A DEEP-LEARNING APPROACH TO UNDERGROUND TARGET DETECTION

Sean P. McKenna
Track 2 Analytics
Melrose, MA, USA
smckenna@track2analytics.com

Lee J. Perren
USACE CRREL
Hanover, NH, USA
lee.j.perren@usace.army.mil

***Abstract*—We investigate the feasibility of a deep-learning approach for subsurface target detection, localization, and characterization (depth estimation). To train and evaluate the three deep neural networks (DNNs) used in this study, we generated a synthetic dataset designed to reflect real-world sensor measurements. Synthetic signals were produced using an analytical model and then corrupted with additive noise of varying structure and amplitude. To produce realistic noise, we adopted a heuristic approach that leverages the scale invariance of many geophysical variables. A long short-term memory (LSTM)-based binary classifier is used to indicate target presence. If a target is detected, a second LSTM-based DNN estimates the along-track target location. A convolutional neural network (CNN)-based model is used to estimate target depth. Overall, this approach achieves accuracy comparable to a model-based correlator used in prior work, while dramatically improving computational speed—by almost three orders of magnitude. These results suggest that AI-based methods are promising and may provide a much-needed boost to the field.**

***Index Terms*—subsurface sensing, target detection, deep learning, synthetic data.**

I. Introduction

Detecting, localizing, and characterizing subsurface targets of interest (e.g., underground infrastructure, mines, cavities, UXO) remains a difficult problem for military, homeland security, and civil works organizations around the world. Applications include underground mine and facility mapping, sinkhole detection, levee or dam assessment, and urban infrastructure mapping. Many different sensing techniques have been applied to this problem over the years, including the use of traditional geophysical methods such as gravity, ground-penetrating radar, electromagnetic, seismic, and synthetic aperture radar methods (e.g., [1–7]). Emerging methods, such as quantum gravimetry [8] and muography [9], have begun to show great promise in their application to subsurface detection and intelligence. Beyond new sensing modalities, an exciting area of development is the use of machine learning and deep learning techniques to improve target detection and characterization [10–16].

In this work, we train a series of three deep neural networks (DNNs) to detect and localize subsurface targets—specifically, long, linear conductors that are typical of the subsurface infrastructure found in many environments. Using an analytical model for an electromagnetic induction (EMI) sensor and a geostatistical approach to produce additive noise, we generated the training data needed by the DNNs. By varying the nature of the background noise and its strength, we were able to produce a range of realistic signals for training and evaluation. The observed impressive performance suggests that the use of AI-based methods is a viable path to pursue and may deliver a timely lift to the field [17].

This material is based upon work supported by the Broad Agency Announcement Program and the Cold Regions Research and Engineering Laboratory (ERDC-CRREL) under Contract No. W913E524C0008.

II. Sensor System

The sensor considered is an electromagnetic gradiometer (EMG) system that relies on the principles of EMI [18]. The idea is to illuminate an electrically conductive target below the surface with an oscillating magnetic field. This time-varying field induces currents in the target(s) that radiate secondary magnetic fields, which can produce a detectable response given a suitable sensor. For the application considered here, we are interested in the response of the EMG system to long, linear underground conductors, such as pipes, conduits, and cabling.

The EMG is a frequency-domain system that consists of an active coil transmitter (usually in a fixed location and oriented as a vertical magnetic dipole) and a mobile sensor, which uses a pair of receivers that produce a gradiometric measurement. Figure 1 shows the man-portable configuration of the EMG sensor. The receivers can be oriented either vertically (as shown) or horizontally in order to measure different components of the electromagnetic field. The sensor is outfitted with a GPS unit that yields centimeter-level position accuracy. The system can also be excited by remote transmitters of opportunity such as nearby AM radio stations. The fundamental measurement made by the system consists of the complex (inphase and quadrature) responses at the two receivers. These are used to compute the magnitude of the gradient, which becomes the measurement we use for detection. Figure 2 shows a sample EMG measurement when a target is present.

