

Research Report

Artificial intelligence and machine learning for efficient minefield clearance

Bruckbauer, A. and Grout, V.

This report was commissioned by **Counter-IED Report** and is published by **Delta Business Media Limited**.

Copyright of the author(s). Reproduced here with their permission and the permission of the publisher.

Recommended citation:

Bruckbauer, A. and Grout, V. (2020) 'Artificial intelligence and machine learning for efficient minefield clearance', Counter-IED Report, Spring/Summer 2020, pp 41-46, published online 27 May 2020. Available at: <https://issuu.com/deltabusinessmedialimited/docs/counter-ied-report-spring-summer-2020?fr=sOTUyYzQ1MDA3NQ>

ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING FOR EFFICIENT MINEFIELD CLEARANCE

By Alexander Bruckbauer and Vic Grout
Wrexham Glyndŵr University

ABSTRACT

Landmines, particularly anti-personnel mines, are dreadful instruments of war. Mines can remain in the ground for decades and injure or kill long after the original conflict. Clearing mines is a dangerous, time consuming and expensive task. Fortunately, mine clearing already has well established and documented processes. To further support these efforts a new research project has started at Wrexham Glyndŵr University to explore the use of machine learning to create a prediction model able to better suggest the positions of hidden landmines based on locations of those already found. Research in psychology and computer science demonstrates the difficulty for humans and machines to create true randomness in their actions. The project will investigate whether it is possible to discover hidden patterns or sequences in mine deployment that could give hints where to look for more. The advantage of the envisioned technique is a lightweight data set only comprising numerical values and their simple acquisition in the field. The proposed system will support – not replace – conventional technology. Although machine learning and A.I. can discover structures, patterns and sequences in a huge data set, that humans cannot, it remains a form of prediction. The aim is therefore not to declare the ground safe ('cleared') but to give suggestions where additional explosives are likely to be found and thus, it is proposed, help to direct mine clearing resources better.

UNCOVERING CONCEALED PATTERNS

Randomness seems trivial but it is hard for humans to produce. Even if we deliberately try to produce it and really think our output is random, it probably isn't. In a study conducted in 2012, subjects were asked to create several random number sequences. Over the course of the project, an algorithm was developed, that was able, when provided with two sequences of one subject and a third from another, to identify the foreign sequence with 88% accuracy⁽¹⁾. Moreover, even machines struggle with true randomness, which is one of the reasons good encryption is hard to implement. So often there is a pattern to things we do, or a machine does, even though it wasn't intended.

One of the many ugly features of an anti-personnel mine field is that the placement of single mines appears to be random, to make it harder for the enemy to overcome or clear it. However, while the location of any single mine might appear random, it clearly didn't just materialize there. There was a strategy to use this weapon system in the first place, plans for where to use it, the tactical deployment, a choice of deployment means, service regulations on how to do that and, if they were laid by hand, a conscious decision to place it at that point. All that considerably limits the space for randomness. Even if the mines were laid without greater tactical considerations, true

randomness is hard to achieve. With this in mind there is good reason to consider that there may be a structure or pattern in the mine locations. This pattern is probably not obvious and very different depending on how the mines were deployed; manually in teams or with an immobile or mobile deployment system. But they might be recognizable using machine learning.

WHAT IS MACHINE LEARNING?

In simple terms, machine learning is a technique to feed data to an algorithm that tries to uncover structures or patterns within the data and so to form a model. After the training, depending on the kind of machine learning and the kind of data used, the algorithm can make predictions based on that model. These predictions could be adding a label to a new unknown data point, a numeric value or the next value in a sequence. A common application for

classification is image recognition, where the machine labels all the things it can recognize in an image. An example for a numeric value could be predicting the price of a house, based on its features. The most commonly known is probably sequence prediction, used in smart phones to predict the next word that will be typed.

