

# Leveraging Environmental Phenomenology to Improve Buried Object Detection using Longwave Thermal Infrared Imaging

**Sophia Potoczak Bragdon, Vuong H. Truong, Jay L. Clausen**

Cold Regions Research and Engineering Laboratory  
Engineer Research and Development Center  
U.S. Army Corps of Engineers  
72 Lyme Road, Hanover, New Hampshire  
UNITED STATES

[Sophia.P.Bragdon@usace.army.mil](mailto:Sophia.P.Bragdon@usace.army.mil)

## **ABSTRACT**

*The ability to reliably detect and classify buried objects using longwave infrared (LWIR) imaging is a challenging military research topic. The LWIR sensor performance is directly impacted by the environmental conditions at the time of imaging. This research explores the complex relationship between the visibility of buried objects and the ambient environmental conditions. By using interpretable deep neural networks, we have identified major environmental factors that influence buried object visibility when using LWIR in warm humid continental and hot desert climate types. These environmental factors include soil moisture content, time of day, air temperature, relative humidity, and incoming shortwave solar radiation. We have incorporated the salient environmental features into a two-step automatic target recognition (ATR) algorithm to improve the detection of buried objects. The first step of the ATR algorithm produces a series of regions of interest using a computer vision approach. Next, we filter out the false positives in a second classification step, which is based on an environmentally informed deep learning framework. We leverage the fact that the environmental conditions are related to the visibility of the object from the background noise to better filter out the false positives. This is accomplished by coupling the regions of interest in the LWIR image with the environmental conditions and we process the hybrid data in a machine learning algorithm that combines a fully connected neural network for the environmental data with a convolutional neural network for the LWIR image data to produce a classification. When compared to simply using the images for distinguishing the false alarms from the true positives, the environmentally informed approach out-performs the image-only approach. We also compare our algorithm to other common machine learning-based ATR algorithms, such as DeepTarget and Faster R-CNN.*

## **1.0 INTRODUCTION**

Buried object detection using longwave infrared (LWIR) imaging is based on identifying surface thermal anomalies and having the ability to determine if the anomaly is due to background noise or the presence of a subsurface object. Utilizing thermal sensors for buried object detection allows one to collect imagery day and night, as well as detect objects of different materials and shapes (e.g., metal and plastic). However, the visibility of a buried object using LWIR is affected by the ambient weather and environmental conditions. For a buried object to be visible in a thermal image, there must be a difference in the thermal signature of the soil surface directly above the object when compared to the surrounding earth and the environmental conditions influence the thermal flux. For example, after a rain event the volumetric soil water content (VWC) will increase, and the increased water content will induce thermal equilibrium between the earth and the buried object. On day with high solar loading, the earth may heat unevenly as the soil warms from overnight low temperatures and due to the thermal noise induced by solar loading, it can be impossible to discern object-induced noise from background noise. In this paper, we will leverage the relationship of the environmental conditions to object visibility to better classify the false detections from true detections. Portions of this work have been described in [1], [2], [3], [8].

## 1.1 Dataset

This study uses data collected in Hanover, New Hampshire USA, which is considered to be a warm humid continental climate, and Yuma, Arizona USA, which is considered to be a hot desert climate. The New Hampshire dataset was collected in 2018 as part of a multi-year study conducted at the Engineer Research and Development Center (ERDC) Cold Regions Research and Engineering Laboratory (CRREL) and we will refer to this dataset as the CRREL dataset. For the CRREL dataset, data collection began in July 2018 and concluded in November 2018, the testbed consists of a 3 by 3 meter plot containing four buried objects and various buried sensors, see Clausen et al. 2020 [6]-[8] for a detailed description of the construction of the CRREL testbed. The buried objects consist of two plastic and two metal boxes (40 by 40 by 20 cm) and are filled with a nitrate fertilizer. One plastic and one metal object are each buried at 5 and 25 cm below ground surface, respectively. The four objects are referred to as the deep plastic (DP), deep metal (DM), shallow plastic (SP), and shallow metal (SM). A meteorological station is adjacent to the testbed to collect local weather conditions (precipitation, average air temperature, average air pressure, relative and relative humidity), radiometers are utilized to measure the incoming and reflected shortwave solar radiation ( $W/m^2$ ), and buried sensors record volumetric water content (VWC) of the soil (%), soil temperature, and soil conductivity. The Arizona dataset was collected in February 2020 at the Yuma Proving Ground (YPG) in Yuma, Arizona and we refer to this dataset as the YPG dataset. The testbed constructed at the Yuma Proving Ground utilized for the YPG dataset consists of four identical buried objects to that used in the CRREL dataset along with a mock-mine which is a circular disc plastic (CP) object buried at 15 cm below ground surface. As with the CRREL testbed, the YPG testbed is outfitted with buried sensors measuring the VWC, soil temperature, and conductivity, a meteorological station, and a set of radiometers measuring the solar radiation. For a detailed description of the testbed and data collection for the YPG dataset, see Musty et al. 2020 [16]. For both the YPG and the CRREL datasets, we image the testbed utilizing a FLIR A310 camera to capture longwave infrared (LWIR) imagery of the ground surface. The infrared cameras are oriented nadir to the ground surface. Figure 1 shows the CRREL testbed and a sample thermal image with the buried objects labelled. Figure 2 shows the YPG testbed and a sample thermal image with the buried objects labelled.

