EMPLOYING DEEP Q-NETWORKS FOR ANOMALY DETECTION OF SWARM SATELLITE DATA & BEYOND

Christopher C. O'Neill, Yaxin Bi, Mingjun Huang
School of Computing, Ulster University, Belfast. Email; oneill-c177@ulster.ac.uk

ABSTRACT

This research focuses on a proposed approach using Q-Learning to detect anomalies in space data. The method uses bandpass and notch filters and then employs Q-Learning to search within these bounds for anomalies. In the simplest incarnation, the upper and lower bounds are defined by user input, leading to a program with six parametres. To automate the upper and lower bounds, a second approach employs a series of filters; such as a non-Conformity measure, like kNN or Histogram bins, the allclose function, and the difference blending equation. The average of the differences of the difference blending, as well as the STD, help to obtain the ranges of the high and reduced intensity anomalies. These ranges are passed onto the Q-Learning algorithm as band-pass and notch filters. This method can also be used to automate the input variable of Matrix Profiles. For this reason, the results of the Q-Learning method are compared with Matrix Profile method. The poster also includes a case study using a non-Conformity measure on SWARM satellite data of magnetic anomalies taking place in the lead up to the 2020 Aegean Sea Earthquake. This application employs the Dobrovolsky Radius, as well as the kNN and allclose method outlined in the automated version of the Q-Learning Anomaly Detection Method.

INTRODUCTION

Anomaly detection (or Outlier Detection) has traditionally been solved by mathematical, statistical methods and Machine Learning methods, like kNN, K-means, iForest, and ARIMA. However, the problem of anomaly detection has become increasingly challenging with the drastic shift in the size and data volatility of databases. To address this problem, researchers have turned to Deep Learning methods, such as ANNs, RNNs and LSTMs. However, since these models are all highly domain specific and do not work well with large-scale time series, researchers are now turning to Reinforcement learning and Deep Q-Networks (or Q-Learning). Q-Learning is a branch of Reinforcement Learning, which utilises an agent to take actions (i.e. move position/state) based on its immediate policy to achieve maximum reward.

Previous research in this area (Alavizadeh et al. 2022; Zhang, L. et al. 2022) have used a Deep Q-Learning based approach for cyber security applications. A new time series anomaly detection method, which employs deep reinforcement learning (DRL) and active learning has been developed by Wu, T. & Ortiz, J. (2021) "to efficiently learn and adapt to anomalies in real-world time series data". The proposed Q-Learning method, in this research, is capable of discovering anomalies in time series data, without the need for Active Learning, meaning that it is a marked improvement on previous methods.

OBJECTIVES

The aim of this research is to apply Q-Learning and Deep Q-Networks to detect both high and low-intensity anomalies in time series data. The study area includes detecting anomalies in SWARM satellite data for use in the prediction of earthquakes, but could also concievably be used in other areas.

METHODOLOGY

In the simplest version, a user inputs the range of the notch and band-pass filters, which establishes the extent of the rewards on the Q-Table. When the agent begins to explore the environment, it collects all of the rewards and penalties until it reaches the notch filter, which is the end of that round. At this point, it back propagates using the Bellman Equation and calculates all of the reward values for the QTable. Using the basic Q-Learning algorithm, it was hoped that the Bellman Equation would be capable of detecting the band-pass, as well as the notch filters, creating a kind of planar wave or quadratic curve. However, this was not the case. Instead, the Bellman equation cascades back down the through the reward states and creates an essentially linear gradient, which smooths out the band-pass filter . The Bellman equation is effectively a linear equation, since it contains no exponents.

To eleviate this situation, it is necessary to split the data up into 3 parts (along the y-axis). The first split separates the first gradient (or reward space) from the other two. The Bellman equation can run on this gradient and return the correct values. Similarly the other two gradients can be separated from one another and the same Q-Learning process can run on both to establish their respective gradients. Finally, the three separate tables can be stitched back together to create a kind of planar wave (See Fig 1). The values in the wave can be averaged, so as to become more smooth, or alternatively a single line of data can be processed, which is also much faster. This planar wave is then applied to the data, and this selects for the anomalies. More than this, it offers a kind of fuzzy logic result. Anomalies and normal data are not merely classified as 1s and 0s, respectively, but are given a sliding scale of confidence, which can be leveraged to filter the results still further.

