Deloitte.

Evaluating a frequency-based anomaly detection algorithm on large-scale vehicular CAN data

Vincent Kieberl & Silke Knossen

Supervisor: Colin Schappin, Deloitte

Cars then vs. cars now

Source: abroadintheyard.com

Source: slashgear.com

(Some of) a car’s sensors, actuators and control modules

Transmission fluid pressure
Spark knock
Oxygen sensor
Engine RPM
Vehicle speed
Gear shift control module
Engine speed
Emergency brake assist
Fuel temperature
Seat belt tensioner
Mass air flow
Electronic stability program
Camshaft position
Rain sensor
Transmission friction sensor
Maniford absolute pressure
Coolant temperature
Outside air temperature
Collision sensor
Oil temperature
Coolant fan
Intake air temperature
Rear view camera
Airbag igniter
ABS
Transmission fluid temperature
Parking aid
Battery voltage
Parking brake motor hall sensor
NOX sensor
Transmission control valve
Fuel pump control module
Power steering
AC refrigerant pressure sensor
Tire pressure monitoring
Intake manifold tuning valve
Cylinder glow plug control module
Diesel particulate filter pressure sensor
Coolant circulation pump
Adaptive lighting
NH3 sensor
Exhaust gas temperature
Ozone reduction catalyst temperature