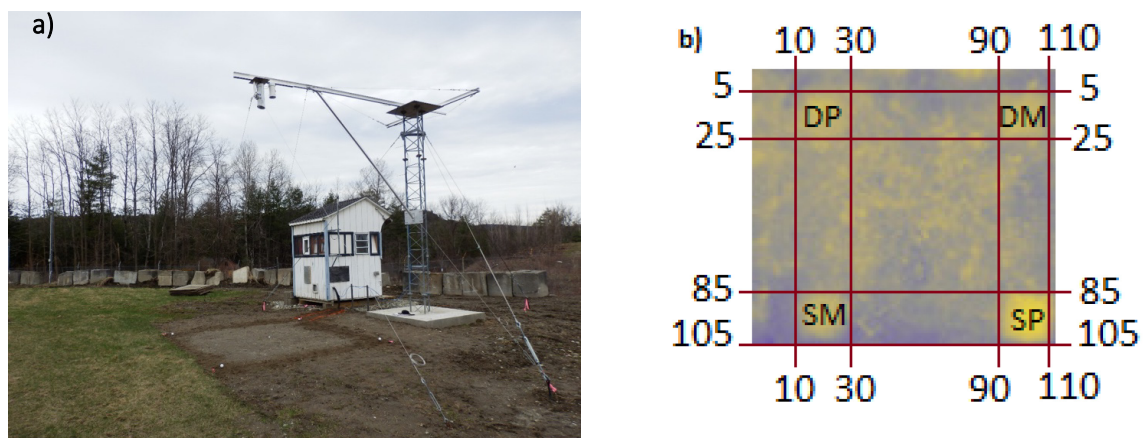


Figure 1: a) The CRREL test bed set-up, and b) a sample LWIR image from 9 September 2018 taken at 1420 that is a cropped to show the test bed where the four buried objects (SP, SM, DP, and DM) are labeled.




Figure 2: a) The YPG test bed set-up, and b) a sample LWIR image from 16 February 2020 taken at 2245 that is cropped to show the test bed where the five buried objects (SP, SM, DP, DM, and CP) are labeled.

## 2.0 ENVIRONMENTAL PHENOMENOLOGY AND OBJECT VISIBILITY

In this section, we study the relationship between object visibility and the environmental conditions to determine the environmental variables that have the greatest impact on object visibility utilizing longwave infrared imaging. To study this complex relationship, we first introduce an object visibility measure that we call *contrast metrics* which serve as a proxy for the visibility of the objects. Then, using a deep learning approach, we study the nonlinear relationship between the object contrast metrics and the environmental conditions. In particular, we train a regression model that has an input given by the environmental variables and outputs the object contrast metric. For this study, we consider a subset of 9 environmental variables that have been shown to impact the visibility of objects [5]-[8]. Specifically, we consider the air temperature (°C), air pressure (mbar), relative humidity (%), 15-minute rain total (mm), wind direction (°), average wind speed (m/s), VWC of the soil (%), incoming shortwave solar radiation (W/m<sup>2</sup>), and reflected shortwave solar radiation (W/m<sup>2</sup>).

### 2.1 Object Visibility: Contrast Metrics

This section introduces the object visibility proxy we use to study the sensitivity of the LWIR imagery and the environmental conditions. The so-called *contrast metrics* are defined in terms of the mean and standard deviation of the object-pixels related to background pixels in the LWIR imagery. We consider a set of four different contrast metrics which are combined into a single metric by projecting the four contrast metrics onto the principal component via principal component analysis. The first metric we consider is called the Michelson contrast metric,  $C_m$ , and it is defined as [15]

