ERDC/CRREL MP-22-17

Cold Regions Research and Engineering Laboratory

US Army Corps of Engineers®
Engineer Research and Development Center

Environmentally Informed Buried Object Recognition

Sophia Potoczak Bragdon, Vuong Truong, and Jay Clausen

November 2022

The US Army Engineer Research and Development Center (ERDC) solves the nation's toughest engineering and environmental challenges. ERDC develops innovative solutions in civil and military engineering, geospatial sciences, water resources, and environmental sciences for the Army, the Department of Defense, civilian agencies, and our nation's public good. Find out more at www.erdc.usace.army.mil.

To search for other technical reports published by ERDC, visit the ERDC online library at https://erdclibrary.on.worldcat.org/discovery.

ERDC/CRREL MP-22-17
November 2022

Environmentally Informed Buried Object Recognition

Dr. Sophia Potoczak Bragdon, Vuong Truong, and Dr. Jay Clausen

US Army Engineer Research and Development Center (ERDC)
Cold Regions Research and Engineering Laboratory (CRREL)
72 Lyme Road
Hanover, NH 03755-1290

Final report

Prepared for US Army Combat Capabilities Command
Night Vision and Electronic Sensors Directorate (NVESD)
Fort Belvoir, VA 22060-5806

Under Program Element 0602712A, Project Number 468349, Task Number A10100

Preface

This study was conducted for the US Army Combat Capabilities Command–Night Vision and Electronic Sensors Directorate (NVESD) under Program Element 0602712A, Project Number 468349, Task Number A1010, "Environmentally Informed Buried Object Recognition."

The work was performed by the Biogeochemical Sciences Branch of the Research and Engineering Division, US Army Engineer Research and Development Center–Cold Regions Research and Engineering Laboratory (ERDC-CRREL). At the time of publication, Nathan J. Lamie was branch chief; and Dr. Caitlin A Callaghan was division chief. The acting deputy director of ERDC–CRREL was Bryan E. Baker, and the director was Dr. Joseph L. Corriveau.

This document was originally published with SPIE Defense and Commercial Sensing Conference Proceedings on 30 May 2022. This research was supported in part by the DoD Research Participation Program administered by the Oak Ridge Institute for Science and Education (ORISE) in conjunction with the US Department of Energy (DOE).

COL Christian Patterson was commander of ERDC, and Dr. David W. Pittman was the director.

Environmentally Informed Buried Object Recognition

ABSTRACT

The ability to detect and classify buried objects using thermal infrared imaging is affected by the environmental conditions at the time of imaging, which leads to an inconsistent probability of detection. For example, periods of dense overcast or recent precipitation events result in the suppression of the soil temperature difference between the buried object and soil, thus preventing detection. This work introduces an environmentally informed framework to reduce the false alarm rate in the classification of regions of interest (ROIs) in thermal IR images containing buried objects. Using a dataset that consists of thermal images containing buried objects paired with the corresponding environmental and meteorological conditions, we employ a machine learning approach to determine which environmental conditions are the most impactful on the visibility of the buried objects. We find the key environmental conditions include incoming short-wave solar radiation, soil volumetric water content, and average air temperature. For each image, ROIs are computed using a computer vision approach and these ROIs are coupled with the most important environmental conditions to form the input for the classification algorithm. The environmentally informed classification algorithm produces a decision on whether the ROI contains a buried object by simultaneously learning on the ROIs with a classification neural network and on the environmental data using a tabular neural network. On a given set of ROIs, we have shown that the environmentally informed classification approach improves the detection of buried objects within the ROIs.

1. INTRODUCTION

Buried object detection and classification is a nontrivial task that has been widely researched. This paper uses thermal infra-red (IR) sensors to detect buried objects which is advantageous as this sensor modality can identify objects composed of both metal and plastic and can operate day and night. The key to successful buried object detection with thermal IR is to identify surface thermal anomalies induced by the presence of a buried object. The thermal signature of the surrounding soil is highly heterogeneous which can lead to a low signal-to-noise ratio (SNR) in a thermal image that contains a buried object. The lower the SNR, the more difficult the task of detection and classification of buried objects. Furthermore, it has been shown in[1–4] that the environmental conditions at the time of imaging affect the thermal signature of buried objects. For example, a rain event will increase the soil's volumetric water content (VWC) and the increased soil moisture leads to a "wash out" phenomena making detection of buried objects difficult. The increased soil moisture masks the buried objects since the moisture seeping into the ground creates a thermal equilibrium between the buried objects and the soil. Additionally, the incoming shortwave solar radiation is an important factor for the ability to detect a buried object with thermal IR. Near dawn as the surrounding soil heats at a different rate from the buried object, there is often a window of increased visibility, but by mid-day a thermal equilibrium is achieved due to the solar-loading of the soil and buried object. After sundown, there is another window of increased visibility where the buried objects retain heat at a different rate compared to the surrounding soil, which makes the object detectable after dark. If one considers the environmental conditions at the time of imaging, this can inform the detection and subsequent classification of buried objects using thermal IR sensors.

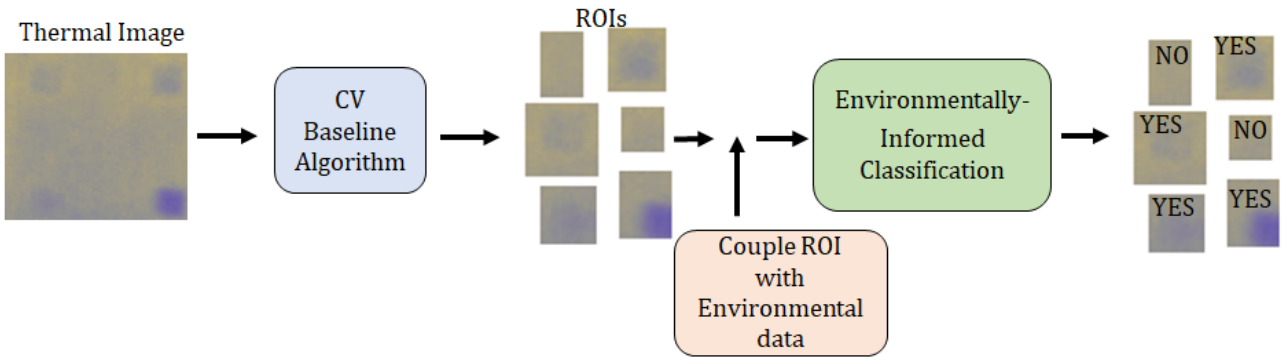

Figure 1. A schematic flow chart representation of the steps of the proposed two-step environmentally informed ATR.