Classification with machine learning is not entirely new in detecting explosives. There have been successful projects in the past where machines were trained to identify mines on different kinds of remote sensing data, such as ground penetrating radar or multispectral images.^[2, 3] But can the position of mines in the mine field be found solely by the positions of the mines previously found there? If we consider the accepted strengths of machine learning and the concept that there is no true randomness, then it seems theoretically possible that a specifically

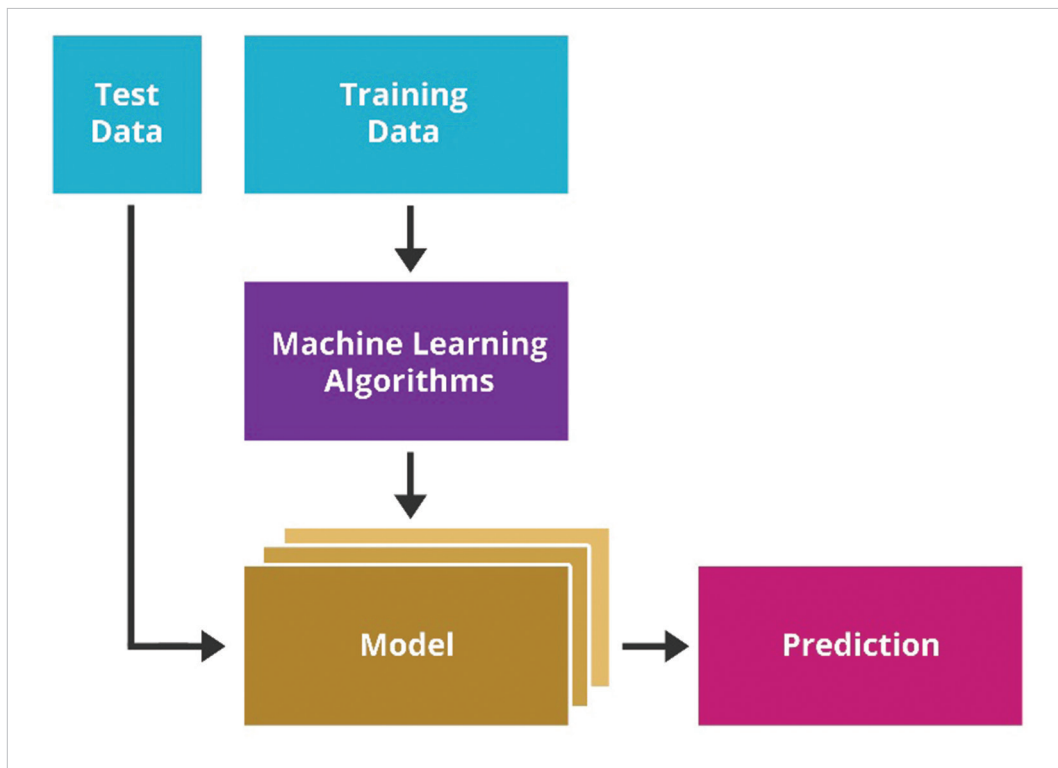


Figure 1: A simplified visualization of a generic machine learning process.

trained model could be able to complete the minefield pattern when enough previous mine positions are known. But how many positions are enough? And can it really work? These are the main questions this new research project at Glyndŵr University is going to investigate. The working hypothesis is that, given enough data, a model can be trained to make helpful predictions about the position of the next mine based on the location of its predecessors.

DATA AND ITS PREPARATION

This might look easy at first glance, feed data into the machine, let it do its magic. But machine learning does not work as it is often portrayed in movies or advertisements. There is no switch that can be flipped and the machine magically begins to learn. Machine learning is a process of multiple, sometimes tedious, steps. Often, and as in this project, the first challenge is the raw data acquisition. As there is no open access to this kind of data, it's not certain that the data exists in a digital format and in the necessary quantities to use it for machine learning. Assuming there is data, the next challenge is preparation. Since we're dealing with location data, these are probably geographic coordinates. This would be a good format because such would be machine readable, but for the purpose of predicting the next mine position based on located mines, it is too specific for a machine to extrapolate general information that it can transfer to a new situation. So the data needs to be generalized for the machine to draw the right conclusions. With this comes the next set of challenges: for example, what will be the units for this generalization? The data needs to be transferred from earth coordinates into a new coordinate system. This could be a Cartesian one. But where to place the origin? The first mine found? The rest of the field could stretch in any direction. How to place the axes? X-axis as west and east and y-axis as north and south, which means geographically oriented? Or maybe the line that connects the first two mines discovered as the x-axis to better reflect the mine field's orientation? The predictions then expected would be regular x and y coordinates. Or the data could be collectively

transformed into a polar coordinate system and would describe mine positions by the length of a vector and an angle, which would also be the values expected from prediction.