$$M_1 = C_m = \frac{\mu_o - \mu_b}{\mu_o + \mu_b} \quad (1)$$

where,  $C_m$  denotes the Michelson contrast metric,  $\mu_o$  the mean value of object region of interest (ROI), which can be as small as one pixel, and  $\mu_b$  the mean value of the background close to the object. The Michelson contrast metric is commonly used in computer vision to assess the level of contrast between different areas of an image, and, in our case, we use it to measure the difference between the object and background pixels. The closer the contrast metric is to zero indicates the difference between the object and the background is minimal making it difficult to distinguish the buried object from the noise. The second contrast metric is the adjusted Michelson contrast metric, which is obtained by normalizing the denominator of equation (1) to enhance the difference between the object and background means. Let  $m$  be the minimum pixel value in the area of interest (AOI) that consists of both the object and the background. Then, the adjusted Michelson contrast metric is given by

$$M_2 = C_{ma} = \frac{\mu_o - \mu_b}{\mu_o + \mu_b - 2m}. \quad (2)$$

Both the Michelson and the adjusted Michelson contrast metrics take on values in the interval  $[-1, 1]$ , although, the Michelson metric,  $M_1$ , is typically much smaller than the adjusted metric,  $M_2$ . Next, we consider two metrics that are versions of the signal to noise ratio [20], [21]. The metrics are given by

$$M_3 = T_{10} = \frac{\mu_o - \mu_b}{\sigma_i}, \quad (3)$$

$$M_4 = T_{10a} = \frac{\mu_o - \mu_b}{\sigma_b}, \quad (4)$$

where  $\sigma_i$  the image standard deviation, and  $\sigma_b$  the background standard deviation. We compute the four contrast metrics for each of the buried objects in the considered dataset. Thus, for the CRREL dataset, for each image, we compute in total 16 contrast metrics, given by  $M_1, M_2, M_3$ , and  $M_4$  for the DP, DM, SP, and SM objects. For the YPG dataset, for each image, 20 contrast metrics are computed, the four metrics for DP, DM, SP, and CP.

To reduce the dimensionality of the object visibility metric associated with each image, we compute the principal component using PCA and project the four metrics for each object type onto the principal component which accounts for 95-98% of the variance in the set of contrast metrics for each object type in each dataset. For example, consider the set of contrast metrics for the CRREL dataset associated with the SP object, we can compute an aggregated metric using the principal component projection as follows:

$$M_{SP} = PCA(M_1, M_2, M_3, M_4)_{SP} := \sum_{j=1, \dots, 4} \frac{\alpha_j (M_j^{SP} - \mu_j(M_j^{SP}))}{\sigma_j(M_j^{SP})} \quad (5)$$

Where  $M_j^{SP}$  is the  $j^{th}$  contrast metric computed for the SP object where  $j = 1, \dots, 4$  and  $\mu_j(M_j^{SP})$  and  $\sigma_j(M_j^{SP})$  is the mean and the standard deviation, respectively, of  $M_j^{SP}$  computed over all images. For the CRREL dataset and the SP object, the projection coefficients for the principal component are given by:

$$\begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_4 \end{bmatrix} = \begin{bmatrix} -0.463219 \\ -0.513571 \\ -0.510056 \\ -0.511388 \end{bmatrix}.$$

Utilizing the PCA projection, we associate four and five contrast metrics with each image in the CRREL and YPG datasets, respectively. Note that in computing the principal component projection using equation (5), we normalize and center the contrast metric datasets to combine the four metrics in a manner that addresses the differing order of magnitude of the contrast metrics. In the next section, we use the contrast metrics computed with the principal component projection to study the relationship of the object visibility with respect to the environmental conditions.

## 2.2 Environmental Conditions to Predict Contrast Metrics

This section focuses on using a deep learning regression model to analyze the sensitivity of object visibility in relation to environmental phenomenology. The relationship is studied separately for the CRREL and YPG datasets due to the vastly different climatic conditions. The CRREL dataset was collected in a warm and humid climate where annual precipitation averages of 111.49 cm and the YPG dataset was collected in a hot arid desert climate with annual precipitation averages of 11.78 cm. In either case, we consider the same nine environmental conditions air temperature, air pressure, relative humidity, VWC, wind speed and direction, incoming and reflected shortwave solar radiation, and 15-minute rain total. For the YPG data collection event, precipitation only occurred once during set-up of the testbed and measured at 2 mm, thus, for YPG dataset, we do not include the 15-minute rain total. We train a regression model based on the TabNet [22] neural network, which utilizes sequential attention transformers to learn the relationship between the input variables (environmental) and the output variable (contrast metric). TabNet has an added feature of built-in feature importance thanks to the sequential attention portion of the architecture, and we make use of this feature to inform which environmental variables have the strongest impact on the prediction of the object visibility. In total, we train nine different TabNet regression models, one for each object type in each dataset with the PCA projection of the contrast metrics representing object visibility. For training and assessing the accuracy of the regression models, we split the CRREL and the YPG datasets into training (65%), validation (15%), and testing (20%). Table 1 and Table 2 show the testing errors for the nine TabNet models trained for this study. Figure 3 shows the feature importance for the SP object in the CRREL and the YPG datasets, respectively. We see that in the CRREL dataset air temperature, incoming shortwave solar radiation, VWC, air pressure, and relative humidity are the most important features. While for the YPG dataset we find that incoming shortwave solar radiation, reflected shortwave solar radiation, VWC, air temperature, and air pressure are the important features. The visibility of the objects in the different climatic conditions are impacted by similar environmental conditions, however, in the desert environment the reflected shortwave solar radiation proves to be more important than relative humidity. Table 1 and Table 2 also show the top three most important environmental features for the different object types in the CRREL and YPG datasets, respectively. Aggregating the most salient environmental features for both datasets, we find that the air temperature, incoming shortwave solar radiation, reflected shortwave solar radiation, VWC, relative humidity, and air pressure impact the prediction of the contrast metric the most across both datasets. In the next section, we will leverage the environmental features identified in the sensitivity study to better filter false positives from true positives.

