Machine learning methods for the classification of unexploded ordnance from electromagnetic data in marine settings

Lindsey J. Heagy[1], Jorge Lopez-Alvis[1], Douglas W. Oldenburg[1], Stephen Billings[2] and Lin-Ping Song[2]
[1]University of British Columbia, Geophysical Inversion Facility, lheagy@eoas.ubc.ca
[2]Black Tusk Geophysics

SUMMARY

Electromagnetic induction (EMI) methods are commonly used to classify unexploded ordnance (UXO). Modern time-domain systems used for classification are multicomponent and acquire many transmitter-receiver pairs at multiple time-channels. Traditionally, classification is done using a physics-based inversion approach where polarizability curves are estimated from the EMI data. These curves are then compared with those in a library to look for a match based on some misfit. In this work, we develop a convolutional neural network (CNN) that classifies UXO directly from EMI data. Analogous to an image segmentation problem, our CNN outputs a classification map that preserves the spatial dimensions of the input. We train the CNN using synthetic data generated with a dipole model considering relevant UXO and clutter objects. We use a two-step workflow. First, we train a CNN to detect metallic objects in field data. From this, we extract patches of data that contain only background signal and use these to generate a new training data set adding this background noise to our synthetic data. A second CNN is trained with these data to perform the classification. We test our approach using field data acquired with the UltraTEMA-4 system in the Sequim Bay marine test site.

Keywords: Time-domain electromagnetics, Convolutional neural networks, Multicomponent data

INTRODUCTION

Electromagnetic induction (EMI) methods are commonly used to detect and classify UXO. While the use of EMI systems on land is well established it was only recently that systems, such as the UltraTEMA-4, were designed for underwater munitions (Funk et al, 2022). Advanced systems aim to reduce costs related to excavation and interrogation by discriminating UXO from non-UXO objects in a single pass. These systems rely on many sources and receivers and create spatially dense data which have high information content since the targets are illuminated from multiple different angles.

The usual workflow for clearing a site consists of generating a map from EMI data from which anomalies of interest are picked, then classification is done for each of these anomalies and a target list is obtained (Beran et al, 2013). Picking anomalies from a gridded image created from sensor data is done by setting a threshold value of amplitude chosen to maximize the detection of ordnance expected at the site. Once the anomalies are picked, classification is done using a physics-based inversion approach where polarizability curves are estimated from the EMI data (Pasion and Oldenburg, 2001; Bell et al, 2001). These curves are then compared with those in a library to look for a match based on some misfit and a class is assigned (e.g. Pasion et al (2007)).

In Heagy et al (2020) we proposed a convolutional neural network (CNN) to classify UXO from multicomponent EM data and demonstrated it on synthetic examples. The workflow presented in this work extends that previous approach in two main ways: (1) the architecture of the CNNs is changed to include only convolutional layers producing classification maps of higher resolution, and (2) our workflow is validated with field data in a marine setting which required a more thorough design of the training dataset. Our workflow is general enough to be applied to both terrestrial and marine data. Marine data should be pre-processed to remove the EM response of the seawater and sediments (Song and Billings, 2020) and our approach is able to account for the remaining spatially correlated noise.

METHODS

We propose a two-stage workflow: (1) detection and binary classification and (2) multi-class classification. The first stage is mainly used to obtain systematic noise from field data in order to train the multi-

class classifier (Figure 1). However, the output of the first stage, which is a detection map of all metallic objects, may be useful on its own. The output of the second stage is a classification map where the detected objects are classified either as clutter or as a UXO of a certain type. This map may be further processed to obtain a diglist. The main reason for developing this two-step workflow was that the addition of systematic noise to train the multi-class CNN was crucial for obtaining accurate classification results. The CNNs are trained with synthetic data and field data are only used for validation and to approximate the systematic noise. This allows us to have a sufficiently large set of examples to train the CNNs.

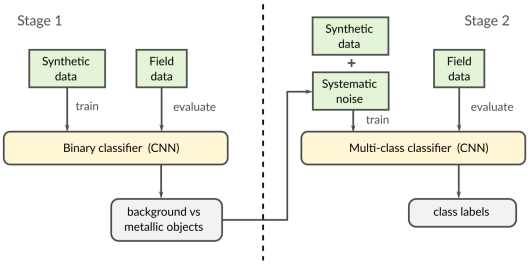

Figure 1: Our proposed two-stage workflow to detect and classify UXOs based on two CNNs.