SCOPE AND SCALE FOR TRAINING

Scale is also important. The size of a single data sample could be a square kilometre. This probably reveals the larger patterns. Do we deal with mine clusters or lines? Maybe natural obstacles, that were included to direct the enemy towards the mines or to construct the field more efficiently, could disrupt the patterns this project is seeking. Choosing a square of side length ten metres, the single cluster or part of a line pattern may become more obvious but the prediction where the next cluster might start will become harder to make. This decision also depends on the resolution of the raw data. At a one-meter accuracy, a 5 x 5 metre square to analyse a mine pattern probably makes no sense, while a 50 x 50 metre well might. It also ties into the question of how geographically accurate the predictions need to be. An output that tries to be accurate to the millimetre has a high probability of being incorrect while predicting a mine within a 10-meter-wide circle would obviously be useless. Nor is data preparation on its own an easy task. All these questions need to be analysed and investigated carefully, since many of the decisions made at that stage will immediately predefine many parameters of the predictions of the trained model later.

CLASSIFICATION, REGRESSION OR SEQUENCE PREDICTION

Another challenge is the choice of mathematical methods for the machine learning process. The most popular application of machine learning is, as already mentioned, putting labels on things (called 'classification'), then returning a numeric value, which is called 'regression'; or predicting the next element in a sequence, unsurprisingly called 'sequence prediction'. Although it probably is possible to classify our pattern in a useful way, this might not be the best suited technique. For example, the machine learning model, which is the result of the machine learning process, could classify the input data as a specific

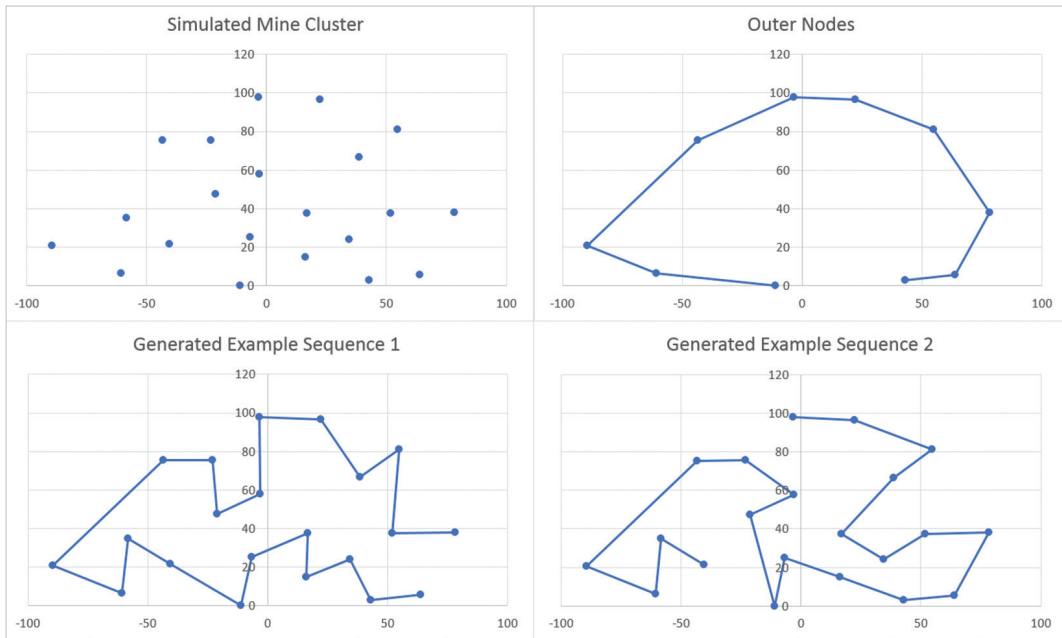


Figure 2: Training data generation: simulating mine distribution (in this case a cluster), determining the outer nodes and generating sequences from the outer nodes through the mine cluster (two shown).

pattern it already knows and the next mine’s position could be extrapolated from there. More fitting for the kind of problem this project hopes to solve, however, seems to be sequence prediction. With this method the model is trained based on sequences.