III. Training Data Generation

We have previously used traditional methods of target detection with the EMG [5]. Recently, the proliferation of machine learning has motivated us to explore new approaches to detection, in particular deep learning. In order to train

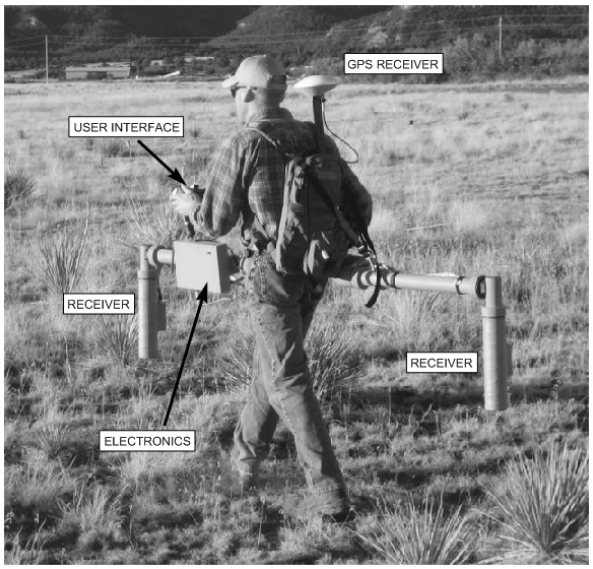

Fig. 1. Photograph of the man-portable electromagnetic gradiometer sensor being used in the field.

the DNNs described in Section IV, a significant amount of training data was needed. Since field data are typically scarce and not well-controlled or lack ground truth, we chose to generate synthetic data. To do this, we relied on an analytical model to produce target signals and then corrupted them with noise of varying structure and level. The model [19] determines the total field as it would be sensed by a triaxial, frequency-domain EMI sensor at an arbitrary location above a homogeneous earth, i.e., the combination of the primary field due to the transmitter, the field due to the presence of the earth, and the field from the subsurface target. In this case, the target is an infinitely long, linear conductor of small diameter. We have implemented the model in MATLAB and use it to simulate the sensing process of the fielded EMG system. Our implementation allows us to vary parameters such as excitation frequency, target depth, and the survey geometry.

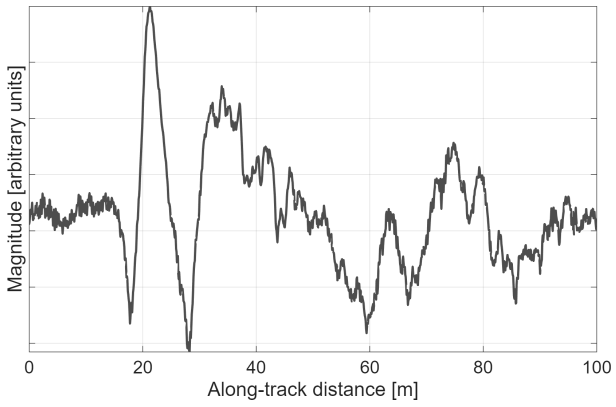

Fig. 2. Sample EMG measurement. Shown is the magnitude of the vertical gradient signal. There is a target approximately 21 m along the survey track.

Running a parametric sweep over the values shown in Table I yielded 336 noise-free target signals, each 200 m in length. The variable x_{TX} reflects the lateral distance between the target and the transmitter, and y_{TX} refers to the distance between the transmitter and the survey line (refer to [18] for these details). The latter is known to the surveyor, the former is not. As such, y_{TX} is used as a feature input into one of the DNNs described in Section IV. A soil conductivity of 50 mS/m was used. For each model run, three frequencies (f_1, f_2, f_3; $f_1 < f_2 < f_3$) were simulated (in the kHz-to-MHz range) to reflect a simultaneous multi-frequency sensor like the EMG.

