Anomaly Detection Framework for Cyber-Security Data

Marina Evangelou

Joint work with Niall Adams
Imperial College London

26 September 2017

Anomaly Detection Systems

As the number of cyber attacks, especially zero-day attacks and the emergence of advanced persistent threats (APT) is increasing, new approaches are required to **complement** the existing defence systems, for example signature based methods

Such approaches include anomaly detection systems that seek to detect abnormal deviations from the "normal" behaviour of the network

Anomaly Detection Framework for Individual Devices

The aim of the proposed work is to model "normal" device behaviour and construct an anomaly detection framework based on the behaviour of each individual device

The interest is on individual devices as a commonly observed pattern of a cyber-attack starts with the infection of an individual device

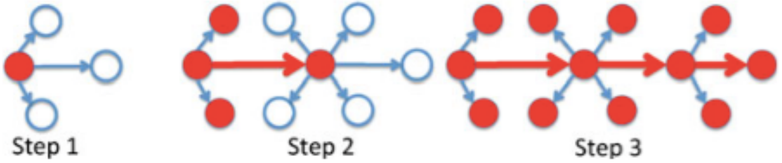

Neil et al. (2013). *Scan Statistics for the Online Detection of Locally Anomalous Subgraphs*.

Technometrics

Device behaviour is defined as the **network traffic** involving the device of interest observed within a pre-specified time period.

Network traffic data are obtained form NetFlow, a protocol operating at the router level which collects flow event logs and is widely used for *auditing* and *monitoring* a network

$$\boxed{\text{time, duration, IP} \rightarrow \text{IP, protocol, ports, packets, bytes}}$$

The data analysed and presented here are part of the anonymised "comprehensive, multi-source cyber-security events" dataset published by Los Alamos National Laboratory in 2015

Device Behaviour

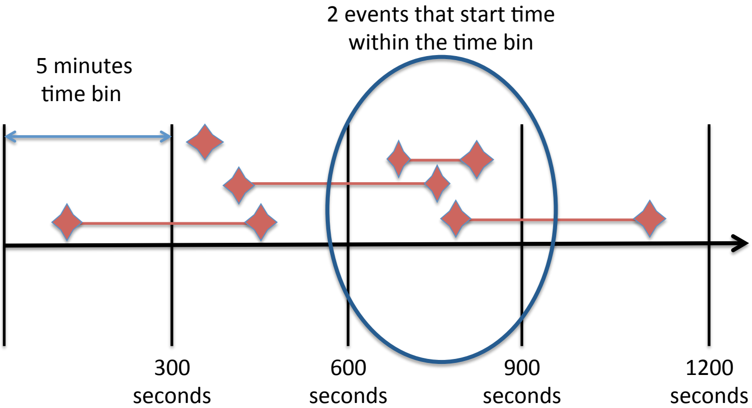

NetFlow event start - end

NetFlow Device Behaviour for 2 devices of the network

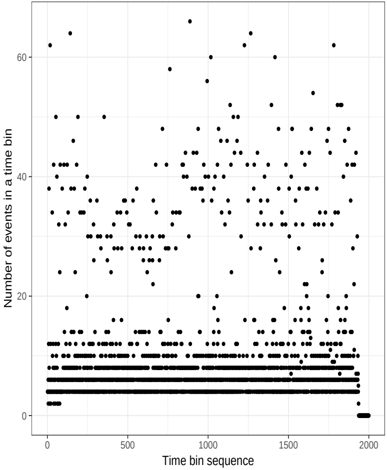

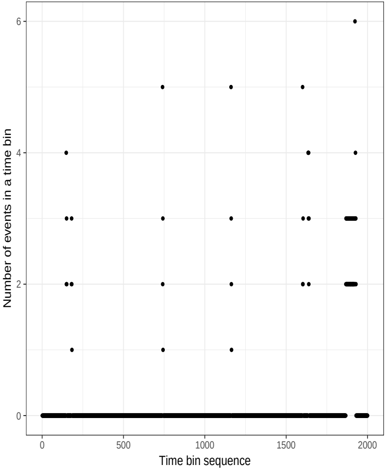

Predicting Device Behaviour

Regression models were built to model the relationship of the response variable Y with a set of constructed features X where:

