

Anomaly Detection with Extreme Value Theory

A. Siffer, P-A Fouque, A. Termier and C. Largouet

April 26, 2017



CONTENTS

Context

Providing better thresholds

Finding anomalies in streams

Application to intrusion detection

A more general framework

Context

GENERAL MOTIVATIONS

—o Massive usage of the Internet



GENERAL MOTIVATIONS

- o Massive usage of the Internet
 - More and more vulnerabilities



1 Tbps DDoS Attack

Powered By 150,000 Hacked IoT Devices

GENERAL MOTIVATIONS

How the Carbanak cybergang stole \$1bn A targeted attack on a bank

- o Massive usage of the Internet
 - More and more vulnerabilities
 - **More and more threats**



1 Tbps DDoS Attack

Powered By 150,000 Hacked IoT Devices

GENERAL MOTIVATIONS

How the Carbanak cybergang stole \$1bn A targeted attack on a bank

- o Massive usage of the Internet
 - More and more vulnerabilities
 - More and more threats
- o Awareness of the sensitive data and infrastructures



1 Tbps DDoS Attack

Powered By 150,000 Hacked IoT Devices

⇒ Network security :
a major concern

- o IDS (Intrusion Detection System)
 - Monitor traffic
 - Detect attacks

A SOLUTION

- o IDS (Intrusion Detection System)
 - Monitor traffic
 - Detect attacks
- o Current methods : rule-based
 - Work fine on common and well-known attacks
 - Cannot detect new attacks



A SOLUTION

—o IDS (Intrusion Detection System)

- Monitor traffic
- Detect attacks

—o Current methods : rule-based

- Work fine on common and well-known attacks
- Cannot detect new attacks



—o Emerging methods : anomaly-based

- Use the network data to estimate a normal behavior
- Apply algorithms to detect abnormal events (→ attacks)



→ Basic scheme



- Basic scheme




- Many "standard" algorithms have been tested

OVERVIEW

- Basic scheme



- Many "standard" algorithms have been tested
- Complex pipelines are emerging (ensemble/hybrid techniques)



- Algorithms are not magic
 - They give some information about data (scores)

—o Algorithms are not magic

- They give some information about data (scores)
- But the decision often rely on a human choice

```
if score > threshold then trigger alert
```


—o Algorithms are not magic

- They give some information about data (scores)
- But the decision often rely on a human choice

`if score > threshold then trigger alert`

—o The thresholds are often hard-set

- Expertise
- Fine-tuning
- Distribution assumption

INHERENT PROBLEM

—o Algorithms are not magic

- They give some information about data (scores)
- But the decision often rely on a human choice

`if score > threshold then trigger alert`


—o The thresholds are often hard-set

- Expertise
- Fine-tuning
- Distribution assumption


—o **Our idea:** provide dynamic threshold with a probabilistic meaning

Providing better thresholds

MY PROBLEM

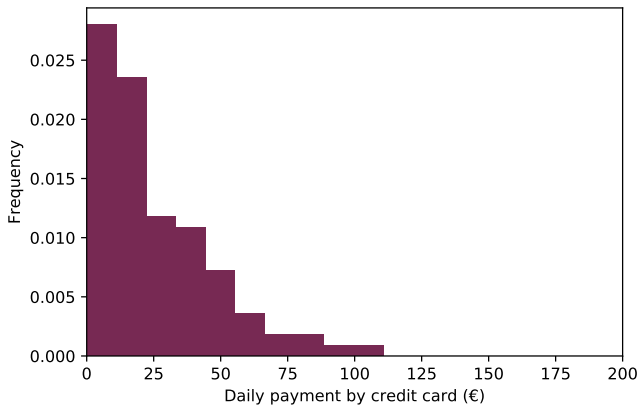


MY PROBLEM

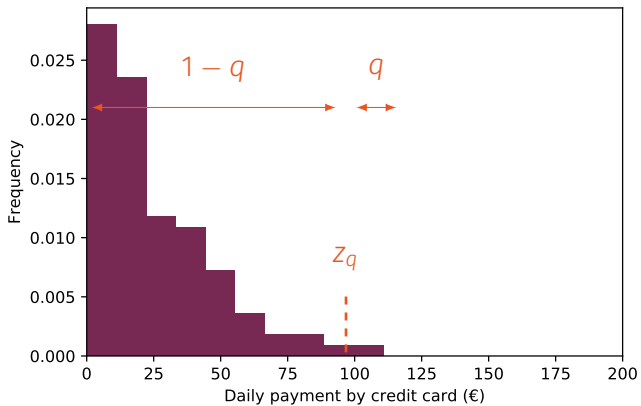


→ How to set z_q such that $\mathbb{P}(X \text{€} > z_q) < q$?