This paper focuses on utilizing the environmental conditions at the time of imaging to inform the recognition of buried objects. Automatic target recognition (ATR) algorithms are comprised of two steps: the detection and the classification of objects. The two classes of ATR algorithms are two-step and one-step algorithms. In a two-step ATR, one algorithm is used to detect a region of interest (ROI) within an image and a second classification algorithm is used to classify the type of object detected in the ROI. In a one-step ATR, one algorithm tackles both detection and classification steps.[5] We present a two-step approach to object detection and classification. The thermal images are processed using a computer vision approach to identify ROIs and, given the low SNR, the ROIs produce many false positives. We then process and classify the ROIs for whether they contain a buried object using an environmentally informed deep learning approach. We augment the ROI with environmental conditions at the time of imaging and process the coupled data into both a tabular neural network for the environmental data and a convolutional neural network for the image data. The results of the combined networks are concatenated and passed through a final classification layer which produces the environmentally informed prediction. Fig. 1 shows a schematic representation of the different steps of the two-step environmentally informed ATR algorithm.

The paper is organized as follows. In Sec. 2, the data used for the two-step algorithm is discussed. In particular, the data collection and the pre-processing that were performed prior to utilization in the detection and classification algorithms. Sec. 3 introduces the computer vision algorithm used for identifying ROIs in the thermal images. In Sec. 4, the environmentally informed classification algorithm is introduced. The architecture of the algorithm is discussed and we compare the environmentally informed classification of the ROIs to the classification of the ROIs using only the CNN. Lastly, in Sec. 5, we discuss the findings and future directions of the work.

2. DATA DESCRIPTION

The data used in this paper is from a multi-year study conducted by the Engineering Research and Development Center (ERDC) and descriptions of this study can be found in Refs. 1,2,6,7. The data was collected using the ERDC / Combat Capabilities Development Command (CCDC) Command, Control, Communications, Computers, Cyber, Intelligence, Surveillance, and Reconnaissance (C5ISR) Test Bed constructed at the Cold Regions Research and Engineering Laboratory (CRREL) in Hanover, NH. Data used in this study was collected from May to November 2018 in the Phase II of the multi-year study.[2] The test bed is 3.05 m by 3.05 m and is composed of a sandy-loam soil, which is kept free of vegetation. The construction of the test bed involved removing the soil and homogenizing the substrate prior to reapplying it with various buried sensors. The homogenizing process was used to minimize background noise in the thermal variations observed on the soil surface without the presence of buried objects. The data collected consists of two stages: pre-emplacement of objects and post-emplacement of buried objects. A forward looking infra-red (FLIR) A310 camera is used to capture long-wave infra-red (LWIR) (7.5 - 13 microns) data of the test bed surface. The FLIR camera records on a five-minute interval.

Four equal-sized rectangular buried objects are located within the test bed and each measure 40 by 40 by 20 cm and are filled with nitrogen fertilizer. There are two plastic (HPDE) objects and two metal objects (aluminum). A shallow plastic (SP) object and shallow metal (SM) object were emplaced in the test bed with

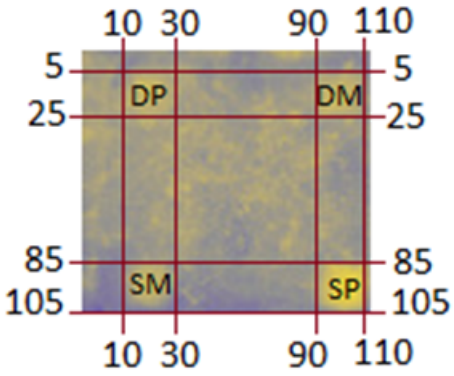

Figure 2. A sample thermal image of the test bed with the buried objects identified along with the corresponding pixel locations.

the top buried at a depth of 5 cm below round surface (bgs) and a deep plastic (DP) object and deep metal (DM) object were buried at a depth of 25 cm bgs. In Fig. 2, a sample thermal image of the test bed is shown and the pixel locations of each object are labeled. The test bed images have dimension 115×105 pixels.

In the test bed, CS655 sensors are used to measure the temperature, electric conductivity, and moisture (volumetric water content) of the soil. The sensors are buried at various depths throughout the test bed (5, 15, and 35 cm bgs). In addition to below-ground measurements, the surface temperature (measured with heat plates), heat flux, and incoming and reflected short wave solar radiation (W/m^2) are recorded. To record local meteorological conditions, there is a meteorological station installed at the CRREL/C5ISR test bed and the measurements include average air temperature (C), average pressure (mbar), average wind speed (m/s), rain total (mm), snow total (mm), and average wind direction (angle). The sensors and meteorological station at the test bed record the sensor values every 15 minutes. See Ref. 2 for more details regarding the construction of the test bed and sensor placement.

The environmental data (including meteorological) are considered to be independent and identically distributed (IID) data. The IID interpretation of the environmental data simplifies the analysis since we are assuming that only the environmental conditions at the time of imaging effect the resulting thermal image. This study focuses on a subset of five environmental conditions: the volumetric water content (VWC), incoming solar radiation, average air temperature, average air pressure, and relative humidity. A previous study (Ref. 7 Sec. 6 and Ref. 8) found that these five environmental conditions have the largest impact on buried object visibility when using thermal IR sensors for detection. Note that the FLIR imagery is recorded on five-minute intervals while the environmental conditions are recorded on 15-minute intervals. To align the full FLIR image data set with the environmental data, we approximate five-minute resolution environmental data by taking a simple average for the considered subset. For example, on 1 July 2018 at 0900 the air temperature was 25.98 (C) and at 0915 it increased to 26.65 (C). These values are averaged to obtain the five-minute resolution for 0900, 0905, and 0910 to be given by 26.3 (C). For a more detail description on the preparation of the environmental data and selection of the most influential variables in buried object visibility see Sec. 6 of Ref. 7 . In Sec. 4 , we will use this environmental data to inform the recognition of buried objects in the thermal imaging.

3. DETECTION OF BURIED OBJECTS: COMPUTER VISION APPROACH

The detection step of the algorithm is accomplished using a computer vision approach which produces a collection of ROIs for each image. The ROIs may or may not contain a buried object which is determined by processing the ROIs using the environmentally informed classification algorithm introduced in Sec. 4.