The need for automotive networking

Central control unit

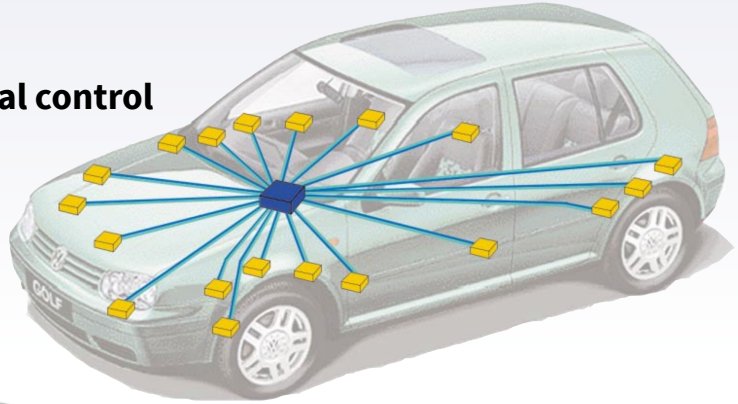

3 control units

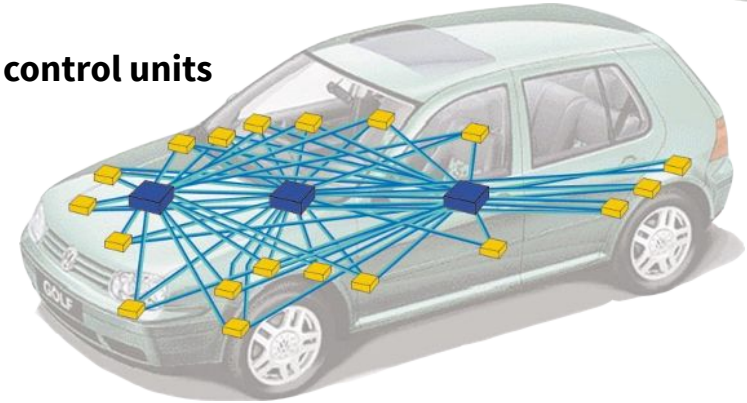

Bus system

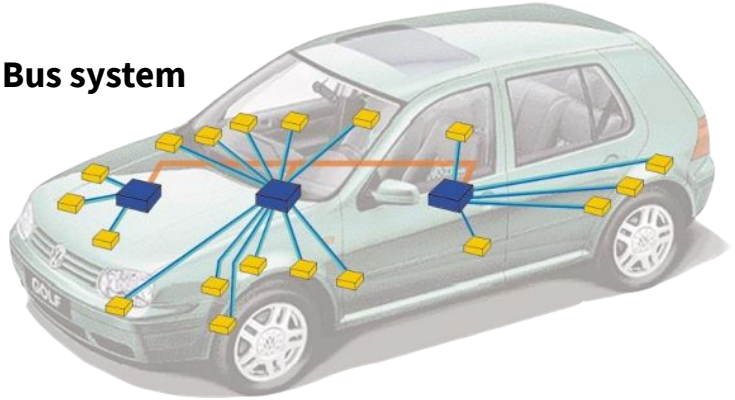

Source: Volkswagen AG, *Data Exchange on the CAN bus I*

Context

The CAN bus

- Controller Area Network (CAN)
- Interconnects Electronic Control Units (ECUs)
- Bus system, broadcast
- CAN IDs for identification
- Read out through OBD-2 port (On-Board Diagnostics)
- Only standardized in OSI layers 1 & 2

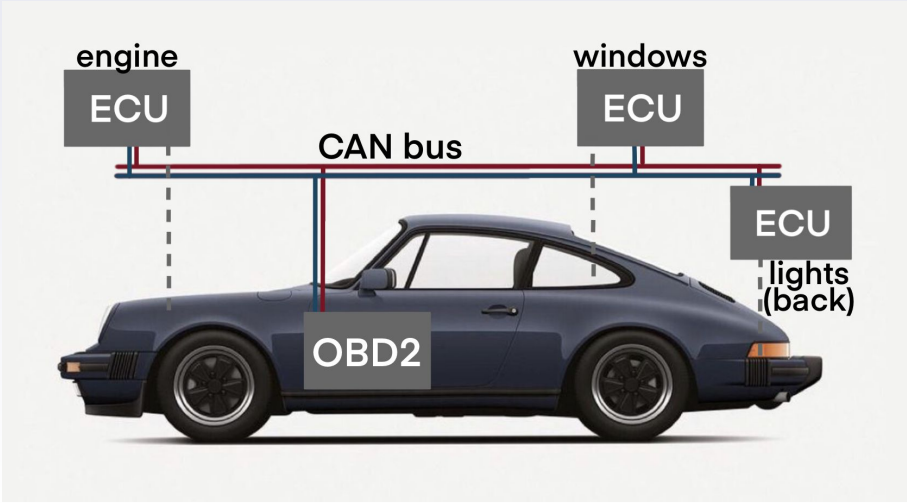

Source: Silke

Source: aimtechnologies.com

Context

Hacking a car using CAN

- Miller & Valasek's Jeep hack
- Inserting, modifying, or deleting frames
- Every ECU has one specific frequency
- Frequency changes when adding/removing frames

Related work

Taylor et al. 2015

- Frequency-based anomaly detection
- Inter-packet time (interval) best feature
- Only used insertion attacks

