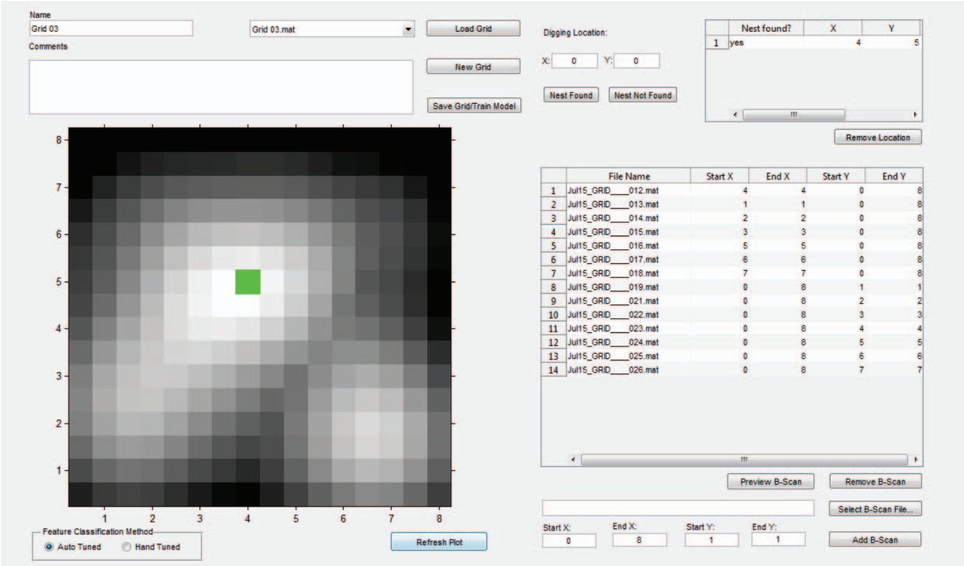


Figure 3. Screen-shot of our field test application. The application is rendering a heat-map of likelihood of discovering a nest. Heat-map was generated from WORMs selected using spectral classification method, learned from data collected in the field.

labeled accordingly in data. In the course of data collection 60 holes were dug, and 18 nests were discovered. The dataset consisted of more than 460 individual B-scans.

### *B. Results and Discussion*

To evaluate potential utility of our approaches we used two metrics computed via cross-validation. The first is mean localization error measured as the average distance between the predicted and the actual locations of nests used in testing. Another metric that is important is the rate of nest recall for certain distances, such as 0.5, 1.0, and 1.5 feet radii. It is measured as the fraction of true nests within the specified radius of the model maximum for each grid. This metric is reflective of the amount of effort necessary to confirm existence of a nest by digging in the sand. The higher the recall rates at small radii the more accurate the localization and the more manual labor can be saved by using our automated approach. Both of these metrics were measured against scan grids which contained confirmed nests.

In our experiments, we evaluated the three different approaches for WORMs classification. The results are summarized in Table 1. “Hand-tuned” refers to the approach from section III-A, “Learned” refers to the approach in section III-B and “Spectral” refers to the approach from section III-C. Fig. 4 shows results of the KDE estimation of likelihood on all of the sections of the beach using the hand-tuned approach. The colored symbols represent the ground truth: green circles show the locations where digging took place and nests were found, and red squares are locations where digging took place and no nests were discovered.

Interestingly, the manually designed decision rules turned out to be quite accurate on average, slightly but insignificantly outperforming the spectral method. Potential utility of machine learning approach is shown by the best recall rates at the smallest of the tested radii of localization tolerance. Weaker mean scores of the machine learning method result from it missing a few true positives. We believe that additional tuning and incorporation of additional features could substantially boost performance of machine learning and spectral methods. An important limiting factor that favored the hand-crafted design was the small number of labeled nest data available to us at the time of writing this paper. We plan to remedy those limitations in the near future.

TABLE I. RESULTS OF EMPIRICAL EVALUATION

| Classification Method | Mean Distance to Nest | % Recall | | |
|---|---|---|---|---|
| | | 0.5 ft | 1.0 ft | 1.5 ft |
| Hand-tuned | 0.83 (+/-0.098) | 33 | 72 | 94 |
| Learned | 1.14 (+/- 0.499) | 50 | 63 | 75 |
| Spectral | 0.87 (+/- 0.160) | 38 | 67 | 89 |

## VI. FUTURE WORK

In our current approach, local maxima of the three-dimensional kernel are interpreted as potential nest locations. We found that in some cases, these maxima were found in the areas where no nests were present. Every time, however, we found sufficient evidence in the B-Scans to suggest that these detections were indeed reasonable and hyperbolic reflections were present, suggesting that other subsurface obstacles were in those locations. Often, these locations coincided with places where experts in the field believed nests to exist (experts were not only looking for surface signs but at the raw GPR data as well), but, upon examination, other artifacts such as clumps of sea shells or air pockets were found.