The computer vision algorithm takes in an U8 image that is converted from the raw I16 thermal radiance value acquired with the FLIR A310 sensor. The heterogeneous nature of the thermal signature of soil and the compounding environmental factors mean that the SNR for the object of interest can be less than two, which makes standard detection techniques difficult to implement. The computer vision (CV) approach utilizes

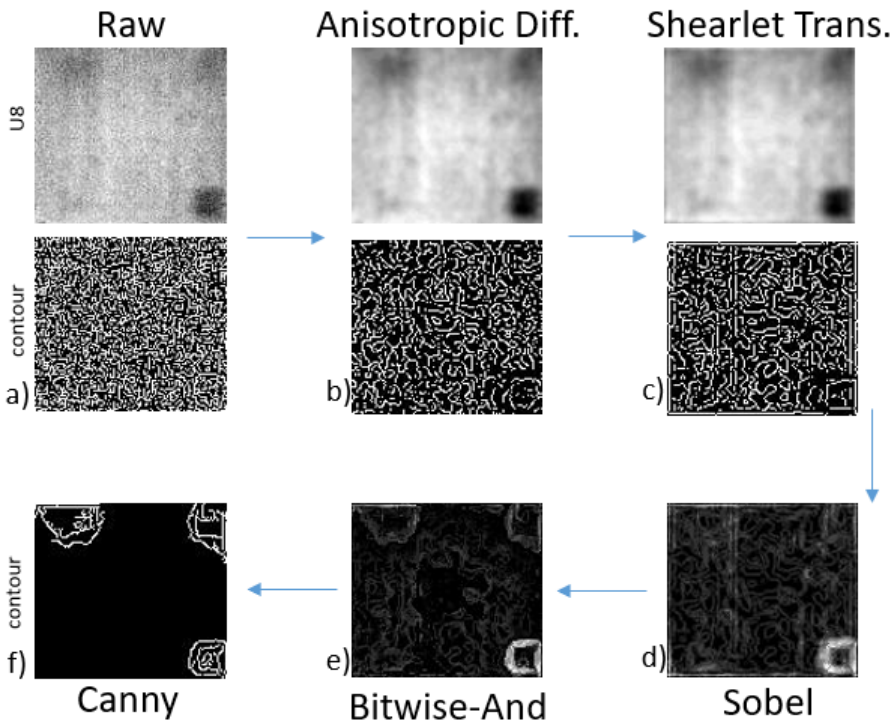

Figure 3. The stages of the computer vision baseline edge detection algorithm with each stage enhancing the SNR.

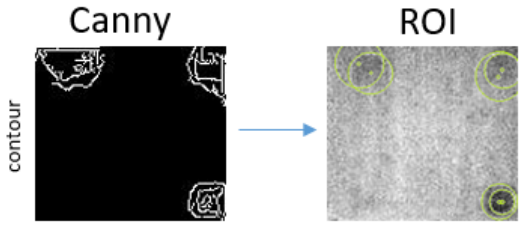

Figure 4. The final step of the CV baseline algorithm produces the ROIs from the Canny edge detection step.

filtering and transformation techniques to enhance the SNR for optical detection which localizes potential objects for classification. The CV detection algorithm takes in the raw U8 image, which is then smoothed using an anisotropic diffusion filter to reduce background noise The smoothed data is then processed using a Shearlet transformation to filter out the high frequency noise. Then to amplify the edges a Sobel operator is convolved with the anisotropic diffused and shearlet transformed image. Lastly, the processed image is convolved with a Bitwise-And operator to increase the SNR and Canny edge detection is applied to generate the ROIs. We call this procedure the CV baseline algorithm. Fig. 3 shows the progression of the CV algorithm from the anisotropic diffusion filter to the Canny edge detection. The goal of the algorithm is to enhance the SNR at each stage to improve the detection of buried objects in a thermal image. The contour plots in Fig. 3 are produced using the Canny edge detection algorithm.[9] Note that Canny edge detection is used in two different ways: without threshold limits to create the contour plots and with threshold limits to identify edge that may correspond to ROIs (Fig. 3f). Following the Canny edge detection step (with thresholds), a probabilistic Hough transform[10] is applied to produce the ROIs which are used to identify buried objects, see Fig. 4. The steps of the CV baseline algorithm are described in more detail in the following sections.

3.1 Smoothing and Noise Reduction

Since the thermal images used to detect buried objects can have a very low SNR, we need to filter the images to smooth and filter the noise levels. Using Canny edge detection, we produce contour plots of the images to empirically evaluate the noise level in the images. To see why the filtering steps are necessary, the obscuration of the signal can be observed in Fig. 3 a which shows the unprocessed raw thermal image and the corresponding contour representation and we can see that the signal of the buried objects is difficult to discern from the background noise. We mitigate this problem by smoothing the image and filtering out salt and pepper noise.

The image is smoothed by applying an anisotropic diffusion transformation which is inspired from biomedical imaging techniques that involve minimizing the energy function using gradient descent to reduce noise in an image.[11] The nonlinear and space-variant transformation technique was originally introduced by Perona and Malik in Ref. 9 as a way to reduce noise while maintaining edge information in the image. This approach is advantageous for our application since we are relying on using edges in a thermal image as an indicator of a surface temperature anomaly due to the presence of a buried object. For a greyscale image, $I(x, y, t)$, where (x, y) represent the image pixel locations and t is the variance, then the anisotropic diffusion is defined as

$$\frac{\partial}{\partial t} I(x, y, t) = \nabla \cdot (c(x, y, t) \nabla I) \,. \tag{1}$$

Above, c is the diffusion coefficient and if c is taken to be a constant then (1) reduces to the isotropic diffusion equation. The diffusion coefficient is taken to be a monotonically decreasing function, $g(\cdot)$, of the magnitude of the gradient of the image. For example, Perona and Malik in Ref. 9 define the diffusion coefficient to be given by a Gaussian,

$$c(x, y, t) = g(\|I\|) = e^{-(\|I\|/K)}, \tag{2}$$

or, to be given by the following function,

$$c(x, y, t) = g(\|I\|) = \frac{1}{1 + (\|I\|/K)^2}. \tag{3}$$

If the diffusion coefficient is taken to be (2), then the application of the anisotropic diffusion privileges high-contrast edges over low-contrast ones. On the other hand, if the diffusion coefficient is given by (3), then the anisotropic diffusion weighs wide-regions over narrow. In the diffusion coefficient, the constant K controls the sensitivity and we take it to be given by twice the global standard deviation, $K = 2\sigma_{I_o}$.[12] In smoothing a thermal image using anisotropic diffusion, the Gaussian diffusion coefficient is applied in non-trivial regions while edges are left un-altered and thus, maintaining their integrity. An edge-estimator, $E(x, y, t)$ is used to approximate the location of edges to inform the filtering and it is given by the gradient of the image, $E(x, y, t) = \nabla I(x, y, t)$. We use a Python implementation of anisotropic diffusion found in Ref. 13. Next, it is necessary to denoise the image by using a high-pass filter.