TABLE I
PARAMETERS OF MODEL SIGNALS GENERATED

Parameter	Values
Depth [m]	3, 5, 7, 9, 12, 15, 20
Crossing angle [deg]	$0, 15, 30, 45$
x_{TX} [m]	$-10, -5, 5, 10$
y_{TX} [m]	50, 75, 100

To generate realistic noisy signals, we followed an approach inspired by the work of Everett and Weiss [20], which takes advantage of the fact that many geophysical variables are universally scale invariant or fractal [21]. This implies that similar features occur over a wide range of scales, and that such variables exhibit a scaling behavior, meaning that their power spectra $P(\nu)$ are proportional to some power of the spatial frequency ν, i.e., $P(\nu) \sim 1/\nu^{\beta}$, where β is the spectral exponent. Results from [20] for one-dimensional EMI sensor profiles support the idea of a power-law dependence for soil conductivity and also show that a typical spatial response profile has stationary increments (the difference between two successive data values). These two properties are typical of fractional Brownian motion, a self-affine fractal signal. For a self-affine data series, the spectral exponent β is bounded: $1 < \beta < 3$. Empirical evidence suggests that the distribution of electrical conductivity in the subsurface follows a lognormal probability density function [22, 23]. Given these observations, we generated 2D random fields of soil conductivity with log-normal statistics and power spectra that satisfy $P(\nu) \sim 1/\nu^{\beta}$ ($1 < \beta < 3$) using the iterative technique outlined by Lewis and Austin [24]. By varying the prescribed spectral scaling exponent, different types of noise signals can be generated, e.g., rapidly varying or slowly varying, both of which we have encountered in the field.

Figure 3 shows a two-dimensional (2D) map of an arbitrary geophysical variable generated using this approach. We have extended this approach to handle a full three-dimensional (3D) domain. While directly incorporating 3D subsurface variability of conductivity and other geophysical parameters would be preferable since it would intrinsically produce noisy measurement signals via simulation, this was not possible since our model uses an analytical integral transform approach. To directly incorporate subsurface variability, a finite-element model or similar would be required. Instead, we found that by sampling from distributions like that shown in Figure 3 and

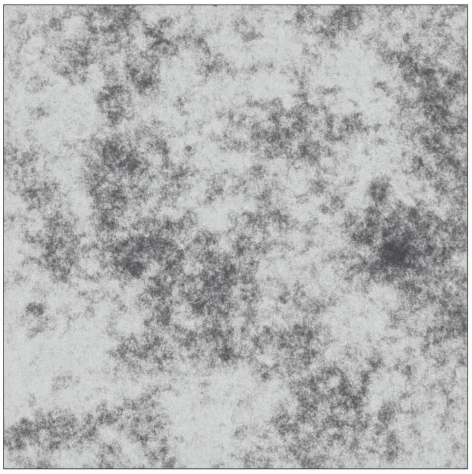

Fig. 3. Synthetically-generated lognormal scale-invariant mapping of an arbitrary geophysical variable.

adding them to the noise-free target signal yielded noisy target signals that resembled those seen in practice.

To each complex noise-free signal we added a noisy real signal and a noisy imaginary signal. Each of the two noisy signals was obtained by sampling (i.e., a horizontal slice) from a 2D array, which we refer to as a "noise map." Figure 3 is an example of such a map. Four noise maps were considered, each having a different spectral exponent ($\beta = 0.5, 1, 1.5, 2$). In all cases, the relationship between the mean of the variable to its standard deviation was 10:1. In addition to varying β, we varied the signal-to-noise ratio (SNR), described below. To produce a sufficiently large and varied dataset, for each spectral exponent and SNR combination, we generated 20 random realizations. To yield these realizations, the noise maps were sampled pseudo-randomly, using the concept of correlation length, where the correlation length represents the spatial extent over which a physical property is correlated. For each realization, a random row was selected, represented by R. We handled the real part of the noise and the imaginary part separately; the approach for both was the same. For each signal, we determined a row to sample by randomly choosing a row that was $\pm$1–5 correlation lengths away from R. This ensured the noise signal added to each measurement was statistically independent yet not unrelated. For the imaginary part, we looked $\pm$5–10 correlation lengths away. Once noise sequences were extracted, they were scaled, demeaned, and interpolated to the number of points in a measurement.