SOLUTION 1: EMPIRICAL APPROACH



SOLUTION 1: EMPIRICAL APPROACH




SOLUTION 1: EMPIRICAL APPROACH




→ Drawbacks: stuck in the interval, poor resolution


SOLUTION 2: STANDARD MODEL




SOLUTION 2: STANDARD MODEL



SOLUTION 2: STANDARD MODEL







SOLUTION 2: STANDARD MODEL



→ Drawbacks: manual step, distribution assumption

REALITIES

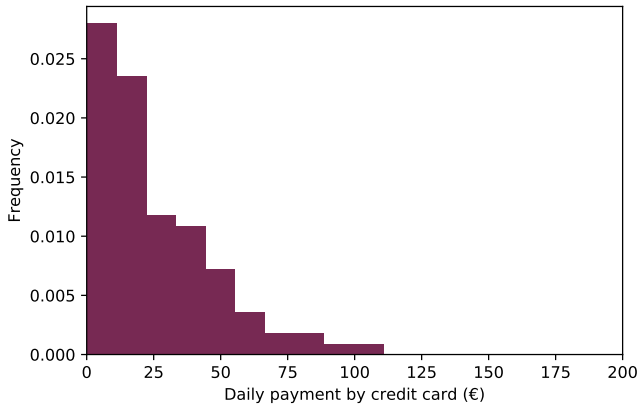


→ Different clients and/or temporal drift


RESULTS

PROPERTIES	Empirical quantile	Standard model
<i>statistical guarantees</i>	Yes	Yes
<i>easy to adapt</i>	Yes	No
<i>high resolution</i>	No	Yes

INSPECTION OF EXTREME EVENTS



INSPECTION OF EXTREME EVENTS



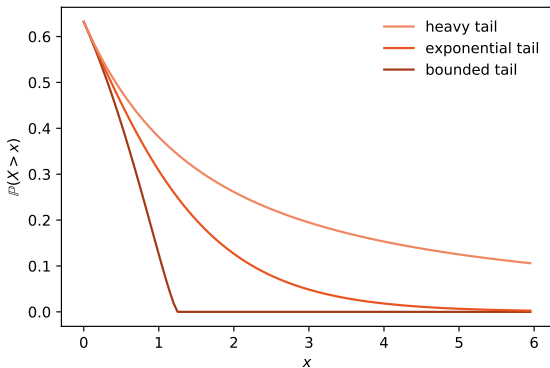
- Main result (Fisher-Tippett-Gnedenko, 1928)

The extreme values of any distribution have nearly the same distribution (called Extreme Value Distribution)

EXTREME VALUE THEORY

- Main result (Fisher-Tippett-Gnedenko, 1928)

The extreme values of any distribution have nearly the same distribution (called Extreme Value Distribution)



→ Let X_1, X_2, \dots, X_n a sequence of i.i.d. random variables with

$$S_n = \sum_{i=1}^n X_i \quad M_n = \max_{1 \leq i \leq n} (X_i)$$

→ Let X_1, X_2, \dots, X_n a sequence of i.i.d. random variables with

$$S_n = \sum_{i=1}^n X_i \quad M_n = \max_{1 \leq i \leq n} (X_i)$$

→ Central Limit Theorem

$$\frac{S_n - n\mu}{\sqrt{n}} \xrightarrow{d} \mathcal{N}(0, \sigma^2)$$

→ Let X_1, X_2, \dots, X_n a sequence of i.i.d. random variables with

$$S_n = \sum_{i=1}^n X_i \quad M_n = \max_{1 \leq i \leq n} (X_i)$$

→ Central Limit Theorem

$$\frac{S_n - n\mu}{\sqrt{n}} \xrightarrow{d} \mathcal{N}(0, \sigma^2)$$

→ FTG Theorem

$$\frac{M_n - a_n}{b_n} \xrightarrow{d} \text{EVD}(\gamma)$$

- Second theorem of EVT (Pickands-Balkema-de Haan, 1974)