The Shearlet transformation uses the frequency domain to filter the thermal image for high frequency noise.[14] Comparing the contour plots in Fig. 3 b and Fig. 3 c, we see that the Shearlet transformation reduces the noise that causes speckling in the contour plots which obscure the buried object signal. The implementation of the Shearlet transformation used in this project is ShearLab3D.[15] For our purposes, the 2-dimensional Shearlet transformation is sufficient since the thermal images are naturally represented in grey scale since the raw image values correspond to radiance counts. The algorithm begins with using a parabolic scaling matrix which is followed by translation and shearing operations. The Shearlet coefficients are obtained with a point-wise multiplication of the Shearlets with the image in the frequency domain. We then threshold the Shearlet coefficients such that coefficients below a specific threshold are set to zero– this denoises the Shearlet transformation. Lastly, we invert the Shearlet transformation to obtain the filtered image data in the spatial-domain rather than the frequency domain. Fig. 5 shows a comparison of an unfiltered thermal image with the high frequency noise filtered out via the Shearlet transformation. The corresponding contour plots produced using Canny edge detection show that the noise reduction using the Shearlet transformation highlights the signal of the buried objects, particularly the deep objects, from the background noise. The next section discusses the steps used to produce the ROIs used in the classification algorithm.

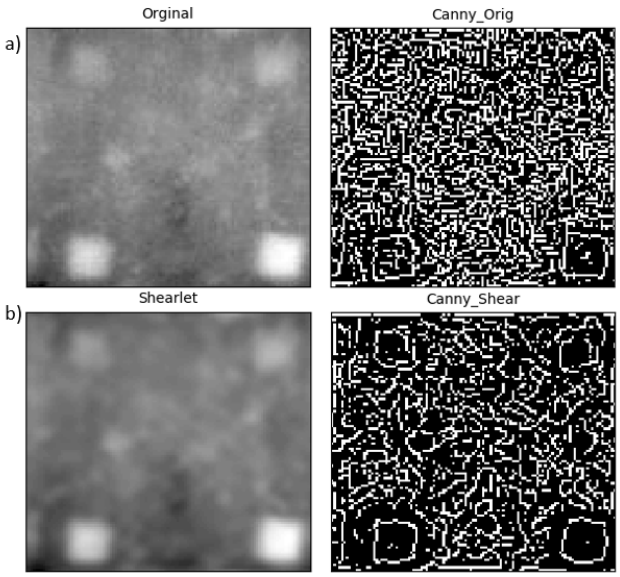

Figure 5. A comparison of the thermal image with the high frequency noise filtered out using the Shearlet transformation along with the corresponding contour plots obtained with Canny edge detection.

3.2 Edge Detection and ROIs

In the previous section, we discussed how we denoised the thermal images to enhance the buried object signal through smoothing and using a high pass filter via the Shearlet transform. Now that the edge-features in the images have been highlighted through the transform process, we need to separate the remaining background noise from the buried objects and this is done via convolution. In particular, we use the Sobel transform and the Bitwise-And operation to combine the filtered images. Bitwise operations are arithmetic operators, i.e. $\vee, \wedge, \otimes, \neg$, are conducted often throughout the CV baseline algorithm. These operations calculate the per-element bitwise conjunction or disjunction of two or more images, depending on the logical argument executed. Lastly, we use the Canny edge detection once again paired with the Hough transform[10] to select ROIs to detect the buried objects.

The Sobel operator for edge detection uses two 3×3 kernels convolved with the image intensity to approximate the first derivatives. In particular, we approximate the horizontal and vertical changes where sharp changes in the image (i.e. edges) corresponding to large jumps in the values of these derivative approximations. Again, we let I represent the image intensity and we approximate the horizontal and vertical first derivatives as follows:

$$G_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} * I, \qquad G_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix} * I, \tag{4}$$

where $*$ is denoting the two-dimensional convolution. The gradient approximations at each point of the image can then be combined using the magnitude to obtain

$$G = \sqrt{G_x^2 + G_y^2} \tag{5}$$

which can be computed at each point of the image. Places where an abrupt change occurs due to an edge in the image will correspond to a larger gradient magnitude compared to background noise. We will use the convention that an edge feature corresponds to a local maximum of the gradient magnitude. After using the Sobel operator to create a gradient magnitude mapping of the image intensity, we use the Bitwise-And operator to combine the result of the Sobel transformation with the Shearlet image. See Fig. 3d for the result of the Sobel operation and Fig. 3e for the combined image obtained from the Bitwise-And operation. The Canny edge detection algorithm is applied to the filtered image obtained in the Bitwise-And operation step of the CV baseline algorithm to

identify the potential locations of buried objects in the image. Finally, the probabilistic Hough transform is used to detect areas containing straight lines after the Canny edge detection to produce the ROIs, see Fig. 4. ROIs deemed to possibly contain a buried object must satisfy the following condition

$$\text{mean}(I) - \text{mean}(\text{ROI}) > \alpha(\sigma(I)), \tag{6}$$

where $\alpha \in (0, 1]$ is a sensitivity threshold constant that can be tuned for sensitivity and $\sigma(I)$ is the standard deviation of the image. If $\alpha = 1$, then the algorithm will produce fewer ROIs which will ultimately reduce false positives, but at the cost of greatly increasing false negatives. Since we can mitigate the false positives, it is better to use a sensitivity threshold that produces more false positives while reducing false negatives. For formulating the training data set in the next section we use $\alpha = 0.5$ to produce sufficient ROIs that balances the false positives with the false negatives. For more information about the CV baseline algorithm see Ref. 7 Sec. 5.

Note that the detection of buried objects using the CV baseline algorithm can produce both false positives and false negatives. The reduction of false negatives requires tuning the parameters of the algorithm, while the false positives can be filtered out using the environmentally informed classification of ROIs. The classification step improves the false alarm rate, but does not improve the false negative rate and improving the false negative rate of the CV baseline algorithm for detection is a topic for future research.