Five different SNRs were considered: 0.3, 1, 3, 10, 30. We defined the SNR as the ratio of the integrated signal energy to the integrated noise energy, where the energy is computed from the squared magnitude of the complex signal. For each SNR, it was necessary to scale the noise sequence such that, when taken in conjunction with the target measurement signal, the prescribed SNR was realized. This was done using a nonlinear optimization routine that yielded a noise scaling factor that would achieve the desired SNR. The final step was to locate the target along each survey. Unlike the model signals where the target is always located at the survey midpoint, our realistic data looks to mimic conditions in the field where, if a target exists, it can be crossed at any point along a survey. We simulated this by randomly shifting signals in space such that the target could be encountered anywhere along the survey track. In total, we produced 134400 target signals. For DNN training and evaluation, we used these target signals and their corresponding noise-only signals, resulting in a dataset of 268800 signals.

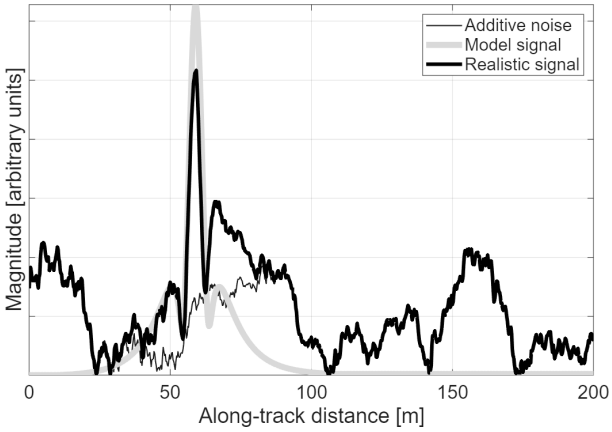

Fig. 4. Example of synthetic data showing the model signal, the additive noise signal, and the resulting realistic signal (c.f., Figure 2). The target is located at 59 m along the collect. Spectral exponent, $\beta = 2$, SNR is 1.0.

IV. Deep Learning Approach

The nature of our detection problem lends itself to two archetypal DNN architectures: the Long Short-Term Memory (LSTM) network [25] and the convolutional neural network (CNN) [26]. (Other architectures are certainly possible, but these two are widely used in timeseries problems.) The LSTM network is a type of recurrent neural network that is a common choice for timeseries data since it is designed to learn and remember patterns in sequential data. A CNN, specifically, a 1D CNN, is a type of neural network that analyzes sequential data (like text, audio, or sensor measurement) by sliding a small filter along the input to detect, and learn, local patterns. For this work, we explored three DNNs: one for detection, one for localization, and one for depth estimation. Figure 5 shows the basic architecture of the three networks. All three are relatively shallow and use standard DNN layers [27–29]. DNNs 1 and 2 have approximately 67k parameters and DNN 3 has approximately 100k. We used a 70/15/15 training/validation/test split. Models were trained for each of the three frequencies considered; however, we found detection and localization performance were largely insensitive to frequency, so we chose to use only one of the frequency models throughout the analysis. Using a single model is also the most pragmatic approach since, in practice, the frequencies collected can vary and may not be known a priori.

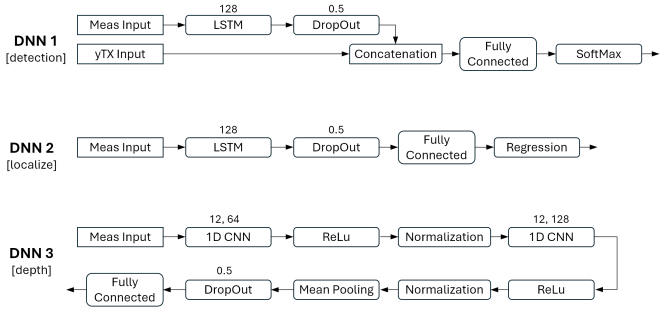

Fig. 5. Overview of the architectures of the three DNNs used in this study. Values above components are as follows. LSTM: number of hidden units; DropOut: dropout probability; 1D CNN: filter size, number of filters.