One pragmatic way to mitigate this issue is to leverage user interaction to collect a richer sample of training data, and to enable the system to learn on-the-fly from immediate user feedback right when the outcome of digging becomes available. This way, the system should be able to incrementally

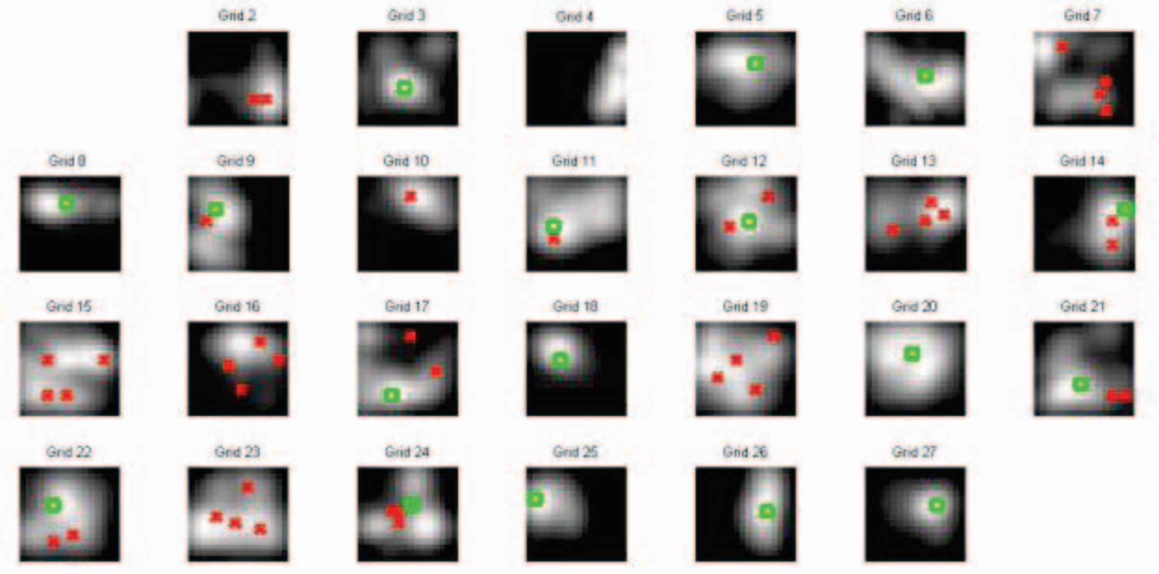


Figure 4. Heat-map representing the likelihood of discovering a nest, generated from test data. Green circles indicate confirmed nest locations. Red crosses indicate locations where no nests were found by digging.

learn over time to ignore detected anomalies that are caused by objects of types that are not of interest. This would also make it adaptable to the specifics of individual locations that could substantially vary from one beach site to another.

There is also room for improvement in data processing and we could incorporate additional potentially informative dimensions into data, such as spectral features from Fourier or Wavelet transforms. We also plan to experiment with a Bayesian approach to aggregating evidence from multiple B-scans and to use active learning and information theory to guide the users as to where to take the next scan to best inform the nest localization analysis. It should cut down the time and effort needed to robustly estimate the place to dig. .

## VII. Conclusion

We presented an amplitude-independent approach for extracting features and several techniques for classifying and aggregating these features for the practical application of detecting and localizing sea turtle nests using Ground Penetrating Radar. Our goal was to develop an easy to use application that could learn which features of the recorded signal are characteristic of subsurface objects of interest, without requiring the operator to have an expert knowledge of GPR signal processing techniques. Despite the difficult data with low signal to noise ratios we were able to demonstrate practically meaningful capability of extracting useful information. We plan to explore these ideas further with additional turtle nest data as well as seek use in a wider set of scenarios where available GPR data is noisy and contains frequent misleading artifacts.


## Acknowledgement

This material is based upon work partially supported by the National Science Foundation under Grant No. IIS-0911032.


## References


[1] L. Xu, E. Oja, and P. Kultanan, "A new curve detection method: Randomized Hough transform (RHT)", Pattern Recog. Lett. 11, 1990, 331-338.

[2] D.H. Ballard, "Generalizing the Hough Transform to Detect Arbitrary Shapes", Pattern Recognition, Vol.13, No.2, p.111-122, 1981

[3] H. Chen and A. G. Cohn, “Probabilistic robust hyperbola mixture model for interpreting ground penetrating radar data” in IEEE World Congress on Computational Intelligence, p.3367–3374, 2010

[4] J.H. Friedman, N.I. Fisher, "Bump Hunting in High Dimensional Data", Statistics and Computing, 9, p.123-143, 1999

[5] Harry M. Jol, "Ground Penetrating Radar: Theory and Applications", First Edition, 2009

[6] L. van Kempen, H. Sahli, E. Nyssen, J. Cornelis, “Signal Processing And Pattern Recognition Methods for Radar AP Mine Detection and Identification” in IEE Conf. Publ. no 458, 2nd Int’l. Conf. on Detection of Abandoned Landmines, p.81-85, 1998

[7] N. Milislavljevic, I. Boch, and M. Archeroy, "Application of the Randomized Hough Transform to Humanitarian Mine Detection", in Proceedings 7th IASTED International Conference on Signal and Image Processing, p.149-154, 2001

[8] Luc van Kempen, H. Sahli, "Ground Penetrating Radar Data Processing: A Selective Survey of the State of the Art Literature", Technical Report IRIS-tr-0060, 1999

[9] W. Gander, G.H. Golub, and R. Strebel. "Least-squares fitting of circles and ellipses", BIT Numerical Mathematics, 34(4):558-578, 1994

[10] A.P.Annan, “Practical Processing of GPR Data”, Proceedings of the Second Government Workshop on Ground Penetrating Radar, 1993

[11] A. Dell’Acqua, A. Sarti, S. Tubaro, and L. Zanzi. “Detection of linear objects in GPR data”, Signal Processing, 84(4):785–799, 2004.

[12] C. Bruschini, B. Gros, F. Guerne, P. Y. Pi`ece, and O. Carmona. “Ground penetrating radar and imaging metal detector for antipersonnel mine detection”. Journal of Applied Geophysics, 40(1-3):59–71, 1998.

[13] Efroymson, MA, “Multiple regression analysis.” In Ralston, A. and Wilf, HS, editors, Mathematical Methods for Digital Computers. Wiley. 1960