4. ENVIRONMENTALLY INFORMED CLASSIFIER

In this section, we will classify the ROIs produced using the CV approach to determine whether a buried object is present in a given ROI. The ROIs produced by the CV algorithm will include both true positives and false positives and it is important to filter out the false positives to ensure the detection of buried objects is reliable. Note that the CV algorithm will also have false negatives, meaning, some buried objects may not be detected in an ROI. The classification step of the algorithm does not mitigate the presence of false negatives. Since it is known that the environmental conditions at the time of imaging can affect the visibility and detection of buried objects,[1,2,6,7] we couple the environmental conditions with the ROIs to better inform the classification of the ROI. The approach used here follows the work by Gessert et al. in Ref. 16, where the authors used tabular data to help inform the classification of images in a medical imaging application.

4.1 Training and Testing Data

Given the ROIs produced with the CV baseline algorithm described in Sec. 3, we need to create labels for the ROI data set and formulate a test set. The ROIs are selected to be rectangular bounding boxes that surround the circular regions identified by the Hough transform step of the CV baseline algorithm and the output is the top left pixel location $(x_{\text{TL}}, y_{\text{TL}})$ with the corresponding width and height of the bounding box. Given the top left pixel locations,

$$x_{\text{TL}} \in [0, 115) \subset \mathbb{Z}, \quad y_{\text{TL}} \in [0, 105) \subset \mathbb{Z},$$

and the height h and the width w of the bounding box, we take $w_{\text{sq}} = \max(h, w)$ to produce a square ROI for a more uniform training and testing data set. Then, the ROI is defined to be a 2-D pixel array subset given by

$$\text{ROI} = [x_{\text{TL}}, x_{\text{TL}} + w_{\text{sq}}] \times [y_{\text{TL}}, y_{\text{TL}} + w_{\text{sq}}] \subset \mathbb{Z}^2.$$

Recall, the top left corner of the test bed image corresponds to pixel location $(0, 0)$ and the bottom right corner of the test bed corresponds to pixel location $(115, 105)$. Note that if $x_{\text{TL}} + w_{\text{sq}} > 115$, then truncate the width such that $x_{TL} + w = 115$, and similarly, we truncate the height if $y_{\text{TL}} + w_{\text{sq}}$ exceeds 105. Using the pixel locations for each ROI we then obtain sub-images by cropping the associated full 115×105 pixel image to the appropriate ROI size. The cropped ROI images are square with varying dimensions, except in the few cases the height or width of the ROI is truncated to not exceed the boundaries of the image. To create a uniform training and testing set, we transform the ROIs to have the size 50×50 pixels.

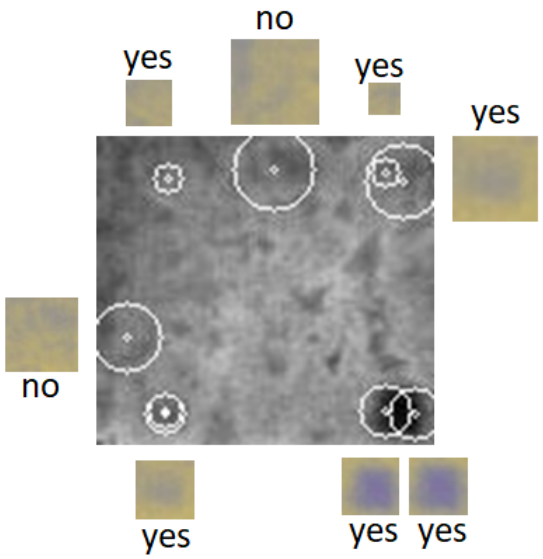

Figure 6. An example thermal image processed with the CV algorithm and the resulting ROIs. Each ROI is labeled "yes" of $\ell(\text{ROI}) = 1$ and labeled "no" if $\ell(\text{ROI}) = 0$. This image was captured 19 September 2018 at 21:45.

To create labels for the training and testing data, we utilize the fact that we know the pixel locations of the buried objects in the thermal images. For each ROI, we create a label, ℓ, as follows:

$$\ell(\text{ROI}) = \begin{cases} 1 & \text{if } \text{ROI} \cap \omega \neq \emptyset, \text{ for all } \omega \in \Omega \\ 0 & \text{if } \text{ROI} \cap \omega = \emptyset, \text{ for } \omega \in \Omega \end{cases} \tag{7}$$

where ROI is the two-dimensional array of pixel locations (prior to transformation), as described above, and Ω is a two-dimensional array of pixels that correspond to a particular object. Recall, for the data set used in this study, there are four buried objects, and we take ω to be an element of the set of four objects Ω:

$$\omega \in \{\text{SP}, \text{SM}, \text{DM}, \text{DP}\} = \Omega,$$

where $\text{SP} = [90, 110] \times [85, 105]$, $\text{SM} = [10, 30] \times [85, 105]$, $\text{DM} = [90, 110] \times [5, 25]$, and $\text{DP} = [10, 30] \times [5, 25]$. Fig. 6 shows an example thermal image processed using the CV algorithm to produce a collection of ROIs with the corresponding computed labels. In this example, all four buried objects were identified by ROIs via the CV algorithm, however, there are two ROIs identified that correspond to false positives. Note that for a given thermal image, the CV algorithm does not necessarily identify all buried objects. There are cases when the identified ROIs neglect to identify an object which would correspond to a false negative.

Lastly, to prepare the dataset for the environmentally informed classification algorithm, we need to collect the environmental metadata for each ROI. Each image is associated with a timestamp, thus each ROI inherits a timestamp from the original thermal image which can be matched with the timestamp for the environmental data. As noted in Sec. 2, previous studies have identified that the environmental conditions that have the greatest impact on object visibility are the soil moisture content (%), average air temperature ($^\circ C$), incoming shortwave solar radiation (W/m^2), average relative humidity (%), and average atmospheric pressure (mbar). For the ROIs in the sample image shown in Fig. 6, the corresponding environmental metadata is shown in Tab. 1.

To create a more balanced training and testing data set, we processed the full 2018 data set which includes images collected prior to object emplacement from 11 May to 11 July 2018 and the after emplacement data from 11 July to 15 November 2018. The combined data consists of 51,240 thermal images processed in the CV algorithm which produced 271,673 ROIs. To create a testing data set sufficiently distinct from the training data, we group the ROIs from a single day to all be included in either the training or testing data set. Meaning, for the 186 days of imaging, 131 days are randomly selected to formulate the training data set and the remaining

Table 1. Environmental metadata corresponding to the example thermal image in Fig. 6.