V. Results and Discussion

We first evaluated the ability of DNN 1 to accurately declare the presence or absence of a target, i.e., detection. Detection performance, measured by accuracy and false-alarm rate, was excellent: accuracies were consistent across SNRs and frequencies, exceeding 92% in all cases (Figure 6). Performance was strong across all SNRs, which was surprising. Unexpectedly, we observed the best performance at the lowest and highest SNRs. This behavior may be related to the way the network learns features; in particular, signals with richer spectral/temporal content due to high noise levels may provide more learnable structure. While the accuracy of this deep-learning approach is comparable to a model-based correlator used in prior work, it dramatically outperforms the correlator in terms of speed, being almost three orders of magnitude faster (e.g., 1493 inferences per second versus 2.9 inferences per second). It also removes the need for an explicit signal model, making the approach more extensible to other sensor configurations and potentially to other sensing modalities.

Next, we evaluated localization (DNN 2). For a location estimate to be deemed correct, we required it to be within 2 m of the true position (based on the geometry of certain targets of interest). For all SNRs save 30, the accuracy for f_1 was about 70% and for f_2 and f_3, it was 78–82%. For SNR $= 30$, the accuracy for f_1 was 74%, for f_2, 85%, and for f_3, 87%. In all cases, the root mean square (RMS) error was near 0.5 m. Thus, on any given survey where a target is present, there is roughly a 75% chance the target will be localized to within 2 m, the actual localization error being about a half-meter.

Finally, we looked at DNN 3 and depth estimation. In general, estimating depth from measurements at the surface is extremely challenging. Here, we considered all SNRs collectively and examined the relationship between accuracy and target depth. To define accuracy, we considered what typically happens in practice: if a probable target is found, the suspected location is excavated. Since such excavation can be costly, one only excavates as deep as necessary, so ideally, information about the anticipated depth to the target is available. With this in mind, we defined a correct depth prediction as $0 \leq d - D \leq 3$, where d is the prediction and D is the truth (here, depth is taken as positive). Practically, this ensures that one excavates deep enough, but also that one is not discouraged from excavating to a very (incorrectly predicted) target. Figure 7 shows the accuracy results versus depth using measurements at each of the three sensed frequencies. As might be expected, the highest frequency does well at the shallowest depth. At intermediate depths (7–12 m), the lowest frequency shows the best performance. Interestingly, at the greatest depth, the highest frequency does best. This observation should not be overinterpreted, since at such depth, the signal becomes dominated by the transmitter rather than the target.

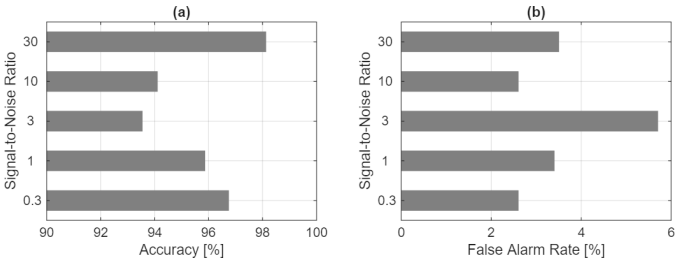

Fig. 6. Detection performance versus signal-to-noise ratio for f_2: (a) accuracy, (b) false alarm rate. Results for f_1 and f_3 are comparable.