Architecture of the proposed neural networks

We use a segmentation architecture for our neural networks. CNNs used for segmentation are often fully convolutional which means they only include convolutional layers, and all of the matrices involved are sparse. The proposed architecture is flexible and may be adjusted to any multi-component EM system but it is easier to understand by taking the UltraTEMA-4 system as an example. The UltraTEMA-4 system has 4 transmitters, 12 receivers each with 3 components, and acquires 27 time channels. Figure 2 shows the UltraTEMA-4 system.

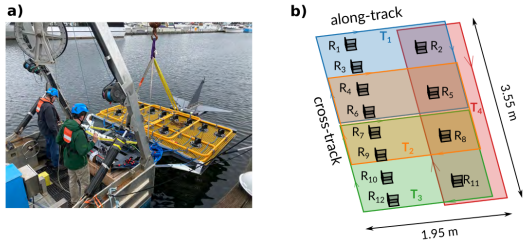

Figure 2: (a) UltraTEMA-4 system, (b) physical dimensions of its layout.

Figure 3 shows our proposed architecture. Defining the input of our CNN requires some choices on how to order the EM data. Following a similar approach to Heagy et al. 2020, we process the EM data using a fixed spatial window in the along-track dimension to create our input. We apply the same normalization and data scaling steps described in Heagy et al., 2020 where we scale the data as a function of time to account for the exponential decay of EM signals, and second, we scale by the maximum amplitude of each sample so that the data all lie within [-1, 1]. Information regarding receiver components and transmitters define the channels of the input. Our input is three-dimensional with multiple channels. However, to create a classification map we are only interested in preserving spatial dimensions, so within the architecture, we eventually collapse the temporal dimension and end up with a 2D output. This output preserves the spatial dimensions of the input, e.g. the number of along-line positions (nx) by the number of receivers (nrx). The number of channels of the output is equal to the number of classes we want to distinguish.

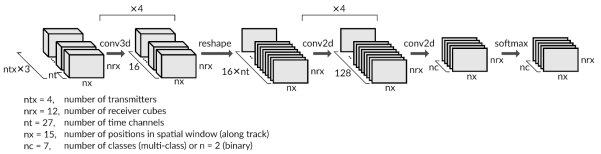

Figure 3: Architecture of CNNs for classification and detection.

The activation functions involved in all layers are rectified linear units (or ReLU) and are applied element-wise and the output of the last convolutional layer is translated to probability values using a softmax function. In this way, the output of the CNN is given in terms of a series of probability maps equal to the number of classes. Each of these maps gives the probability of a certain location to be classified as a particular type of object.

The UltraTEMA-4 system acquires data from each different transmitter sequentially, but for our CNN, we wish to use all of this information together so we use the midpoint of all 4 transmitter positions as the reference point for our classification. A fixed number of transmitter cycles equal to 15 was chosen for our CNN input, which approximately spans 3 m in the along-track direction. The number of classes for the last layer of the detection CNN is two (metallic vs. background) while for the multi-class CNN, the output will depend upon the number of UXO types present. Implementation and training of the proposed CNNs was done using the open-source library Pytorch.

Generating the training data: dipole model

Since UXOs are compact objects, their secondary magnetic fields $\mathbf{B}_s$ can be described as caused by

an induced magnetic dipole moment $\mathbf{m}$. The UXO electromagnetic response is measured in terms of $d\mathbf{B}_s/dt$:

$$\frac{d\mathbf{B}_s}{dt} = \frac{\mu_0}{4\pi}\left[\frac{3\hat{r}\left(\hat{r}\cdot\frac{d\mathbf{m}}{dt}\right) - \frac{d\mathbf{m}}{dt}}{r^3}\right] \quad (1)$$

For an ordnance object, $d\mathbf{m}/dt$ is approximated using a polarizability tensor $\mathbf{L}(t)$. The time-derivative of the induced dipole moment is obtained from:

$$\frac{d\mathbf{m}}{dt} = \mathbf{A}^{\mathbf{T}}\mathbf{L}(t)\mathbf{A}\mathbf{h}_{\mathbf{p}} \quad (2)$$

where $\mathbf{A}$ is a rotation matrix, $\mathbf{h}_{\mathbf{p}}$ is the primary magnetic field generated by the transmitter at the location of the UXO just before the current is shut-off, and $\mathbf{L}(t)$ is the diagonal polarizability tensor with elements L1, L2, L3.