Environmental Variable	Value for 19 Sept. 2018 at 21:45.
Soil Moisture Content (%)	0.164462
Average Air Temperature (°C)	11.99
Incoming Solar Radiation (W/m^2)	0
Average Relative Humidity (%)	92.55
Average Atmospheric Pressure (mbar)	1002.5

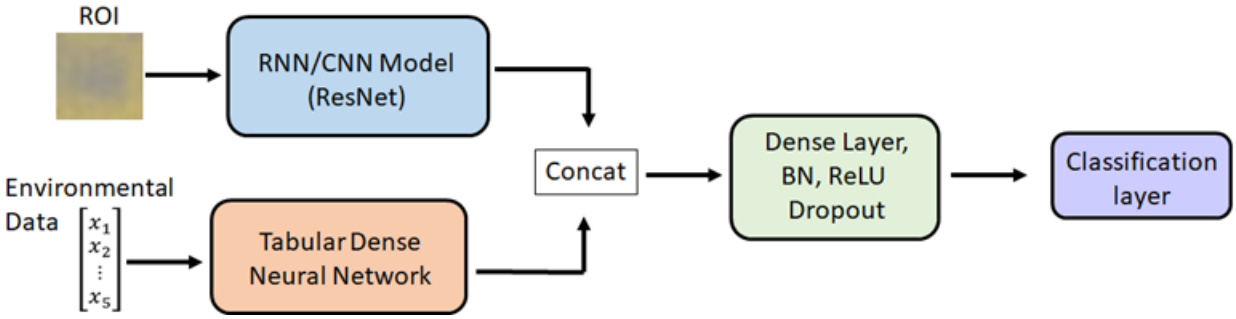

Figure 7. The architecture for the environmentally informed classification algorithm used to process and classify the ROIs produced by the CV algorithm. The architecture is directly inspired by the work in Ref. 16.

55 days formulate the testing data set. Any ROI corresponding to a day that is assigned to the training set will be assigned to the ROI training set and similarly for the testing set. We randomly select the days rather than timestamps or ROIs for the training versus testing sets to ensure the classifier is being trained on distinct data from the testing set. Otherwise, if the timestamps are randomly selected, it would be possible, for example, for ROIs imaged on 19 Sept. 2018 at 21:45 to be in the training data set while the ROIs imaged five minutes later on 19 Sept. 2018 at 21:50 be in the testing data set. In this case, the classifier would be trained on data very similar to the testing set thus resulting in an effective performance evaluation of the classifier on unseen data. Finally, a small portion of the training data set is set aside for a validation set to evaluate the classifier network during training.

4.2 Classification of ROIs

To classify the ROIs, we utilize an environmentally informed neural network which concurrently learns from the ROI image data and corresponding environmental tabular data. The architecture of the classification network is directly inspired by the work of Gessert et. al in Ref. 16. Fig. 7 shows a schematic representation of the environmentally informed classification network. By processing the ROIs produced by the CV algorithm discussed in Sec. 3, we are able to increase the accuracy of the two-step detection algorithm by filtering out false positives. Recall, the false negative rate can only be reduced through further tuning of parameters in the CV algorithm used to produce the ROIs or by investigating alternative approaches to generating the ROIs.

The network takes in an ROI, which has been transformed to a uniform size of 50×50 pixels, and the corresponding environmental conditions represented by a vector $x \in \mathbb{R}^5$. The image data is first processed with an image classification algorithm and for this work we use ResNet34[17] pre-trained on the ImageNet[18] data set. The environmentally informed classification algorithm is implemented using the Fastai (version 1) deep learning library[19, 20] and PyTorch.[21] The network is trained using a cross-entropy loss function and we use the accuracy and area under the receiver operator characteristic (ROC) curve (AUC) metrics to assess the performance. The fastai framework enables simple implementation of transfer learning to take advantage of pre-trained networks, such as ResNet, to customize the optimized weights learned through back propagation and minimization of the loss function to the specific application and, in this case, which is to classify the ROIs. Note that the ROIs in our training and testing data set can contain one of the four buried objects in the test bed, however, to simplify the model we simply identify the ROI as containing an object or not and we do not specify the type of object. For our purposes, the last two layers of the ResNet model are left off to allow for classification using the combined

Table 2. A comparison of the training and testing errors for image-only and the environmentally informed ROI classification.

Classification Model	Training		Testing	
	AUC	Accuracy	AUC	Accuracy
ROI Only	0.901	0.841	0.510	0.448
ROI and Environmental	0.934	0.836	0.922	0.828

ROI and environmental data. The output of the image network is a 100-dimensional vector. The tabular data is processed through a neural network given by the fastai TabularModel. This model is a sequential model that consists of five l ayers: 1) a l inear a ctivation l ayer w hich t akes i n t he fi ve-dimensional ve ctor of environmental variables and outputs an eight-dimensional vector, followed by 2) a layer that passes through a rectified linear unit (ReLu) activation, 3) batch normalization, 4) a dropout layer, and 5) finally another linear activation layer. The final o utput o f t he t abular n etwork is a n e ight-dimensional v ector. A fter t he R OIs a nd e nvironmental data are processed through the image network and the tabular network, respectively, the outputs are concatenated to produce a vector of size 108 combining the results of the image and tabular networks. The concatenated results are passed through five f ully c onnected l ayers w hich d ecrease t he o utput s ize i n e ach l ayer t o ultimately produce the classification. T he fi nal cl assification out put has dim ension two and the firs t valu e indi cates the probability that the ROI contains a buried object and the second value indicates the probability that the ROI was a false positive. To produce the final c lassification pr ediction, we se lect th e gr eater pr obability fr om the two-dimensional output from the final c lassification layer.

We compare the performance of the environmentally informed classification o f t he R OIs t o t he classification using only the ROI image data and ResNet34. In the image-only classification, w e u se t he s ame f ramework as in the environmentally informed classification except we use t ransfer l earning o n t he full R esNet m odel a nd use the classification l ayer i n R esNet t o p roduce o ur p redictions. For t he c omparison, w e u se t he s ame t raining set with the same validation subset and testing set for both classification a lgorithms. O n t he t esting d ata s et the image-only classification achieved an AUC of 0 .51 w hile t he environmentally informed c lassification ac hieved an AUC of 0.92, see Tab. 2. Notably, the environmentally informed classification d ecreases t he o ver-fitting th at is observed with the image-only classification w hich c an b e seen i n Tab. 2 w hen comparing t he A UC and accuracy metrics on the training data set versus the testing data set. Fig. 8 shows a comparison of the ROC curve computed on the testing data set where the false positive rate is plotted against the true positive rate. Overall, the environmentally informed classification of the ROIs outperforms the image-only ROI classification.