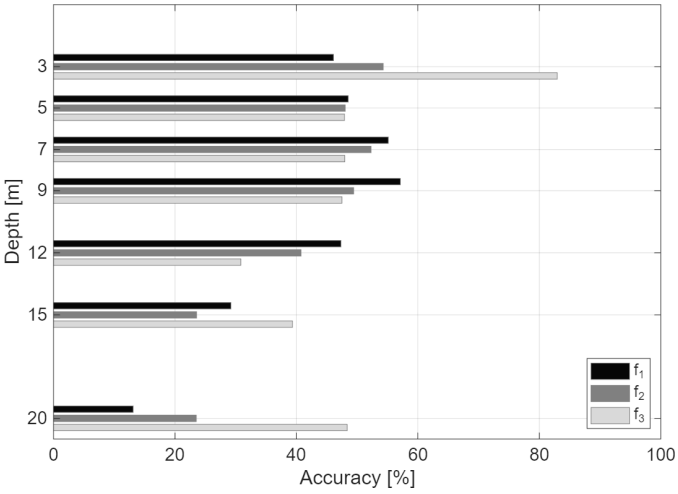

Fig. 7. Depth prediction accuracy for measurements at each of the three sensed frequencies.

VI. Concluding Remarks

We have presented preliminary work that looks at the feasibility of a deep-learning approach to subsurface target detection, localization, and depth estimation. While initial results are quite favorable, there is much more to explore. As with the field of machine learning in general, explainability remains a salient challenge. The DNN architectures we employed were relatively simple and more complex architectures are certainly possible and may improve performance. Lastly, evaluation using field data—currently in progress—is critical to truly assess the practical merits of these techniques.

REFERENCES

[1] D. K. Butler, "Microgravimetric and gravity gradient techniques for detection of subsurface cavities," *Geophysics*, vol. 49, no. 7, pp. 1084–1096, 07 1984.

[2] T. Counts, G. Larson, A. C. Gürbüz, J. H. McClellan, and W. R. S. Jr., "Investigation of the detection of shallow tunnels using electromagnetic and seismic waves," in *Detection and Remediation Technologies for Mines and Minelike Targets XII*, R. S. Harmon, J. T. Broach, and J. H. H. Jr., Eds., vol. 6553, International Society for Optics and Photonics. SPIE, 2007, p. 65531G.

[3] J. A. Martinez-Lorenzo, C. M. Rappaport, and F. Quivira, "Physical limitations on detecting tunnels using underground-focusing spotlight synthetic aperture radar," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 1, pp. 65–70, 2011.

[4] S. D. Sloan, J. J. Nolan, S. W. Broadfoot, J. R. McKenna, and O. M. Metheny, "Using near-surface seismic refraction tomography and multichannel analysis of surface waves to detect shallow tunnels: A feasibility study," *Journal of Applied Geophysics*, vol. 99, pp. 60–65, 2013.

[5] S. P. McKenna, K. B. Parkman, L. J. Perren, and J. R. McKenna, "Automatic detection of a subsurface wire using an electromagnetic gradiometer," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 51, no. 1, pp. 132–139, 2013.

[6] S. D. Sloan, "A history of tunnels and using active seismic methods," *Geophysics*, vol. 86, no. 3, pp. WA49–WA65, 2021.

[7] C. Peng, C. Wang, and Z. Li, "Review of geophysical data acquisition methods for underground feature detection and future trends," *Tunnelling and Underground Space Technology*, vol. 147, p. 105986, 2025.

[8] B. Stray, A. Lamb, A. Kaushik, J. Vovrosh, A. Rodgers, J. Winch, F. Hayati, D. Boddice, A. Stabrawa, A. Niggebaum, M. Langlois, Y.-H. Lien, S. Lellouch, S. Roshanmanesh, K. Ridley, G. Villiers, G. Brown, T. Cross, G. Tuckwell, and M. Holynski, "Quantum sensing for gravity cartography," *Nature*, vol. 602, pp. 590–594, 02 2022.

[9] L. Gebhart and D. Snowden-Ifft, "Muon tomography for detection of dynamic border tunnels," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1053, p. 168383, 2023.

[10] V. Puzyrev, "Deep learning electromagnetic inversion with convolutional neural networks," *Geophysical Journal International*, vol. 218, no. 2, pp. 817–832, May 2019.