A library of polarizabilities for a large collection of ordnance objects is available from the DoD and was queried to generate the synthetic data used to train our neural networks. To run the forward simulations, we use the simulation code that is a part of BTInvert by Black Tusk.

Parameters for the training set and clutter class

For each training sample, we need to define the type of object (UXO, clutter or no object), its location, and orientation. We also need to specify the background noise. Since the neural networks process the data using a fixed spatial window, the training set needs to include examples that contain: a UXO, a clutter object, and only background signal. The type of UXO is taken from a set of possible ordnance objects that may be defined by means of historic records or any other field-specific prior information.

Defining the clutter class presents a challenge as we do not necessarily know what types of clutter objects we will encounter. We have developed a strategy that uses a physically-based parameterization to define 9 parameters per clutter object. Each of the polarizabilities are modelled by:

$$L(t) = kt^{-\beta}\exp\left(-t/\gamma\right) \quad (3)$$

where k, β and γ control the behavior of the EM decay associated with one of the axes (Pasion and Oldenburg, 2001). We estimate the 9 parameters for each of the UXO objects by performing a non-linear parametric inversion to fit this parameterization to the polarizability curves from the library. Clutter objects are then designed by randomly sampling the remainder of the parameter space and selecting only the samples that are sufficiently from the UXO in this parameter space.

Defining labels for the training dataset

The CNN output has the same spatial dimensions as the input, therefore we need to define labels with these dimensions (nx by nrx) to be used as target labels during training. We choose a single time channel (time channel 6, in our case) and choose a threshold based on the amplitude of the signal across each transmitter and the three components of each receiver, as shown in Figure 4. Locations with amplitudes above the threshold value are labelled as their specific UXO class while remaining pixels are simply labeled as background pixels (blue in Figure 4). We have only considered a single object per fixed spatial window, therefore our current workflow is not yet able to account for multiple objects close to each other.

When performing classification with a trained network, a full acquisition line is processed by sliding a window with a defined step and performing classifications at each step. To generate the classification map we use a voting scheme of all overlapping outputs and assign the label that is voted the most.

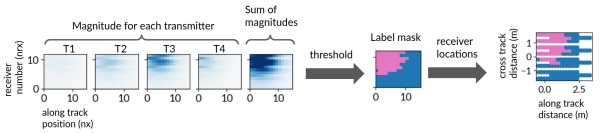

Figure 4: Assigning segmentation-like labels to train our classification CNNs. Points in pink denote signal from a metallic object (UXO or clutter) while points in blue denote background.

RESULTS

We demonstrate our approach using field data acquired with the UltraTEMA-4 system in the Sequim Bay marine test site in 2022. The dataset includes a calibration grid and a blind grid. Here, we focus on the calibration grid, which is composed of 7 acquisition lines and includes 21 seeded objects whose types are known as shown in Figure 5. A preprocessing step was used to remove the EM response of conductive seawater and sediments (Song and Billings, 2020) but some spatially correlated noise still remains in the preprocessed data.

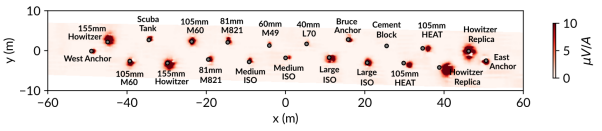

Figure 5: Z-component of pre-processed EM data (time channel 5) collected for calibration grid 2022.

Once the detection CNN is trained it produces a classification map from which we can crop out the

zones surrounding detected objects and take the remaining data as background signal (Figure 6a). Then, we randomly sample pieces of this background signal (Figure 6b) and estimate the systematic noise to be added to the training dataset of the second CNN.