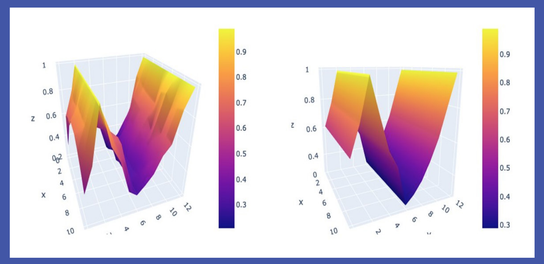

Figure 1: : The Amplitude Index Reward Space that results from the Q-Learning algorithm (left) and the averaged or homogenised version (right). Amplitude is in the y-axis, index is the x-axis, and reward is in the z-axis.

AUTOMATING THE RANGES

Having the user set the ranges for the Q-Learning process is not ideal, because there is nothing to say that this user will even have the appropriate skillset to be able to determine those ranges adequately. Therefore, we can employ a variety of statistical and machine learning methods to carry out this task for us. The first step is to use some form of non-Comformity measure. A Histogram Bin approach is advised for larger data sets (over 10,000 observations), but for smaller datasets we can use kNN. In this example, we apply kNN to our normalised dataset (A) and then find the location of the anomalies using the allclose function. The data set in question is randomly generated consisting of 8 high intensity anomalies and two reduced intensity anomalies.

Next the difference blending equation (abs(1 - (1 - A)) is applied. This allows the kNN algorithm to detect the singularities (beginnings and endings) of the reduced anomalies with a high degree of confidence. This dataset is referred to as k. The next phase is characterised by the discovery of the reduced intensity anomalies, as pretext for the creation of the bounds of the bandpass filter. To begin with, the difference between each adjacent element in k are taken. The mean (v) and standard deviation (z) of this new array p is then calculated. Boolean conditions are established to meet criterion: i - z > z, where i is the value being evaluated and z is the standard deviation of the differences in k, and this creates a list vg. A similar boolean condition criterion: i > (v + z), where i is the value being evaluated, creates vg1. The two lists vg and vg1 are then combined using the AND operator to create vg2, which is then NOTed and used as a mask on list p to create mx1. Finally, the mean value (y) of the masked array is obtained. This value can be used to automatically tune Matrix Profiles for the optimum detection of anomalies and works especially well for ECG data (See Fig 2).

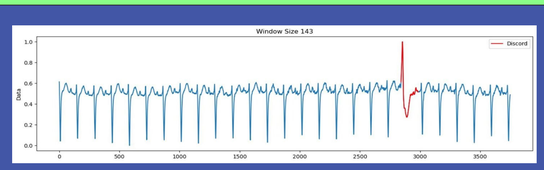

Figure 2: This figure shows the automatic optimisation results of the Matrix Profile algorith using ECG benchmark data.

The next step is the initialisation of a numpy array of zeros: val = 0_n, where n is the length of the random dataset. A double for loop through the elements of the array containing the index values of the anomalies seen in k creates slices equivalent to the magnitude of the mean v, which is the average of the differences in k. These slices or arrays are called mp1 and mp2. It then compares these values using the numpy function np.allclose(), which results in a boolean array named mp3. The boolean values are converted to integer values:

$$mp4_{(i,c)} = \left\{ 1 \; if \; mp3_{(i,c)} = False \; 0 \; otherwise \right\}$$

The integer boolean values are repeated, so that they are the same length

$$mp5_{(i,c,j)} = mp4_{(i,c)}, \forall j \in 0,1,...,N-1$$

where N is the length of the arrays mp1 and mp2 . The various arrays are then cumulatively added in an element-wise fashion:

$$val_{c:(i+int(v)+1)} += mp5_{(i,c)}, \forall c \in index1$$

where v is the average of the differences in k. This adequately detects the reduced intensity anomalies. From here, the range of the reduced intensity anomalies can be discovered and added to our ranges of bandpass/notch filters.