**Table 1: For the CRREL dataset, the test errors for mapping environmental variables to the PCA contrast metric and the top three most important features for each target.**

CRREL	Testing Errors (RMSE)	1st Important Feature	2 <sup>nd</sup> Important Feature	3 <sup>rd</sup> Important Feature
SP	0.39	Air Temperature	Incoming Shortwave Solar Radiation	VWC
SM	0.50	Incoming Shortwave Solar Radiation	Air Temperature	VWC
DP	0.60	Relative Humidity	VWC	Air Pressure
DM	0.48	VWC	Incoming Shortwave solar radiation	Air Temperature

Table 2: For the YPG dataset, the test errors for mapping environmental variables to the PCA contrast metric and the top three most important features for each target.

YPG	Testing Errors (RMSE)	1st Important Feature	2 <sup>nd</sup> Important Feature	3 <sup>rd</sup> Important Feature
SP	0.31	VWC	Reflected Shortwave Solar radiation	Air Pressure
SM	0.21	Incoming Shortwave Solar Radiation	Reflected Shortwave Solar Radiation	VWC
DP	0.33	Incoming Shortwave Solar Radiation	Reflected Shortwave Solar Radiation	VWC
DM	0.29	Reflected Shortwave Solar Radiation	Incoming Shortwave Solar Radiation	VWC
CP	0.14	VWC	Incoming Shortwave Solar Radiation	Air Temperature

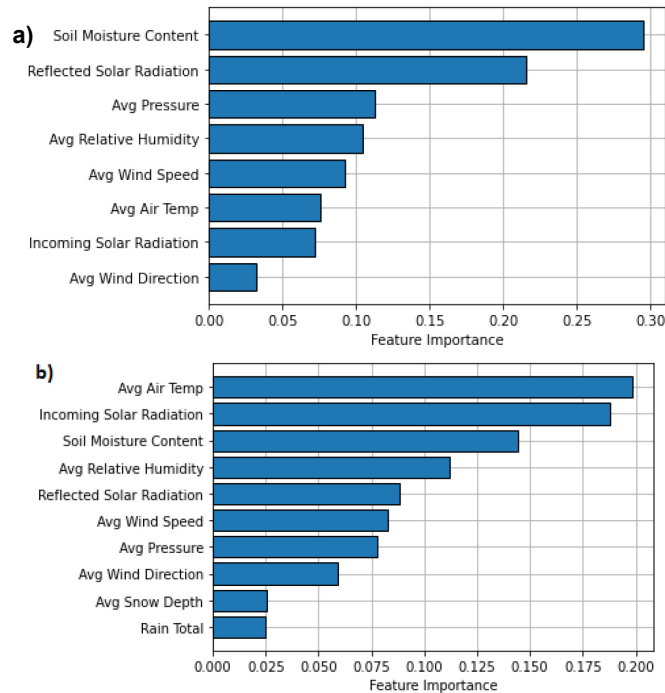


Figure 3: a) The feature importance for the SP object in the YPG dataset, and b) the feature importance for the SP object in the CRREL dataset.

### 3.0 TWO-STEP AUTOMATIC TARGET RECOGNITION ALGORITHM

In this section, we introduce the environmentally informed two-step automatic target recognition (ATR) algorithm developed for the CRREL and YPG datasets. The first step of the ATR is the detection step, which identifies thermal anomalies as a region of interest (ROI) using a computer vision approach. The ROIs produced in the detection step contain both false and true detections, thus, we process the ROIs through the second step to filter out the false detections. The classification step is based on a deep learning approach coupling image chips corresponding to the ROI with the environmental conditions identified in the previous section. In Subsection 3.1, the detection step is described, in Subsection 3.2 the classification step is described, and in Subsection 3.3 we discuss results comparing the two-step ATR to two common deep learning ATR algorithms. For additional details on this approach see [1], [2], [3], [8].