* Y is the number of events assigned to time bin $t+1$
* X represents the features constructed from the observed data of time bin t

Feature construction

Time	Duration	Source Device	Source Port	Dest. Device	Dest. Port	Protocol	Packets	Bytes
44369	13	C66	N23785	C585	445	6	14	3390
44370	0	C66	N978	C5721	445	6	8	3062

- Event related features:
  Number of events, Number of events with duration more than 5 minutes
- Nature related features:
  Number of events with specific protocol numbers
- Summary statistics of Duration, Bytes, Packets
- Time related features
  Working hours indicator, Working days indicator

- Quantile Regression aims to estimate the conditional quantiles from the data
- The τ^{th} conditional quantile minimizes the expected loss such that:

$$\min_{\beta} \sum_{i} \rho_{\tau}(y_i - X\beta)$$

where:

  * $\rho_{\tau}(\cdot)$ is the quantile regression function

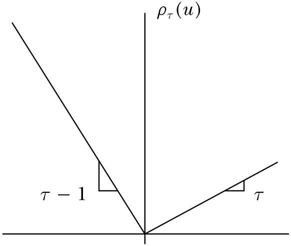

Koenker, R. and Hallock, K.F. (2001). *Quantile Regression*. Journal of Economic Perspectives

Quantile Regression Forests

- Quantile Regression Forests (QRF) proposed by Meinshausen (2006) combine the ideas of Quantile Regression with Random Forests (a collection of regression trees)
- QRF in contrast to Random Forests keep the values of all observations in each node, not just their mean and assess the conditional distribution based on this information

Prediction Intervals

- Let $Q_\alpha(x)$ be the α-quantile such that $Q_\alpha(x) = \inf\{y : F(y|X = x) \geq \alpha\}$ where $F(y) = P(Y \leq y|X = x)$

- A $\eta\%$ prediction interval for the value of Y is given by:

$$I(x) = \{Q_{(1-\eta)/2}(x), Q_{(1+\eta)/2}(x)\}$$

  such that a 95% prediction interval is $\{Q_{0.025}(x), Q_{0.975}(x)\}$

- There is a high probability that a new observation of Y given $X = x$ will lie in the prediction interval

- The width of the prediction interval depends on the observed feature vector

Predictions Intervals: 2.5% and 97.5% Conditional Quantiles

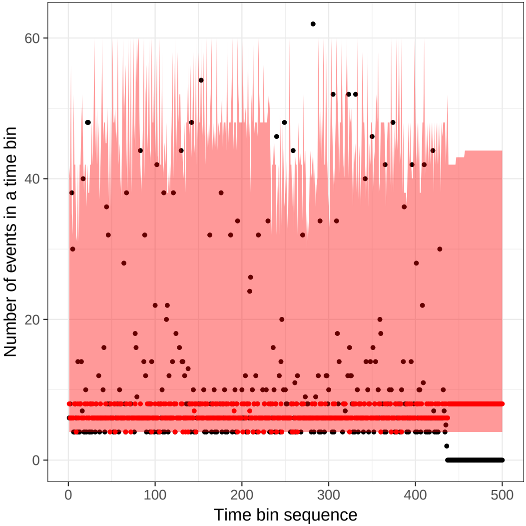

Proposed Anomaly Detector

Observed device behaviour can be characterised as anomalous if it lies outside of the constructed prediction intervals of the QRF models, such that

$$\begin{cases} y_{observed} > Q_{(1+\eta)/2}(x) \\ y_{observed} < Q_{(1-\eta)/2}(x) \end{cases}$$

where $y_{observed}$ is the recorded device behaviour

Validation of the Proposed Anomaly Detection Framework

- The proposed anomaly detection framework was validated through a series of experiments
- The anomaly detector is compared to:
  - Benchmark Anomaly Detector: any observed values outside of the 95% (unconditional) quantiles of device behaviour are classified as abnormal

Validation of the Proposed Anomaly Detection Framework

- The proposed anomaly detection framework was validated through a series of experiments
- The anomaly detector is compared to:
  - Benchmark Anomaly Detector: any observed values outside of the 95% (unconditional) quantiles of device behaviour are classified as abnormal
  - Pruned Exact Linear Time (PELT): Change-point detection approach proposed by Killick *et al.* (2012)

