Anomaly Detection with Bayesian Neural Networks

Experts: Alastair Hamilton, Gordon Baggott (Lloyds Banking Group)
Session: Session 1(31 May - 2 July)

Anomaly detection (also outlier detection) is the identification of items, events or observations which do not conform to an expected pattern or other items in a data set [1]. It has a wide number of use cases in industry, including fraud detection [2], malicious content removal [3], etc. Anomaly detection is used in Lloyds for many of these applications, but cyber and other e-crime threat detection [4] is of particular interest to make the bank as secure as possible for our customers and colleagues from criminal threats. We achieve a major element of this security by this using large amounts of internal and external system log data.

A key component of e-crime threat detection is making sure predictions are assigned a confidence to allow threat modellers to adjust thresholds to limit false positive results, reducing the workload of manual investigations. A drawback of using neural networks for predictions is that they do not inherently support confidence intervals. It has been demonstrated that by considering the weights in these models as distributions you can use Bayesian Inference and statistical programming techniques to solve for the weight distributions [5]. By using these confidences we rank order investigations by importance, making sure we investigate the most likely threats first.

Given this you might answer one or more of the following questions:

1. Can we build a Bayesian Neural Network to find anomalies with confidence intervals in tabular data? How does this compare to the results of other anomaly detection algorithms (such as Local Outlier Factor, DBSCAN, etc.)?
2. Can we build Bayesian Convolutional Neural Networks that find anomalies with confidence intervals in spatially dependant data such as images and how does this compare to the results of other anomaly detection algorithms (such as Local Outlier Factor, DBSCAN, etc.)?
3. Can we build Bayesian LSTM networks (long/short-term memory neural networks) that find anomalies with confidence intervals in serialised data such as time-series data? How does this compare to the results of other anomaly detection algorithms (such as Local Outlier Factor, DBSCAN, etc.)?

We recommend using UberAI's Pyro library [6] for the Bayesian Inference as this integrates with Facebook's PyTorch deep learning framework [7] for building the neural nets. The Pyro documentation [9] contains useful guides on how to use the library with a lot of examples. It would also be possible to use R, for example the keras package.

There are a number of useful anomaly detection datasets here (http://odds.cs.stonybrook.edu/). We recommend the **Wine** dataset for question (1), **MNIST** for (2) and the **NYC Taxi** time series from the **NAB** data inventory for (3).

We are hopeful that your research will contribute to the continuous evolution of Lloyds banking Group's security operations programme.

References:

[1] https://cumulocity.com/guides/machine-learning/anomaly-detection/
[2] [2009.07578] Anomaly and Fraud Detection in Credit Card Transactions Using the ARIMA Model (arxiv.org)
[3] https://towardsdatascience.com/fake-news-classification-via-anomaly-detection-765c4c71d539
[4] [2011.02308] Database Intrusion Detection Systems (DIDs): Insider Threat Detection via Behavioural-based Anomaly Detection Systems -- A Brief Survey of Concepts and Approaches (arxiv.org)
[5] [1709.01907] Deep and Confident Prediction for Time Series at Uber (arxiv.org)
[6] http://pyro.ai/
[7] https://pytorch.org/