#### 3.1 Detection Step

The ability to detect buried objects using LWIR imagery of the ground surface depends on identifying thermal anomalies at the ground surface induced by the presence of a buried object. This is a difficult task due to the low signal-to-noise ratio often observed in LWIR containing buried objects. To address this challenge, we approach the detection step with a computer vision methodology, which is based on filtering the image to reduce noise while preserving edge-like structures in the image. First, we smooth the image using anisotropic diffusion with a Gaussian diffusion coefficient [18]. Then, we apply the shearlet transform to filter out the high frequency noise in the smoothed image [14]. The shearlet transform can be thought of as an anisotropic version of the wavelet transform which allows for preservation of edge-like structures. Next, we create a so-called edge-image utilizing the Sobel operator on the smoothed and de-noised image. The edge-image is produced by computing the magnitude of the discrete gradient using a 3 by 3 finite difference discrete derivative stencil on each pixel in the image [10]. This highlights the edge structures in the image because the magnitude of the gradient is greater where edges are present. Then, we combine the edge-image produced with the Sobel operator with the smoothed image from the first two steps via the bitwise-and operation [9]. This step essentially creates a new image that is a combination of the smooth and denoised image with the edge image to amplify the edge structures in the image which correspond to thermal anomalies. Lastly, we apply Canny edge detection coupled with the Hough transform to produce ROIs based on the presence of closed-loop edges in the image [4], [13]. Figure 4 shows the flow of the detection step through the filtering, edge enhancement, and edge detection to produce the ROIs. Inherently, the detection of ROIs includes false negatives (missed detections), false positives, and true positives. The false positives can be addressed in the classification step, thus, to reduce the false negatives, we have a low threshold to determine whether a detected area is an ROI. By lowering the threshold, we decrease the false negative rate at the cost of increasing false positives.

#### 3.2 Classification Step

The classification step incorporates the environmental conditions at the time of imaging to improve the accuracy for filtering out the false positives produced in the detection step. For the classification of the ROIs, we incorporate the environmental conditions to better filter false positives from true positives. Figure 5 shows the workflow of the two-step ATR, with the architecture of the classification step highlighted. First, we create a labelled set of ROIs by creating image chips containing just the pixels in the identified bounding box from the detection step. Then, we associate with each ROI the set of six environmental conditions at the time of imaging identified in Section 2. Recall, the environmental conditions consist of the air temperature, air pressure, relative humidity, incoming shortwave solar radiation, reflected shortwave solar radiation, and volumetric water content of the soil. As shown in Figure 5, the ROI image chip is first processed through an image classifier network. We choose to use ResNet34 [12] and the choice of image classifier and the size of ResNet is discussed in [3]. We utilize a version of ResNet that is pre-trained on ImageNet and we remove the classification layer of the network. Next, the vector of environmental conditions is processed through a fully-connected feed-forward neural network, which is called a tabular network in Figure 5. After the

environmental information and the ROIs are processed through the tabular network and the image classifier, respectively, the results are concatenated into a vector, which is then processed through another shallow fully-connected feed-forward neural network with a final classification layer. The final output is a classification that utilized both the environmental and ROI inputs to reason whether an ROI corresponds to a true detection. The architecture of the classification network is inspired by the work in [11].