Lastly, we compare the distribution of the environmental conditions when the environmentally informed classification predicts the correct label versus the incorrect label. Note that the incorrect label distributions combine the two possible situations: the correct ROI label is $\ell(\text{ROI}) = 1$ but the classification predicted a label of 0 and, for the opposite when the correct ROI label is $\ell(\text{ROI}) = 0$ while the classification predicted a label of 1. The distribution of the VWC varied the most between the correct and incorrect predicted labels, see Fig 9. The lower the value of the VWC the more likely the environmentally informed classification algorithm predicts the correct label for the ROI. On the other hand, at higher values of VWC there is an increased probability of the classification algorithm selecting the incorrect label for the test data set. This is consistent with findings in previous studies[1, 2, 6, 7] conducted at CRREL to understand the impact of environmental conditions on the visibility of buried objects using thermal IR. In particular, as the VWC increases after a rain event the increased soil moisture acts as a thermal mask for the buried objects, which makes detection with thermal IR difficult. The other four environmental condition distributions were also analyzed. For air temperature, the closer the average air temperature was to freezing the more likely the classification algorithm predicts the incorrect label. The distribution comparison for incoming short wave solar radiation, average pressure, and relative humidity did not illuminate any strong patterns for environmental conditions corresponding to an incorrect label prediction using the environmentally informed classification algorithm.

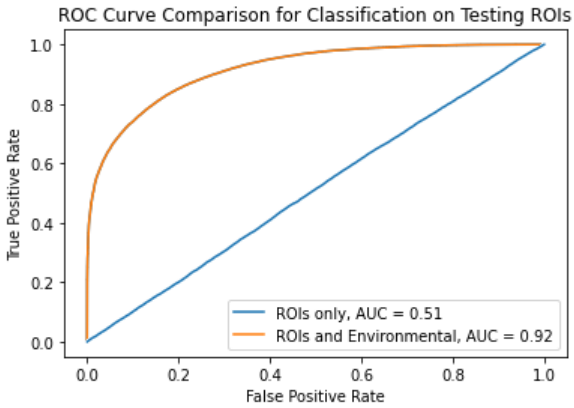

Figure 8. A comparison of the ROC curves for the classification of ROIs using only images versus the coupling of the ROIs with environmental data.

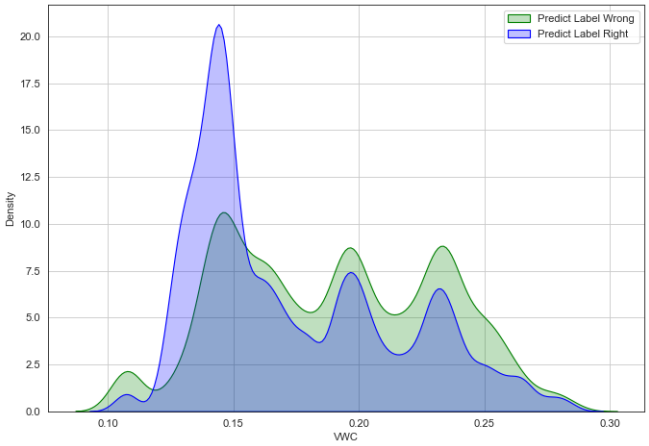

Figure 9. A comparison of the distribution of the VWC associated with the test ROIs where the green distribution corresponds to test ROIs assigned the correct label and the blue distribution corresponds to test ROIs assigned the incorrect label using the environmentally informed classification algorithm.

5. CONCLUSION

This paper introduced a two-step algorithm to detect and classify buried objects using thermal IR sensors. The detection step of the two-step algorithm was implemented using a CV approach to increase the SNR and used edge detection to identify regions of interest within the thermal imagery. The CV approach to produce ROIs results in true positives, false positives, false negatives, and true negatives, which correspond to the detection of a buried object, detection of background noise as an object, object signal is identified as background noise, and background noise that is identified correctly, respectively. Both the true positives and false positives result in the identification of a ROI within the thermal image. The true negatives and false negatives both correspond to the situation where a ROI is not selected by the algorithm. The reduction of the false negative rate in the CV approach for detection of potential buried objects is a topic of future research. However, the false positive rate can be mitigated by the second-step of the algorithm, which classifies the ROIs as containing a buried object or not. We introduced an environmentally informed classification algorithm, which simultaneously learns on the environmental data and the ROIs improve the performance of the classification algorithm. We compared the environmentally informed approach to the image-only approach in which the ROIs are classified using transfer learning on a pre-trained ResNet34 algorithm. The inclusion of environmental conditions greatly improved the performance of the ROI classification algorithm. Future research for the classification algorithm involves training

the algorithm on a diverse data set that includes buried objects of varying shape and dimension and a wider variety of weather conditions.

ACKNOWLEDGMENTS

Permission to publish was granted by Director, Cold Regions Research and Engineering Laboratory. Distribution Statement A: approved for public release; distribution is unlimited. This research was funded by the U.S. Army Combat Capabilities Development Command under PE 0602712A/Project 468349/Task A1010, Analysis of Multi-Modal Sensors. This research was supported in part by an appointment to the Department of Defense (DOD) Research Participation Program administered by the Oak Ridge Institute for Science and Education (ORISE) through an interagency agreement between the U.S. Department of Energy (DOE) and the DOD. ORISE is managed by ORAU under DOE contract number DE-SC0014664. All opinions expressed in this paper are the author's and do not necessarily reflect the policies and views of DOD, DOE, or ORAU/ORISE.

REFERENCES

[1] Clausen, J., Dorvee, J., Wagner, A., Frankenstein, S., Morriss, B., Claffey, K., Sobecki, T., Williams, C., Newman, S., Booker, B., Affleck, R., Smith, C., Maxson, M., Bernier, A., and Jones, B., "Spatial temporal variance in the thermal response of buried objects," Tech. Rep. ERDC/CRREL TR 20-10, U.S. Army Corps of Engineers, Engineering Research and Development Center, Cold Regions Research and Engineering Laboratory, Hanover, NH (2020).

[2] Clausen, J., Frankenstein, S., Dorvee, J., Workman, A., Morriss, B., Claffey, K., Sobecki, T., Williams, C., Newman, S., Booker, B., Affleck, R., Smith, C., Maxson, M., Bernier, A., and Jones, B., "Spatial temporal variance of soil and meteorological properties affecting sensor performance - phase II," Tech. Rep. ERDC/CRREL TR 21-10, U.S. Army Corps of Engineers, Engineering Research and Development Center, Cold Regions Research and Engineering Laboratory, Hanover, NH (2021).