[11] M. Huang, J. Ninić, and Q. Zhang, "BIM, machine learning and computer vision techniques in underground construction: Current status and future perspectives," *Tunnelling and Underground Space Technology*, vol. 108, p. 103677, 2021.

[12] Y.-C. Hung, Y.-X. Zhao, and W.-C. Hung, "Development of an underground tunnels detection algorithm for electrical resistivity tomography based on deep learning," *Applied Sciences*, vol. 12, no. 2, 2022.

[13] S. Yu, Y. Shen, and Y. Zhang, "CG-DAE: a noise suppression method for two-dimensional transient electromagnetic data based on deep learning," *Journal of Geophysics and Engineering*, vol. 20, no. 3, pp. 600–609, 05 2023.

[14] C. Liu, Y. Du, G. Yue, Y. Li, D. Wu, and F. Li, "Advances in automatic identification of road subsurface distress using ground penetrating radar: State of the art and future trends," *Automation in Construction*, vol. 158, p. 105185, 2024.

[15] L. Lin, C. Li, H. Wei, Z. Zhong, X. Wang, Q. Li, and A. Gorman, "An improved parametric 3D geologic modeling framework for seismic structure identification using deep learning in complex geologic settings," *Geophysics*, vol. 90, no. 3, pp. IM81–IM102, 04 2025.

[16] D. Tuia, K. Schindler, B. Demir, X. X. Zhu, M. Kochupillai, S. Džeroski, J. N. V. Rijn, H. H. Hoos, F. D. Frate, M. Datcu, V. Markl, B. L. Saux, R. Schneider, and G. Camps-Valls, "Artificial intelligence to advance earth observation," *IEEE Geoscience and Remote Sensing Magazine*, vol. 13, no. 4, pp. 119–141, Dec. 2025.

[17] S. D. Sloan, S. L. Peterie, R. D. Miller, J. Ivanov, and J. R. McKenna, "Detecting clandestine tunnels using near-surface seismic techniques," *Geophysics*, vol. 80, no. 5, pp. EN127–EN135, 2015.

[18] S. P. McKenna, K. B. Parkman, L. J. Perren, and J. R. McKenna, "Response of an electromagnetic gradiometer to a subsurface wire," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 12, pp. 4944–4953, 2011.

[19] S. P. McKenna and J. R. McKenna, "Modeling and analysis of the response of a triaxial, frequency-domain electromagnetic induction sensor to a buried linear conductor," *Geophysics*, vol. 75, no. 1, pp. F1–F14, 2010.

[20] M. E. Everett and C. J. Weiss, "Geological noise in near-surface electromagnetic induction data," *Geophysical Research Letters*, vol. 29, no. 1, pp. 10–1–10–4, 2002.

[21] D. L. Turcotte, *Fractals and Chaos in Geology and Geophysics*, 2nd ed. Cambridge University Press, 1997.

[22] B. Lampe and K. Holliger, "Effects of fractal fluctuations in topographic relief, permittivity and conductivity on ground-penetrating radar antenna radiation," *Geophysics*, vol. 68, no. 6, pp. 1934–1944, 01 2003.

[23] S. M. Lesch and D. L. Corwin, "Using the dual-pathway parallel conductance model to determine how different soil properties influence conductivity survey data," *Agronomy Journal*, vol. 95, no. 2, pp. 365–379, 2003.

[24] G. M. Lewis and P. H. Austin, "An iterative method for generating scaling log-normal simulations," in *11th AMS Conference on Atmospheric Radiation*, 2002, p. JP4.16.

[25] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[26] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.

[27] L. Deng and D. Yu, "Deep learning: Methods and applications," *Foundations and Trends in Signal Processing*, vol. 7, no. 3–4, pp. 197–387, 2013.

[28] J. Schmidhuber, "Deep learning in neural networks: An overview," *Neural Networks*, vol. 61, pp. 85–117, Jan. 2015.

[29] The MathWorks Inc., "Deep Learning Toolbox version: 25.2," Natick, Massachusetts, United States, 2025.