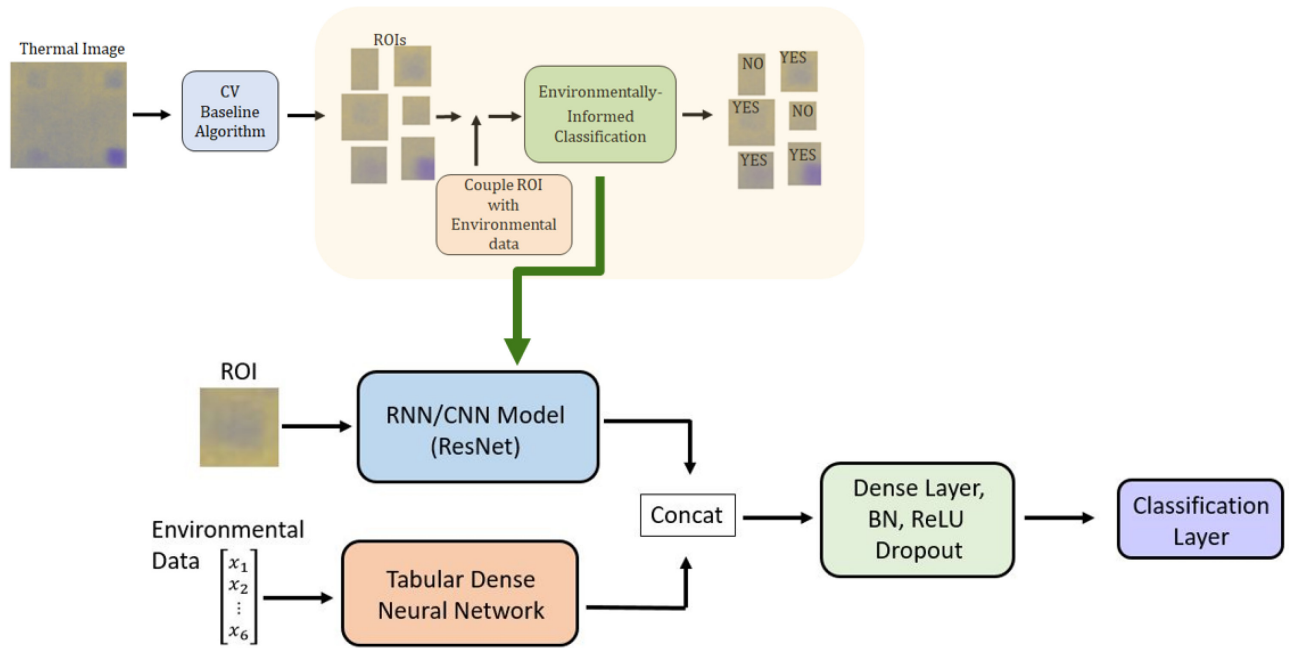


Figure 4: The workflow of the detection step of the ATR algorithm. The detection starts by filtering the raw image using anisotropic diffusion and the shearlet transform, then produces an edge image using the Sobel transform, which is combined with the smoothed image in step c) to create the image shown in step e) which highlights the edges. In step f) canny edge detection is applied to ultimately produce regions of interest in the thermal image. The contour plots in a), b), and c) indicate the reduction of noise edge-structures as the image is smoothed and de-noised.





**Figure 5: The workflow of the two-step ATR with the architecture of the environmentally informed classification step highlighted in more detail.**

For assessing the architecture of the network, we only consider the performance of the two-step ATR on the CRREL dataset. We also perform an ablation analysis where we consider the performance of the classifier algorithm utilizing both the ROIs and environmental data, only the ROIs, and only the environmental variables. To perform this ablation analysis, when assess the classification utilizing only the ROIs, we apply transfer learning on the pre-trained ResNet34 including the original classification layer, and when we consider the environmental input only, we process the environmental input through the tabular network appended with a classification layer. Figure 6 shows the receiver operator characteristic (ROC) curves of the ablation analysis, and we find that the performance on the CRREL dataset is the best when we combine both the environmental and ROI image chips.

Lastly, we assess the performance of the two-step ATR introduced here on the YPG dataset. The first assessment is to directly performance inference on the YPG dataset utilizing the network trained on the CRREL dataset. As we can see in Figure 7, the classification algorithm does not perform well on direct inference to the YPG dataset. However, if we take a small subset of the YPG dataset and perform *transfer learning* on the classifier trained for the CRREL dataset to specialize the weights to the YPG dataset, we find that the accuracy and performance of the classification increases substantially, see Figure 7. In the next section, we will compare the two-step ATR to the performance of two other machine learning ATR algorithms.

### 3.3 Comparison to Machine Learning ATR Algorithms

Lastly, we assess the performance of the two-step ATR algorithm discussed in Section 2 and 3 against the faster R-CNN [17] and the DeepTarget [19] algorithms. Faster R-CNN is a one-step ATR algorithm that is based on a region-proposal convolutional neural network architecture [17]. We assess the comparison algorithms using the CRREL dataset and by applying transfer learning to the YPG dataset. A one-step ATR identifies regions of interest and classifies them in one step. Faster R-CNN is designed for a dynamic dataset, and we are applying it to a static dataset since we are imaging the ground surface. Thus, we apply random cropping to create variability in the CRREL and YPG datasets to avoid overtraining. Note that this algorithm

is not well-suited for our application to static datasets, and thus, the performance on CRREL and YPG datasets is sub-optimal. The DeepTarget algorithm is a two-step ATR developed by N. Nasrabadi [19]. The architecture of DeepTarget is based on the VGG image classifier algorithm for both the detection and classification step. For the detection step of DeepTarget, a VGG-like network is trained on image chips, some of which contain an object and some of which contain only background, and the trained detection network is used as a sliding detector. Performing inference using the sliding detector results in a heat map where the value of the pixel is determined by the probability that the pixel corresponds to a target. ROIs are obtained using a thresholding on the heatmap and the ROIs are classified using another VGG-like network. See [19] for complete details on the DeepTarget ATR algorithm. Table 3 shows the accuracy and the area under the ROC curve (AUC) measurements for the custom two-step ATR, DeepTarget, and Faster R-CNN evaluated on the CRREL dataset and using transfer learning on YPG. In both cases, the environmentally informed ATR outperforms the considered algorithms on the CRREL and YPG datasets.