[3] Frankenstein, S., Wagner, A., and Clausen, J., "Effects of soil properties, emplacement depth, and object composition on thermal signature," in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XXV*], *Proc. of SPIE* **11418** (2020).

[4] Clausen, J., Musty, M., Wagner, A., Frankenstein, S., and Dorvee, J., "Modeling of a multi-month thermal ir study," in [*Chemical, Biological, Radiological, Nuclear, and Explosives (CBRNE) Sensing XXI*], *Proc. of SPIE* **11416** (2020).

[5] Schachter, B., [*Automatic Target Recognition*], SPIE, Bellingham, WA, fourth ed. (2020).

[6] Clausen, J., Felt, C., Musty, M., Truong, V., Frankenstein, S., Wagner, A., Affleck, R., Peckham, S., and Williams, C., "Spatial and temporal variance of soil and meteorological properties affecting sensor performance - phase III," Tech. Rep. *In Press*, U.S. Army Corps of Engineering, Engineering Research and Development Center, Cold Regions Research and Engineering Laboratory, Hanover, NH (2022).

[7] Clausen, J., Truong, V., Bragdon, S., Frankenstein, S., Wagner, A., Affleck, R., and Williams, C., "Buried object detection improvements incorporating environmental phenomenology into signature physics," Tech. Rep. *In Review*, U.S. Army Corps of Engineering, Engineering Research and Development Center, Cold Regions Research and Engineering Laboratory, Hanover, NH (2022).

[8] Bragdon, S., Truong, V., and Clausen, J., "Environmental effects on contrast metrics using thermal IR sensors," in [*MSS Virtual Parallel (BSD, Materials and Detectors, and Passive Sensors) Conference*], *Proc. of MSS* (2022).

[9] Perona, P. and Malik, J., "Scale-space and edge detection using anisotropic diffusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12**(7), 629–639 (1990).

[10] Matas, J., Galambos, C., and Kittler, J., "Robust detection of lines using the progressive probabilistic hough transform," *Computer Vision and Image Understanding* **78**(1), 485–497 (2000).

[11] Zhu, Y., Yu, X., and Zhang, B., "A nonlinear diffusion model for image restoration," *Acta Math. Appl. Sin. Engl. Ser.* **32**, 631–646 (2016).

[12] Kovesi, P., "Anisotripic diffusion MATLAB code." School of Computer Science and Software Engineering, The University of Western Australia (2000).

[13] Muldal, A., "Anisotropic diffusion python code." Department of Pharmacology, University of Oxford (2012). Accessed: 2021.

[14] Price, S., Anderson, D., and Keller, J., "Design of a buried explosive hazard pre-screener in forward looking imagery based on shearlet filtering and image post-processing," in [*Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XX*], *Proc. SPIE* **9454** (2015).

[15] Kutyniok, G., Lim, W.-Q., and Reisenhofer, R., "ShearLab 3D: Faithful digital shearlet transforms based on compactly supported shearlets," *ACM Transactions on Mathematical Software* **42**(5) (2016).

[16] Gessert, N., Nielsen, M., Shaikh, M., Werner, R., and Schlaefer, A., "Skin lesion classification using ensembles of multi-resolution efficientnets with meta data," *MethodsX* **7**, 100864 (2020).

[17] He, K., Zhang, X., Ren, S., and Sun, J., "Deep residual learning for image recognition," in [*Conference on Computer Vision and Pattern Recognition*], *Proc. of the IEEE*, 770–778 (2016).

[18] Stanford Vision Lab, "Imagenet." https://www.image-net.org (2021). Accessed: 2021-05-01.

[19] Howard, J. and Gugger, S., "Fastai: A layered API for deep learning," *Information* **11** (2020).

[20] Howard, J., "fastai V1." https://fasta1.fast.ai (2020).

[21] PyTorch, "Pytorch." https://pytorch.org (2021).

REPORT DOCUMENTATION PAGE

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
		START DATE	END DATE
November 2022	Final	FY20	FY21

4. TITLE AND SUBTITLE
Environmentally Informed Buried Object Recognition

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT
		0602712A
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**
468349	A1010	

6. AUTHOR(S)
Sophia Potoczak Bragdon, Vuong Truong, and Jay Clausen

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
US Army Engineer Research and Development Center (ERDC)
Cold Regions Research and Engineering Laboratory
72 Lyme Road
Hanover, NH 03755

8. PERFORMING ORGANIZATION REPORT NUMBER
ERDC/CRREL MP-22-17

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
US Army Combat Capabilities Command
Night Vision and Electronic Sensors Directoirate (NVESD)
10221 Burbeck Road
Ft. Belvoir, VA 22060-5806

10. SPONSOR/MONITOR'S ACRONYM(S)
NVESD

ERDC/CRREL
MP-22-17

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
This document was originally published with SPIE Conference Proceedings on 30 May 2022.

14. ABSTRACT

The ability to detect and classify buried objects using thermal infrared imaging is affected by the environmental conditions at the time of imaging, which leads to an inconsistent probability of detection. For example, periods of dense overcast or recent precipitation events result in the suppression of the soil temperature difference between the buried object and soil, thus preventing detection. This work introduces an environmentally informed framework to reduce the false alarm rate in the classification of regions of interest (ROIs) in thermal IR images containing buried objects. Using a dataset that consists of thermal images containing buried objects paired with the corresponding environmental and meteorological conditions, we employ a machine learning approach to determine which environmental conditions are the most impactful on the visibility of the buried objects. We find the key environmental conditions include incoming short-wave solar radiation, soil volumetric water content, and average air temperature. For each image, ROIs are computed using a computer vision approach and these ROIs are coupled with the most important environmental conditions to form the input for the classification algorithm. The environmentally informed classification algorithm produces a decision on whether the ROI contains a buried object by simultaneously learning on the ROIs with a classification neural network and on the environmental data using a tabular neural network. On a given set of ROIs, we have shown that the environmentally informed classification approach improves the detection of buried objects within the ROIs.

15. SUBJECT TERMS
Buried object detection; Machine learning; Thermal imaging; Environmental phenomenology; Deep learning; LWIR

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	C. THIS PAGE		
Unclassified	Unclassified	Unclassified		18

19a. NAME OF RESPONSIBLE PERSON	19b. TELEPHONE NUMBER (include area code)
Jay Clausen	603.646.4597