During data preparation, sequences of positions are generated. These are created from each outer mine position, step by step, following the shortest distance to the next mine, until all positions are added. A machine learning model can be trained using all these prepared sequences. For a prediction, the trained model would then be given the positions from the field, also ordered by the shortest distance between them. The model would try to complete this sequence and offer its prediction for the next position.

HOPES AND EXPECTATIONS

Aside from all the theories, intellectual constructs and capabilities of machine learning and artificial intelligence, what the trained model will actually do is harder to explain and, to an extent, harder still to predict.

Once at this stage, where there is a trained model, its performance needs to be evaluated. In standard machine learning approaches, this is completed against a portion of the original data, withheld from the initial model training. In the ideal case, success would be predictions that were better than mere random guessing. This would be a key milestone indicating that the work could indeed support the mine clearing process.

The Method is not designed to be used in the initial Detection Phase of Mine clearing, rather as a supplementary Support to guide efforts later. Since any guesses (human or AI) are not a safe way to be sure, it is instead envisioned to give suggestions where more mines could be hidden. This could help to direct efforts more efficiently. Specifically, in theory, as more mines are found, the quality of the predictions is expected to increase. As more data is being acquired, further training of the model, based on this new data from these specific locations, allows adaption of the predictions in the operation.

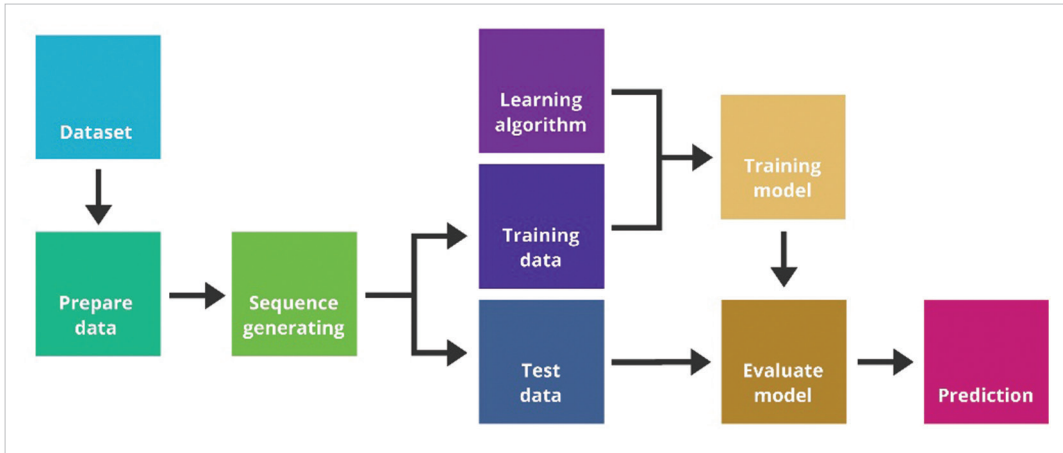


Figure 3: A visualization of the more specific machine learning process in this project.