Validation of the Proposed Anomaly Detection Framework

- The proposed anomaly detection framework was validated through a series of experiments
- The anomaly detector is compared to:
  - Benchmark Anomaly Detector: any observed values outside of the 95% (unconditional) quantiles of device behaviour are classified as abnormal
  - Pruned Exact Linear Time (PELT): Change-point detection approach proposed by Killick *et al.* (2012)
- The Benchmark anomaly detector intervals are the same across all time bins
- Both the Benchmark anomaly detector and PELT do not rely on the feature vector

Validation Experiment

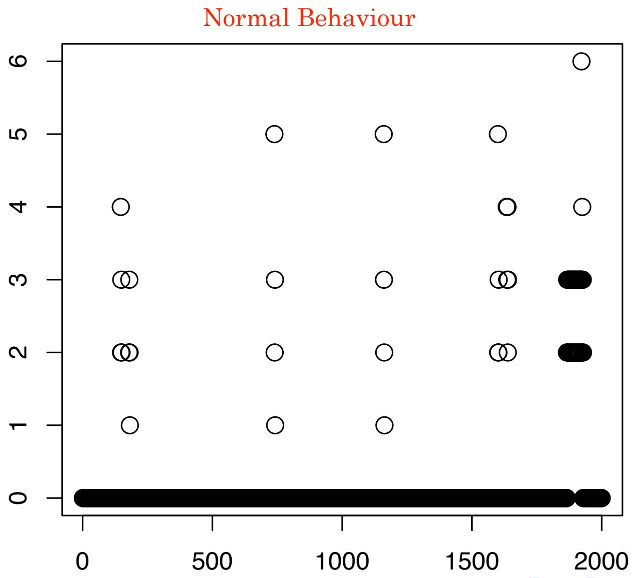

Validation Experiment

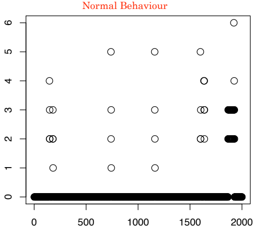

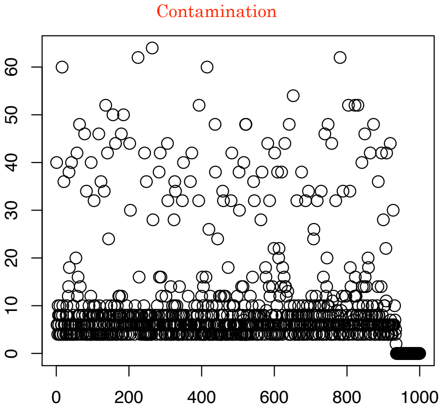

Validation Experiment

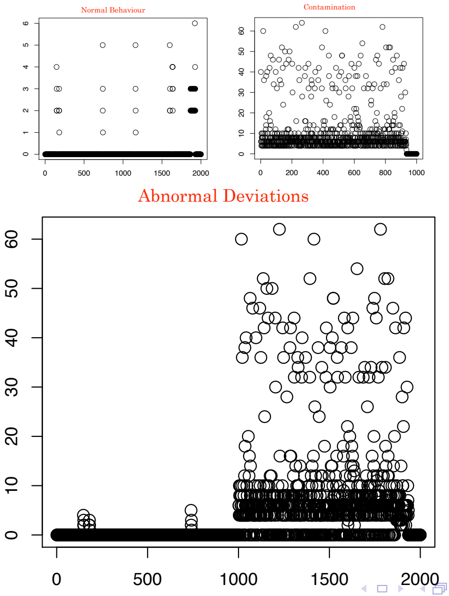

Validation Experiment

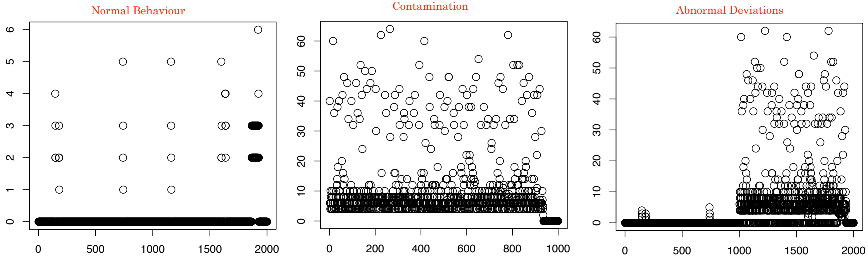

Detector	Accuracy	Sensitivity
QRF	0.959 (0.019)	0.934 (0.012)
PELT	0.544 (0.227)	0.349 (0.062)
Benchmark	0.930 (0.006)	0.935 (0.012)