COMPARATIVE STUDY

In this section, we compare the results of our earlier Q-Learning algorithm using user inputs (top) to the Matrix Profile algorithm (bottom), using the NYC Taxi cab dataset as a benchmark.

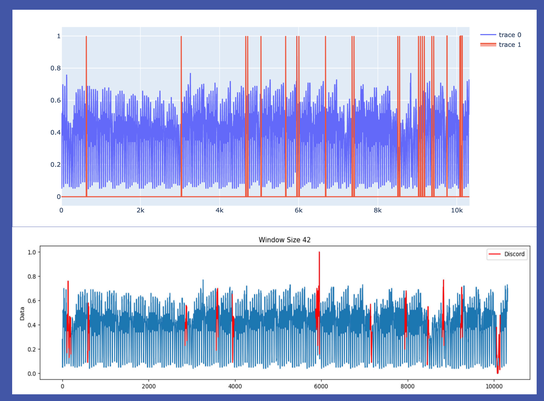

It is quite clear that based on these results there is some conformity, but also some disagreement between the two approaches. The author would suggest using a rolling average on the dataset to improve upon the level of conformity in the results.

2020 AEGEAN SEA CASE STUDY

While reduced intensity geomagnetic anomalies are believed to have some precursive relationship to earthquakes (Li, et al. 2022), for the most part, it is only the high and low-intensity anomalies that are of interest to us. Therefore, for this case study it will be necessary to apply the kNN and allclose function, for anomaly detection and location, respectively. The case study in question focuses on the magnitude 7 Aegean Sea earthquake (also known as the Greek earthquake) , which occurred in 2020. The Dobrovolsky radius (10^{0.43*M}km) was used to set the extent of the area of study (and as a result the amount of data to be considered) and gives us radius of ~1023.2929 km. For the purposes of our study, we will be using VFM product, as it contains the most intact data.

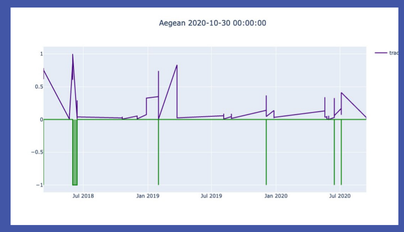

Figure 4: Geomagnetic anomalies from SWARM satellite data collected over a period of 3 years.

In order for progress with geomagnetic anomaly detection for seismic precursors to begin, a number of prerequisites must be established. For instance, tracks must be within ±50° geomagnetic latitude, and the field must be in a state of low activity, and so forth. However, following these guidelines reduces the amount of data significantly, (See Figure 4). To overcome this, the author suggests using a global dataset and/or different parametres, which may not even include geomagnetic data.

Figure 5: The Dobrovolsky radius plot of the Aegean Sea earthquake.

The geomagnetic Dobrovolsky radius plot of the Aegean Sea earthquake reveals a number of anomalies over a 3 year period preceeding the earthquake. These often appear in rapid bursts along satellite tracks. No anomalies appeared directly over the epicentre, but this is likely due to the lack of useable data. Furthermore, it appears that a number of these anomalies occur over populated regions, suggesting a potential anthropogenic source for the electromagnetic signals, in some of those cases.

REFERENCES

Alavizadeh, et al. (2022). Deep Q-Learning Based Reinforcement Learning Approach for Network Intrusion Detection. Computers 11, no. 3: 41. https://doi.org/10.3390/computers11030041

Li, M., et al (2022). Two Large Earthquakes Registered by the CSES Satellite during Its Earthquake Prediction Practice in China. Atmosphere.

Wu, T. & Ortiz, J. (2021). RLAD: Time Series Anomaly Detection through Reinforcement Learning and Active Learning. Available at: https://arxiv.org/abs/2104.00543 [Last accessed: April 12th 2023]

Zhang, L. et al. (2022). A hidden attack sequences detection method based on dynamic reward deep deterministic policy gradient. Secur. Commun. Netw., vol. 2022, pp. 1–13, Jan.