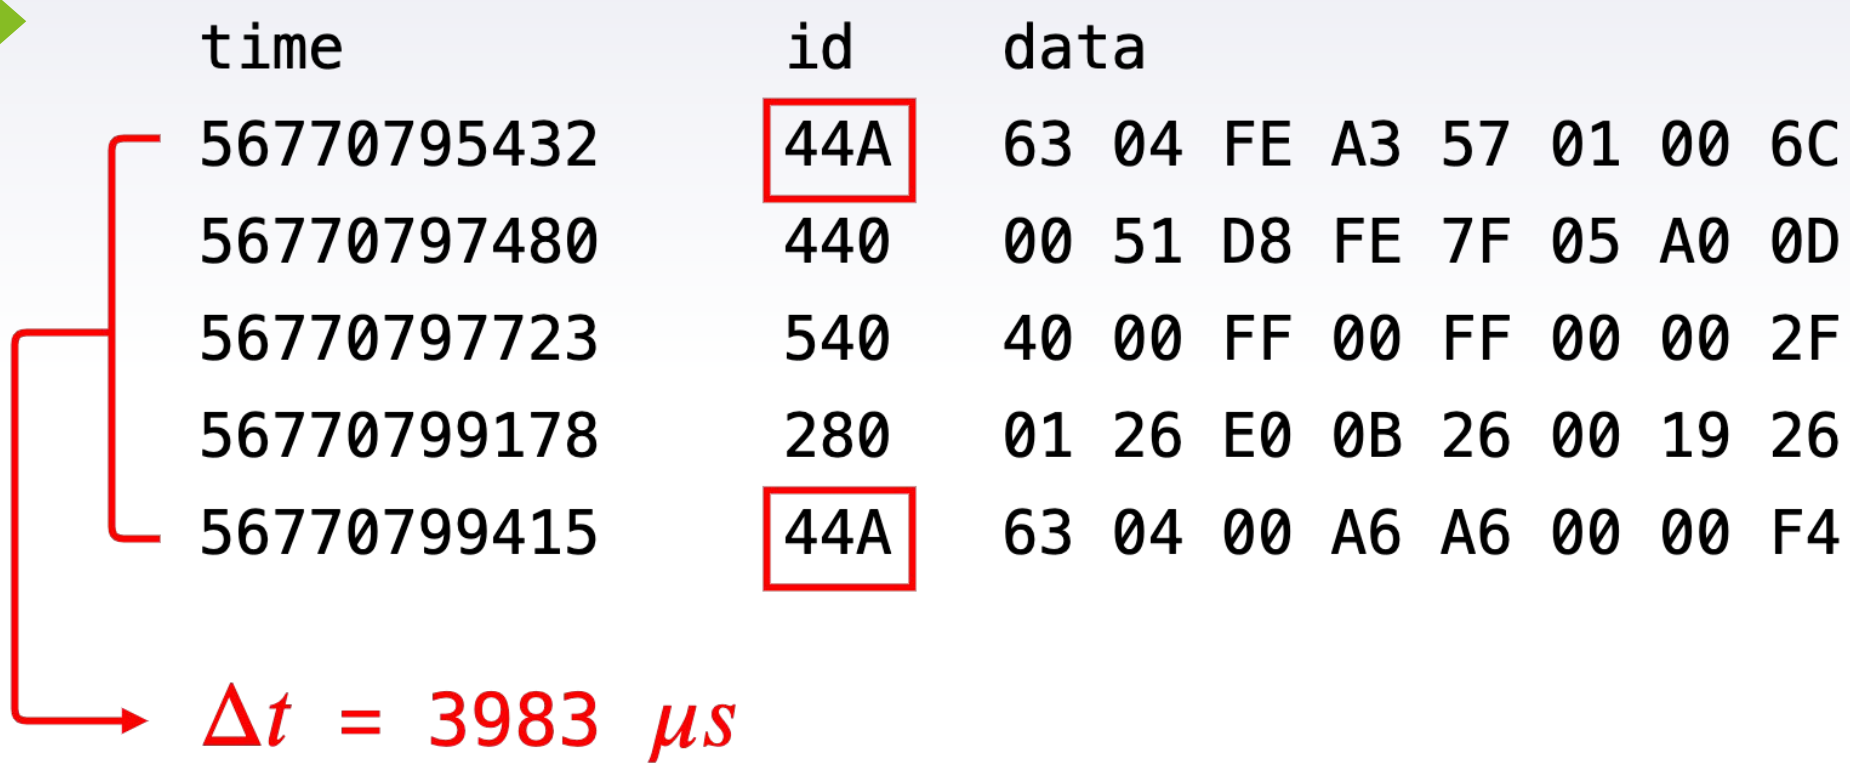
time id data
56770795432 44A 63 04 FE A3 57 01 00 6C
56770797480 440 00 51 D8 FE 7F 05 A0 0D
56770797723 540 40 00 FF 00 FF 00 00 2F
56770799178 280 01 26 E0 0B 26 00 19 26
56770799415 44A 63 04 00 A6 A6 00 00 F4
$\Delta t = 3983\ \mu s$

Related work

Schappin 2017

- Different types of attacks:
  - Fabrication attack: adding CAN messages
  - Suspension attack: deleting CAN messages
  - Masquerade attack: modifying CAN messages by adding them with ID and frequency of another ECU

Related work

Schappin 2017

- Robust Covariance Estimator (RCE)
- Split CAN IDs into 3 groups with 3 separate classifiers: fast/medium/slow
- Data from 2011 Dodge Ram, 4.5 minutes in total, of which 30 seconds test data
- Data may not resemble real-world situations

Research question

To what extent does the amount of training data influence the performance of the model based on the Robust Covariance Estimator (RCE) as proposed by [1] ?