NetFlow and Process device behaviour

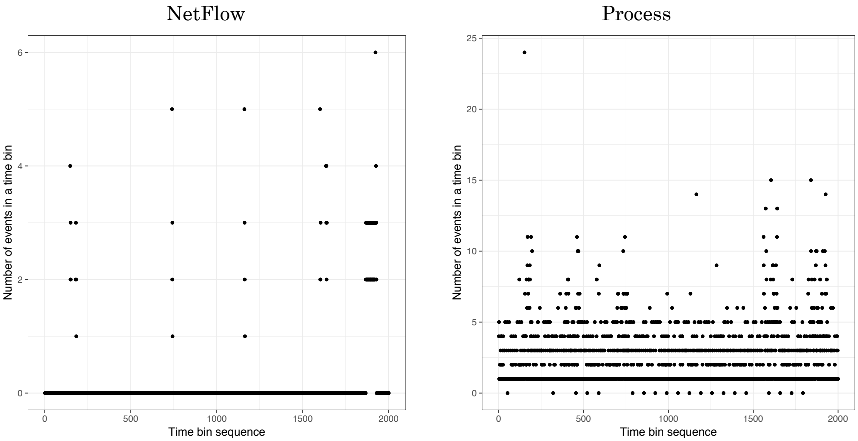

Process data: Feature construction

Time	User	Device	Process	Start/End
2	C66@DOM1	C66	N23785	Start
2	C66@DOM1	C66	N978	Start

- Event related features:
  Number of events, Number of events that only started in the time bin
- Nature related features:
  Entropy of processes
- Time related features
  Working hours indicator, Working days indicator

Process data: Predictions Intervals

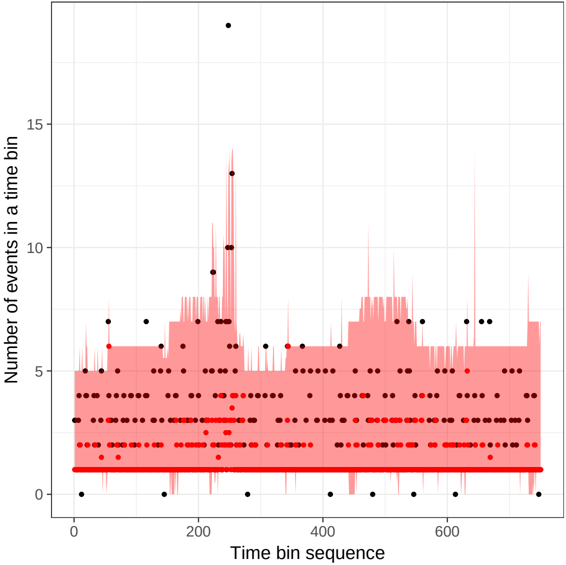

Conclusions

- Diverse device behaviours are observed across the network
- The QRF approach was found to have the best performance across a number of other tested regression models
- A data-driven anomaly detection framework is proposed that is based on prediction intervals of QRF models
- Through a number of validation experiments the proposed framework was found to outperform other detectors
- The anomaly detection framework can be extended for other data sources, *e.g.* process data

Thank you for listening ! Any Questions?

- Data: https://csr.lanl.gov/data/cyber1/
- Adams, N., and Heard, N. (2016). *Dynamic networks and cyber-security*. World Scientific
- Evangelou, M. and Adams, N. (2016). *On the predictability of NetFlow data*. IEEE Information and Security Informatics
- Neil, J. *et al.* (2013). *Scan Statistics for the Online Detection of Locally Anomalous Subgraphs*. Technometrics
- Meinshausen, N. (2006). *Quantile Regression Forests*. Journal of Machine Learning Research
- Koenker, R. and Hallock, K.F. (2001). *Quantile Regression*. Journal of Economic Perspectives
- Killick, R., *et al.* (2012). *Optimal detection of changepoints with a linear computational cost*. arXiv:1101.1438v3