The excesses over a high threshold follow a Generalized Pareto Distribution (with parameters γ, σ)

- o Second theorem of EVT (Pickands-Balkema-de Haan, 1974)

The excesses over a high threshold follow a Generalized Pareto Distribution (with parameters γ, σ)

- o What does it imply ?
 - we have a model for extreme events
 - we can compute z_q for q as small as desired

- Get some data $X_1, X_2 \dots X_n$
- Set a high threshold t and retrieve the excesses $Y_j = X_{k_j} - t$ when $X_{k_j} > t$

HOW TO USE EVT


- Get some data $X_1, X_2 \dots X_n$
- Set a high threshold t and retrieve the excesses $Y_j = X_{k_j} - t$ when $X_{k_j} > t$
- Fit a GPD to the Y_j (→ find parameters γ, σ)

HOW TO USE EVT

- Get some data $X_1, X_2 \dots X_n$
- Set a high threshold t and retrieve the excesses $Y_j = X_{k_j} - t$ when $X_{k_j} > t$
- Fit a GPD to the Y_j (→ find parameters γ, σ)
- Compute z_q such as $\mathbb{P}(X > z_q) < q$

HOW TO USE EVT

- Get some data $X_1, X_2 \dots X_n$
- Set a high threshold t and retrieve the excesses $Y_j = X_{k_j} - t$ when $X_{k_j} > t$
- Fit a GPD to the Y_j (→ find parameters γ, σ)
- Compute z_q such as $\mathbb{P}(X > z_q) < q$




Finding anomalies in streams

STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM



STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM

(initial batch)


$X_1, X_2 \dots X_n$




STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM




STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM



STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM



STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM




STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM



(stream)


$X_{i>n}$

STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM

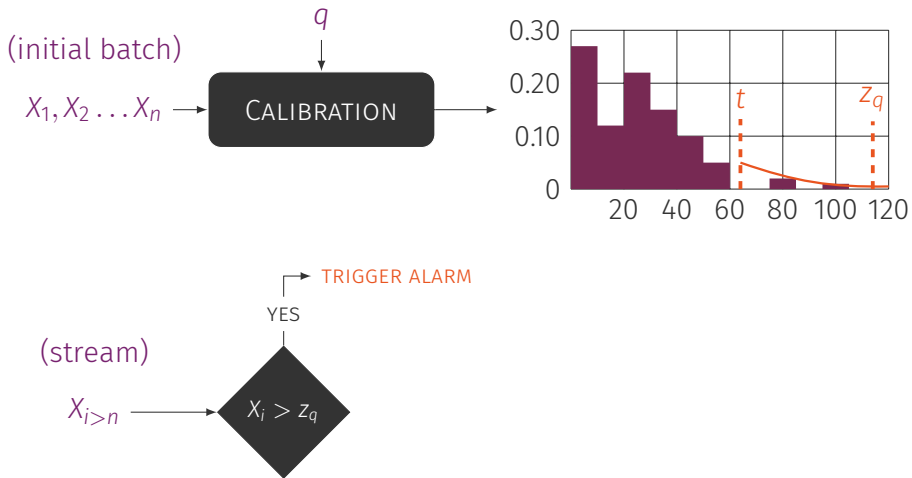


(stream)


