Detecting Anomalies in Smart IoT Environments

Coseners 2019

31st Multi-Service Networks workshop (MSN 2019)

Roman Kolcun

July 4, 2019

Goal of the research

- Design a system capable of detecting anomalies in data communication of IoT devices in home environment
- Able to detect and inform users that a device in their home is misbehaving
- Leverage crowdsourcing to generate models of behaviours (e.g. ML models) (more on this in the next presentation)

Why data from multiple sources is needed

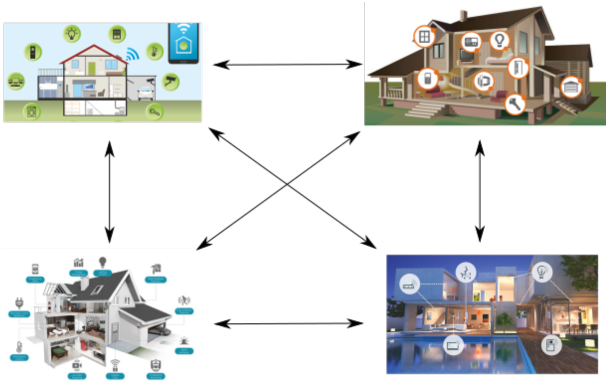

Figure: No two homes are the same.

Crowdsourcing

Figure: A figure depicting crowdsourcing.

Why is it not trivial?

- Communication of a device may depend on
  - region
  - occupancy of the home
  - other devices present on the network
  - installed third-party apps
- Creating models in a privacy-preserving manner
- Create and/or refine models on local router/gateway

How are we going to evaluate it?

There are two test-beds: one in NEU, the other at ICL

Figure: Northeastern University

How are we going to evaluate it?

There are two test-beds: one in NEU, the other at ICL

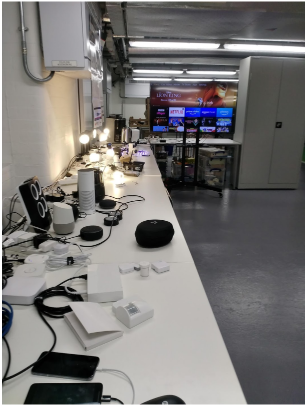

Figure: Imperial College London

Advantages of multiple testbeds

- Data collected in the same way
- It is possible to study differences depending on regions
- Possibility to validate models in multiple locations
- Develop and evaluate algorithms for federated machine learning

What is missing?

- There are very few public data-sets
- Most of the papers use the one published by UNSW
- ML models trained on high-end computers
- Complexity and model size is rarely mentioned
- "Smarter" smart devices (i.e. which support third-party apps) are not considered

Analysis of collected data

- Usage of encryption
- Analysis of the content of network communication
- We also analyse region-based differences

Figure: Figure symbolising encryption

Encryption

- Almost half (46%) of traffic cannot be classified by tools such as WireShark
- This traffic can be classified using entropy analysis (higher entropy suggests encrypted data)
- There are some positive trends where none of the devices send all traffic unencrypted
- However, most of the devices send some traffic unencrypted
- Significant amount of traffic cannot be easily determined and requires further research
- Usage of encryption also depends on region (e.g. a smart TV did not use encryption in the UK)

Analysis of the content

- Personally Identifiable Information (PII)
- Inference of device behaviour

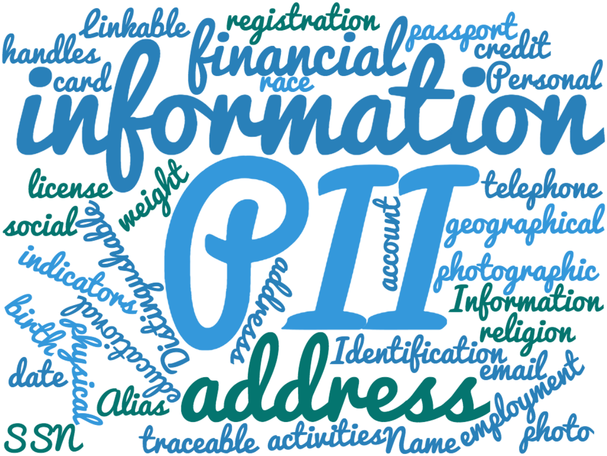

Figure: Figure symbolising PII

PII analysis

- We searched for MAC addresses, UUID, names, emails, etc. in the plaintext communication
- We found several PII exposures (e.g. MAC addresses or device name)
- A camera was sending a notification using HTTP to a server in China every time a motion was detected
- Some devices exposed some PII depending on region (e.g. a smart hub was leaking MAC address in the UK, not in the US)

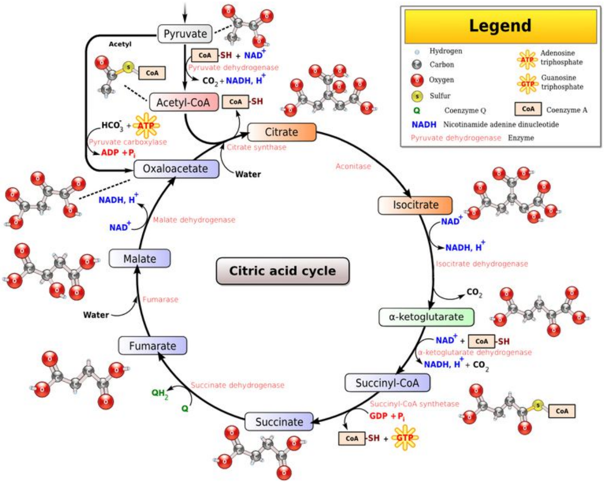

Figure: Krebs cycle (symbolising Inference of device behaviour)

Inference of device behaviour

- We used machine learning to guess the action a device performed
- We were able to predict significant amount of actions such as powering on, issuing a voice command, streaming video, etc.
- There are some regional differences in predictability of actions
- We used these models on "idle" traffic and found that some cameras are triggered by noise or some ambient movement

Questions?