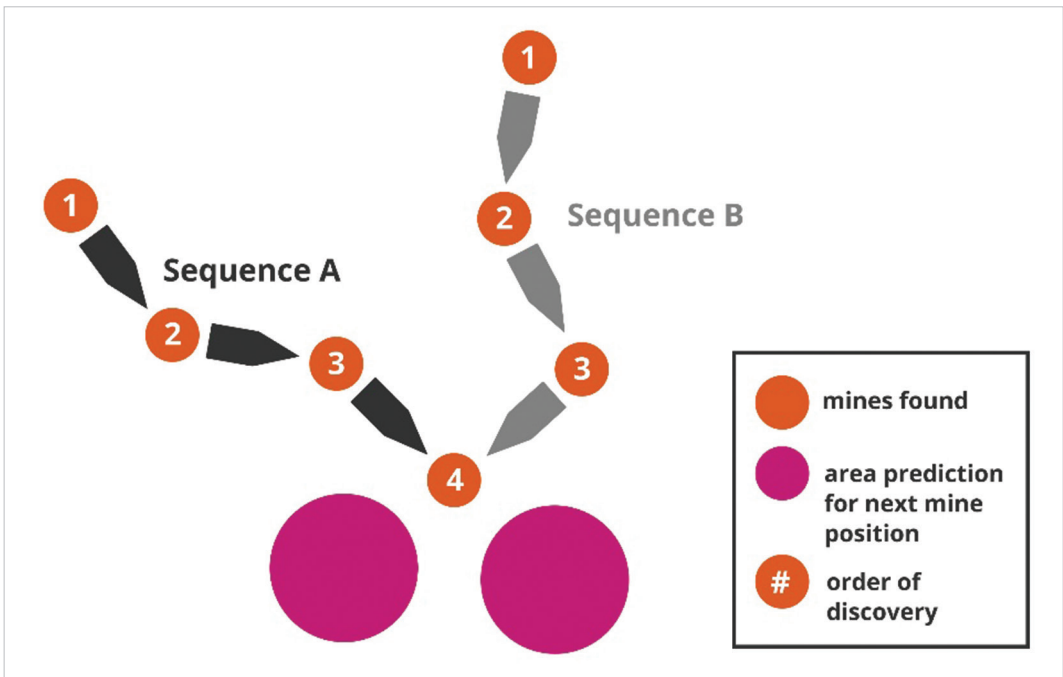


Figure 4: Project goal: predict next mine positions based on the location and order of mines already found.

This is what the project aims to achieve. At the moment it is in the phase of laying the foundations, finding and creating the best data structures,

creating simulation data and determining optimal parameters for data preparation, and starting the machine learning pipeline. Following the established

software development principle of 'separation of concerns', all this can be done based on the best assumptions of how real-world data would look and the use of similarly formed simulated data. At some stage, however, and at the latest when the machine learning pipeline from data preparation over training to evaluation can be shown to work, this project will need actual, real mine location data. At the moment, sources for relevant data are being investigated, but due to the high demand of data in machine learning processes, any additional source would be beneficial for the project and greatly appreciated.

On the way to the project goal of predicting mine positions purely based on the positional data of mines already found there are many challenges. Acquiring data, finding appropriate methods to prepare it for machine learning and good predictions, choosing the right type and algorithm for the task at hand and even parameters that haven't been identified yet. Although we're confident at this point, we can't make assumptions and claim success until we have concrete results. Even if this project does not yield the results we set out to achieve, we are expecting to gain valuable insight into the possibilities and limitations concerning the means at our disposal, creating the opportunity to improve the chances of success in regards to future projects immensely. Either way, this project aims to find out what is possible and we most certainly invite discussion from all quarters. ■

REFERENCES

- [1] <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0041531>
- [2] J. Leiva-Murillo, L. Gómez-Chova and G. Camps-Valls, "Multitask SVM learning for remote sensing data classification", *Image and Signal Processing for Remote Sensing XIV*, vol. 7830, 2010. [Accessed 2 January 2020].
- [3] X. Núñez-Nieto, M. Solla, P. Gómez-Pérez and H. Lorenzo, "GPR Signal Characterization for Automated Landmine and UXO Detection Based on Machine Learning Techniques", *Remote Sensing*, vol. 6, no., 2014. [Accessed 2 January 2020].

ABOUT THE AUTHORS



Alexander Bruckbauer is a PhD student in Computer Science at Wrexham Glyndŵr University and senior software developer at Panasonic Europe. After his military service as Artillery Observer in the German Army he completed several IT apprenticeships in the

network and software development fields. Besides being part of many industry projects, he continued his education, first in Germany, then later graduated from Glyndŵr and was awarded the BCS The Chartered Institute for IT prize. Alexander focuses his research on machine learning and artificial intelligence since his Master's thesis in the field.

Contact: s17001620@mail.glyndwr.ac.uk



Vic Grout is a professional futurist/futurologist at Wrexham Glyndŵr University, investigating emerging and future technologies and their wider social, ethical, legal, political, environmental and economic impact. With degrees in mathematics and computing and

a PhD in engineering, his current role is Professor of Computing Futures, having previously been Professor of network Algorithms. His work ranges from using Internet hardware and software to enable the elderly and disabled to lead longer independent lives, through projects such as effective online age-gating and medical diagnosis, to advising the European Commission on ethical concerns relating to funded research projects. He has authored over 400 research papers and four books on these various topics, edits the 'Turing's Radiator' blog on technological philosophy and appears regularly on radio and TV. His first science-fiction novel, 'Conscious', was published in 2017.