$X_{i>n}$



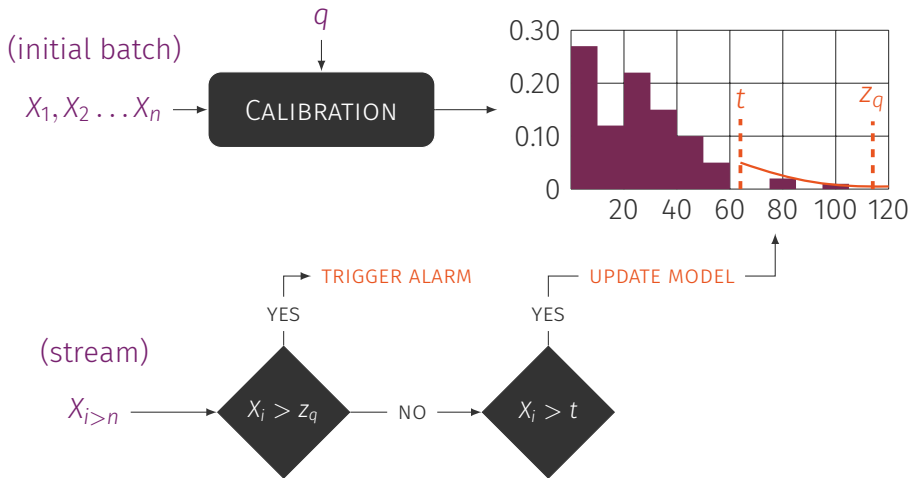
STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM



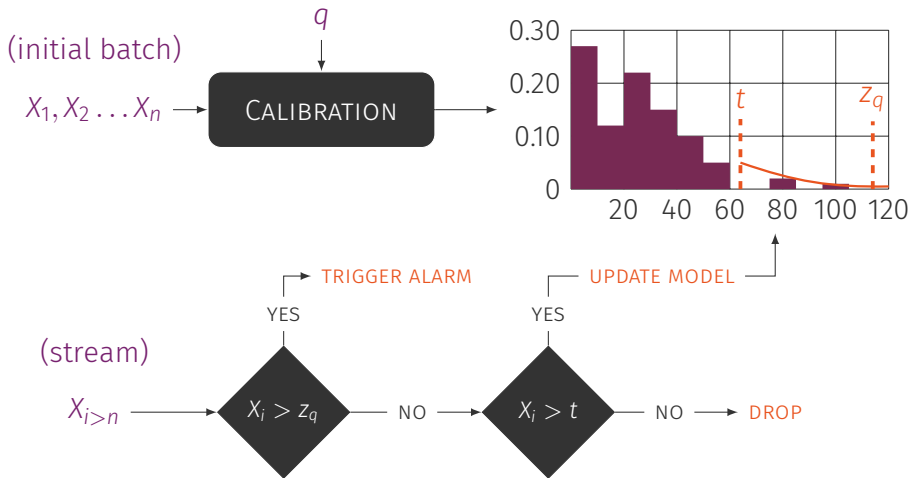
STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM



STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM



STREAMING PEAKS-OVER-THRESHOLD (SPOT) ALGORITHM




CAN WE TRUST THAT THRESHOLD z_q ?

- An example with ground truth : a Gaussian White Noise
 - 40 streams with 200 000 iid variables drawn from $\mathcal{N}(0, 1)$
 - $q = 10^{-3} \Rightarrow$ theoretical threshold $z_{th} \simeq 3.09$

CAN WE TRUST THAT THRESHOLD z_q ?

- o An example with ground truth : a Gaussian White Noise
 - 40 streams with 200 000 iid variables drawn from $\mathcal{N}(0, 1)$
 - $q = 10^{-3} \Rightarrow$ theoretical threshold $z_{th} \simeq 3.09$
- o Averaged relative error



Application to intrusion detection

→ Lack of relevant public datasets to test the algorithms ...

- o Lack of relevant public datasets to test the algorithms ...
- o KDD99 ? See [McHugh 2000] and [Mahoney & Chan 2003]

- o Lack of relevant public datasets to test the algorithms ...
- o KDD99 ? See [McHugh 2000] and [Mahoney & Chan 2003]
- o We rather use MAWI
 - 15 min a day of real traffic (.pcap file)
 - Anomaly patterns given by the MAWILab [Fontugne *et al.* 2010] with taxonomy [Mazel *et al.* 2014]


- o Lack of relevant public datasets to test the algorithms ...
- o KDD99 ? See [McHugh 2000] and [Mahoney & Chan 2003]
- o We rather use MAWI
 - 15 min a day of real traffic (.pcap file)
 - Anomaly patterns given by the MAWILab [Fontugne *et al.* 2010] with taxonomy [Mazel *et al.* 2014]
- o Preprocessing step : raw .pcap → NetFlow format (only metadata)

AN EXAMPLE TO DETECT NETWORK SYN SCAN

- o The ratio of SYN packets : relevant feature to detect network scan [Fernandes & Owezarski 2009]