Figure 6: Receiver operator characteristic (ROC) curves for the ablation analysis of the environmentally informed classification step when evaluated on the CRREL dataset.

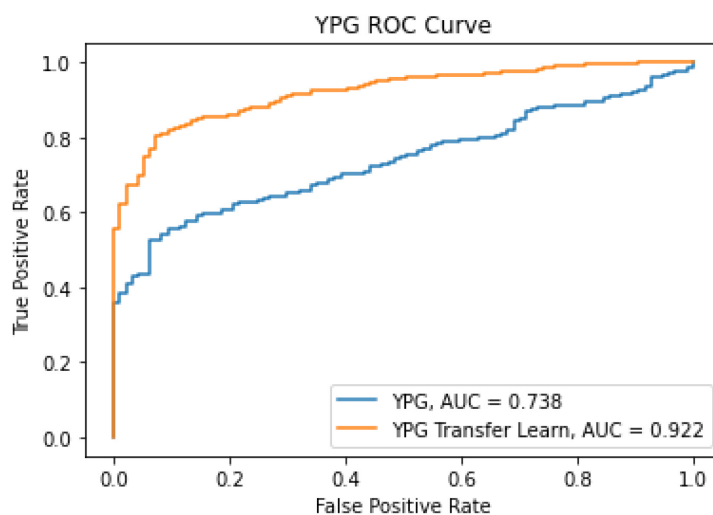


Figure 7: The ROC curves for applying the classification algorithm with and without transfer learning to the YPG dataset.

**Table 3: A comparison of the environmentally informed two-step ATR to DeepTarget and Faster R-CNN on the CRREL and YPG datasets measured with accuracy.**

Algorithm	CRREL Accuracy	YPG Accuracy
<b>Environmental two-step ATR</b>	0.854	0.847
<b>DeepTarget [19]</b>	0.781	0.770
<b>Faster R-CNN [17]</b>	0.723	0.716

## 4.0 DISCUSSION AND CONCLUSION

This paper discussed a two-step ATR algorithm that utilizes filtering and edge detection to produce regions of interest in LWIR imagery containing buried objects and a classification step that leverages the environmental conditions for determining if an ROI corresponds to a true detection. We have shown that, for our considered datasets, when environmental conditions are incorporated, we see an increased accuracy in the classification step when compared to classifying from the ROI-image chips alone. Further, we showed that the algorithm can be extended to distinct climatic conditions via transfer learning to specialize the classification network to new environmental conditions. Finally, we compared the performance of the two-step ATR to DeepTarget and Faster R-CNN and showed that the environmentally informed two-step ATR outperforms the considered algorithms on our datasets. Future directions of this work include extending the classification algorithm to non-binary classification to consider different target types and confusers, such as rocks or disturbed earth without a buried object. We also are considering the utilization of synthetic imagery generated from a physics-based model in the training dataset. First, it would allow us to diversify the training dataset, possibly eliminating the need for transfer learning to other climatic conditions. Second, it could provide away to fine-tune the network to a new location with differing climatic conditions to accurately detect and classify objects obtained with real imagery.

## 5.0 REFERENCES

- [1] Bragdon, S.P., Truong, V.H., Clausen, J.L. (2022). “*Environmental Effects on Contrast Metrics Using Thermal IR Sensors*,” Proceedings of 2022 Virtual Parallel Conference of the Military Sensing Symposia (MSS) Specialty Committees on Battlefield Survivability and Discrimination, Materials, Detectors, and Passive Sensors, (28 February 2022).
- [2] Bragdon, S.P., Truong, V.H., and Clausen, J.L. (2022). “Environmentally Informed Buried Object Recognition,” *Proceedings SPIE 12116: Chemical, Biological, Radiological, Nuclear, and Explosives (CBRNE) Sensing XXIII*, 1211610 (30 May 2022).
- [3] Bragdon, S.P., Truong, V.H., and Clausen, J.L. (2023). “An environmentally aware two-step automatic target recognition algorithm for buried objects,” *Proceedings of 2023 Parallel Conference of the Military Sensing Symposia (MSS) Specialty Committees on Battlefield Survivability and Discrimination*, (23 February 2023).
- [4] Canny, J. (1986). “A Computational Approach to Edge Detection,” *IEEE Transactions on Patterns and Machine Intelligence*, vol. PAMI-8, no. 6, pp. 679-698, doi:10.1109/TPAMI.1986.4767851.
- [5] Clausen, J. L., Dorvee, J. R., Wagner, A., Frankenstein, S., Moriss, B. F., Claffey, K. J., . . . Jones, B. J. (2020). *Spatial and Temporal Variance in the Thermal Response of Buried Objects*. ERDC/CRREL TR 20-10. Hanover, NH: U.S. Army Corps of Engineers, Engineer Research and Development Center, Cold Regions Research and Engineering Laboratory.