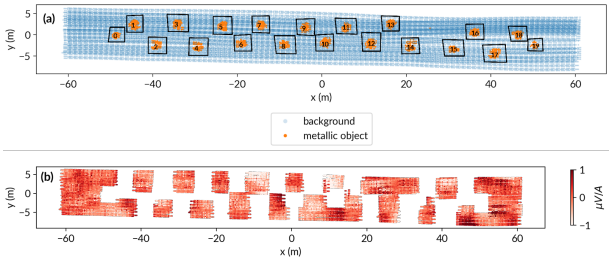

Figure 6: (a) detection map (output of first CNN) and, (b) background signal from which systematic noise is estimated.

Once the multi-class CNN is trained it processes the data per sliding window from which we get a set of probability maps and a single classification map as shown in Figure 7.

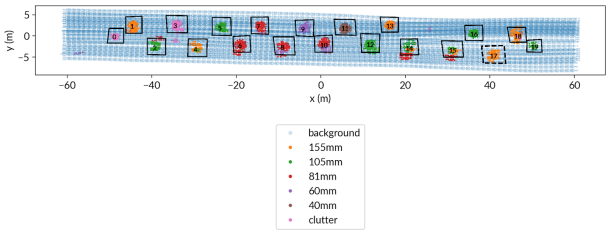

Figure 7: Classification map for Sequim Bay calibration line 2022.

While the classification map is already a result from which we can visually identify the class of object, we can further process it to generate a diglist. We manually select a cell surrounding each of the classified objects. Then, for each cell, we filter out the points classified as background and take the corresponding probability values for each class. All selected probabilities are averaged per class and we choose the class with the highest average probability. The horizontal position is simply estimated by taking an average position of the points within the cell that were not classified as background. In this way, we get a single class per object as shown in Figure 8.

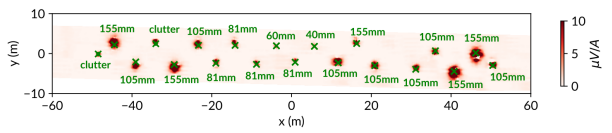

Figure 8: Classes and locations assigned after post-processing; this information together with the average probability values is used to generate a diglist.

Discussion and Conclusions

The workflow we developed is the first that we are aware of that successfully made use of neural networks to classify UXO directly from EMI data. Our workflow was applied to a marine EMI dataset from a test site (Sequim Bay, data acquired in 2021) where it successfully detected all of the UXOs while discriminating ∼70% of the clutter in the blind grid. Naturally, any new method requires development, and one area that requires future work is the need to address multi-target scenarios. However, having an end-to-end workflow will enable us to assess the strengths of a CNN-based approach and identify opportunities where components of our workflow may complement or offer improvements to the standard workflow. Identifying these opportunities will be a focus of future work.

References

Bell T, Barrow B, Miller J, Keiswetter D (2001) Time and Frequency Domain Electromagnetic Induction Signatures of Unexploded Ordnance. Subsurface Sensing Technologies and Applications 2(3):153–175, DOI 10.1023/A:1011978305379, URL https://doi.org/10.1023/A:1011978305379

Beran L, Zelt B, Billings S (2013) Detecting and Classifying UXO. The Journal of ERW and Mine Action 17(1)

Funk R, Gamey TJ, Billings S (2022) ULTRATEMA-4 marine dynamic classification system field trials. In: SAGEEP 2022

Heagy LJ, Oldenburg DW, Perez F, Beran L (2020) Machine learning for the classification of unexploded ordnance (UXO) from electromagnetic data. In: SEG Technical Program Expanded Abstracts 2020, pp 3482–3486, DOI 10.1190/segam2020-3428369.1

Pasion LR, Oldenburg DW (2001) A Discrimination Algorithm for UXO Using Time Domain Electromagnetics. Journal of Environmental and Engineering Geophysics 6(2):91–102, DOI 10.4133/JEEG6.2.91, URL https://doi.org/10.4133/JEEG6.2.91, publisher: Environmental & Engineering Geophysical Society

Pasion LR, Billings SD, Oldenburg DW, Walker SE (2007) Application of a library based method to time domain electromagnetic data for the identification of unexploded ordnance. Journal of Applied Geophysics 61:279–291

Song LP, Billings S (2020) Advanced Marine EMI Processing Techniques for Munitions Detection and Classification. Tech. Rep. MR19-1261