AN EXAMPLE TO DETECT NETWORK SYN SCAN

- The ratio of SYN packets : relevant feature to detect network scan [Fernandes & Owezarski 2009]



AN EXAMPLE TO DETECT NETWORK SYN SCAN

- o The ratio of SYN packets : relevant feature to detect network scan [Fernandes & Owezarski 2009]




- o Goal: find peaks

→ Parameters : $q = 10^{-4}$, $n = 2000$ (from the previous day record)

SPOT RESULTS

—o Parameters : $q = 10^{-4}$, $n = 2000$ (from the previous day record)




DO WE REALLY FLAG SCAN ATTACKS ?

- The main parameter q : a False Positive regulator


DO WE REALLY FLAG SCAN ATTACKS ?

- The main parameter q : a False Positive regulator



DO WE REALLY FLAG SCAN ATTACKS ?

- The main parameter q : a False Positive regulator



- 86% of scan flows detected with less than 4% of FP

A more general framework

- A single main parameter q
 - With a probabilistic meaning $\rightarrow \mathbb{P}(X > z_q) < q$
 - False Positive regulator

- A single main parameter q
 - With a probabilistic meaning $\rightarrow \mathbb{P}(X > z_q) < q$
 - False Positive regulator
- Stream capable
 - Incremental learning
 - Fast (~ 1000 values/s)
 - Low memory usage (only the excesses)

→ SPOT

- performs dynamic thresholding without distribution assumption
- uses it to detect network anomalies

→ SPOT

- performs dynamic thresholding without distribution assumption
- uses it to detect network anomalies

→ But it could be adapted to

OTHER THINGS ?

- o SPOT
 - performs dynamic thresholding without distribution assumption
 - uses it to detect network anomalies
- o But it could be adapted to
 - compute upper and lower thresholds

OTHER THINGS ?

- o SPOT
 - performs dynamic thresholding without distribution assumption
 - uses it to detect network anomalies
- o But it could be adapted to
 - compute upper and lower thresholds
 - other fields

—o SPOT

- performs dynamic thresholding without distribution assumption
- uses it to detect network anomalies

—o But it could be adapted to

- compute upper and lower thresholds
- other fields
- drifting contexts (with an additional parameter) → DSPOT

A RECENT EXAMPLE


→ Thursday the 9th of February 2017

- o Thursday the 9th of February 2017
 - **9h** : explosion at Flamanville nuclear plant


- o Thursday the 9th of February 2017
 - **9h** : explosion at Flamanville nuclear plant
 - **11h** : official declaration of the incident by EDF

- o Thursday the 9th of February 2017
 - **9h** : explosion at Flamanville nuclear plant
 - **11h** : official declaration of the incident by EDF
- o What about the EDF stock prices ?

EDF STOCK PRICES



EDF STOCK PRICES



CONCLUSION

- o Context: A great deal of work has been done to develop anomaly detection algorithms

CONCLUSION

- Context: A great deal of work has been done to develop anomaly detection algorithms
- Problem: Decision thresholds rely on either distribution assumption or expertise

CONCLUSION

- Context: A great deal of work has been done to develop anomaly detection algorithms
- Problem: Decision thresholds rely on either distribution assumption or expertise
- Our solution: Building dynamic threshold with a probabilistic meaning

CONCLUSION

- Context: A great deal of work has been done to develop anomaly detection algorithms
- Problem: Decision thresholds rely on either distribution assumption or expertise
- Our solution: Building dynamic threshold with a probabilistic meaning
 - Application to detect network anomalies

CONCLUSION

- Context: A great deal of work has been done to develop anomaly detection algorithms
- Problem: Decision thresholds rely on either distribution assumption or expertise
- Our solution: Building dynamic threshold with a probabilistic meaning
 - Application to detect network anomalies
 - But a general tool to monitor online time series in a blind way

CONCLUSION

- Context: A great deal of work has been done to develop anomaly detection algorithms
- Problem: Decision thresholds rely on either distribution assumption or expertise
- Our solution: Building dynamic threshold with a probabilistic meaning
 - Application to detect network anomalies
 - But a general tool to monitor online time series in a blind way
- Future: Adapt the method to higher dimensions