- [6] Clausen, J. L., Frankenstein, S., Dorvee, J. R., Workman, A., Moriss, B. F., Claffey, K. J., . . . Jones, B. J. (2021). *Spatial and Temporal Variance of Soil and Meteorological Properties Affecting Sensor Performance - Phase II*. ERDC/CRREL TR 21-10. Hanover, NH: U.S. Army Corps of Engineers, Engineer Research and Development Center, Cold Regions Research and Engineering Laboratory. doi:10.21079/11681/41780.
- [7] Clausen, J.L., Felt, C., Musty, M., Truong, V., Frankenstein, S., Wagner, A., Affleck, R., Peckham, S., and Williams, C. (2022a). *Modernizing Environmental Signature Physics for Target Detection – Phase III*. ERDC/CRREL TR 22-4. Hanover, NH: U.S. Army Corps of Engineers, Engineer Research and Development Center, Cold Regions Research and Engineering Laboratory. doi: 10.21079/11681/43442.
- [8] Clausen, J., Truong, V., Bragdon, S.P., Frankenstein, S., Wagner, A.M., Affleck, R.T., and Williams, C.R. (2022b). *Buried-Object-Detection Improvements Incorporating Environmental Phenomenology into Signature Physics*. ERDC/CRREL TR 22-19. Hanover, NH: U.S. Army Corps of Engineers, Engineer Research and Development Center, Cold Regions Research and Engineering Laboratory.
- [9] Dmitrovic, S. (2021). Bitwise Operators. In: *Modern C for Absolute Beginners*. Apress, Berkley, CA.
- [10] Duda, R. and Hart, P. (1973). *Pattern Classification and Scene Analysis*. Wiley.
- [11] Gessert, N., Sentker, T., Madesta, F., Schmitz, R., Kniep, H., Baltruschat, I., Werner, R., Schlaefer, A. (2019). “Skin Lesion Classification using CNNs with Patch-Based Attention and Diagnosis-Guided Loss Weighting,” *IEEE Transactions on Biomedical Engineering*, vol. 67, no. 2.
- [12] He, K., Zhang, X., Ren, S., and Sun, J. (2016). “Deep Residual Learning for Image Recognition,” *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition 2016*, pp 770-778.
- [13] Illingworth, J. and Kittler, J. (1987). “Adaptive Hough Transform,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, no. 5, pp. 690-698, doi:10.1109/TPAMI.1987.4767964.
- [14] Kutyniok, G. and Labate, D. (eds), (2012). *Shearlets: Multiscale Analysis for Multivariate Data*. Springer, New York.
- [15] Michelson, A. A. (1927). *Studies in Optics*. Chicago, Illinois: The University of Chicago Press.
- [16] Musty, M., Truong, V.H., Clausen, J.L., Frankenstein, S., Williams, C.R., Trautz, A., Howington, S., and Kala, R.V. (2022). *Thermal Infrared Comparison Study of Buried Objects between Humid and Desert Test Beds*. ERDC/CRREL TR 22-10. Hanover, NH: U.S. Army Corps of Engineers, Engineer Research and Development Center, Cold Regions Research and Engineering Laboratory. doi: 10.21079/11681/45064.
- [17] Nasrabadi, N.M. (2019). “DeepTarget: An Automatic Target Recognition Using Deep Convolutional Neural Networks,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 55, no. 6, pp. 2687-2697.
- [18] Perona, P. and Malik, J. (1990). “Scale-Space and Edge Detection Using Anisotropic Diffusion,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 7, pp. 629-639.
- [19] Ren, S., He, K., Girshick, R., and Sun, J. (2015). “Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks,” *Proceedings of NIPS 2015: Advances in Neural Information Processing Systems*.

- [20] Schachter, B. 2013. "Target-Detection Strategies." *Optical Engineering* 54(4):041102.
- [21] Schachter, B. 2020. *Automatic Target Recognition*. (4 ed.). Bellingham, WA: SPIE.
- [22] Arik S. and T. Pfister. 2020. "TabNet: Attentive Interpretable Tabular Learning." Pre-print, arxiv:1908.07442v5.