[1] Schappin, 2017.

Subquestions

- How can we collect a dataset from a real vehicle that contains over 40 minutes of CAN data with microsecond accuracy?
- What are the differences in data characteristics in data from an Audi and a Ford vehicle?
- What is the influence of the amount of training data on the performance of the RCE on fabrication, suspension, and masquerade attacks?

Approach

Data acquisition

- PCAN USB FD connected to OBD2 port
- Tried on six cars of which two were successful
  - Audi A4 2006
  - Ford Fiesta 2017
- Min. 70 minutes of data

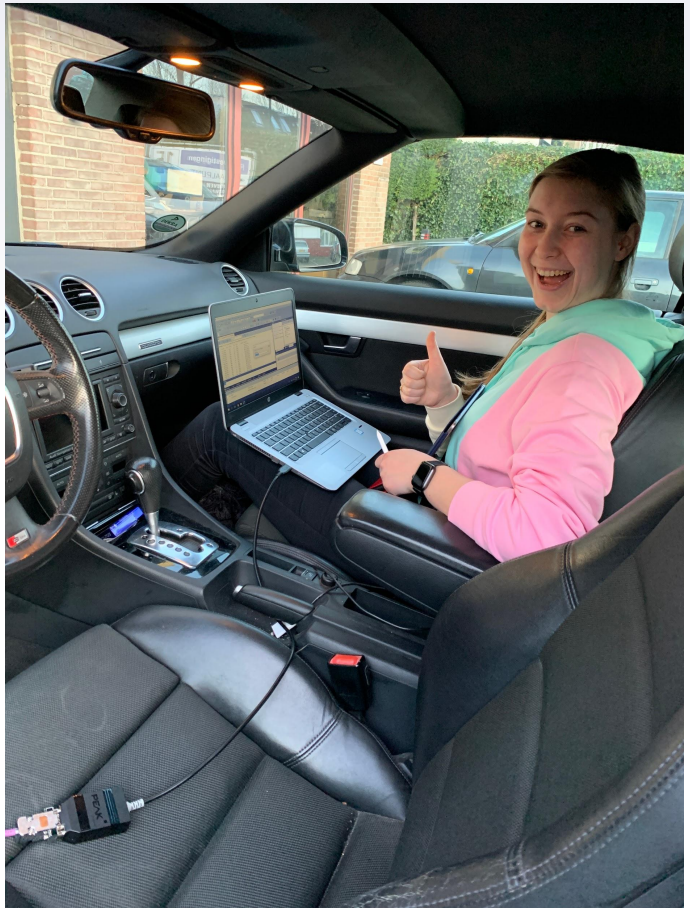

Approach

The data

- Audi A4 (2006)
  - 31 different CAN IDs
  - Interval range 10ms - 1s
  - All IDs throughout whole dataset
- Ford Fiesta (2017)
  - 51 different CAN IDs
  - Interval range 10ms - 10s
  - Two IDs only present in the first 5 minutes

Approach

The RCE algorithm

- One-class classification algorithm
- Three classifiers for different interval ranges
- Preprocessed data
  - Three matrices for the interval ranges
- Classify data per window

Data matrix for specific interval range

	ID 1	...	ID n
Window 1	**mean interval**	...	**mean interval**
...	...	...	...
Window n	**mean interval**	...	**mean interval**

Approach

Experiments

- Different sizes of training sets
  - 2; 5; 10; 20; 30; 45 minutes
- Simulating attacks by altering the testsets
  - Fabrication, suspension, masquerade
- Different attack sizes per attack
  - Small, medium, and large attacks
  - 1 frame; 25 frames; ⅓ of all frames

Fabrication attack - 1 message

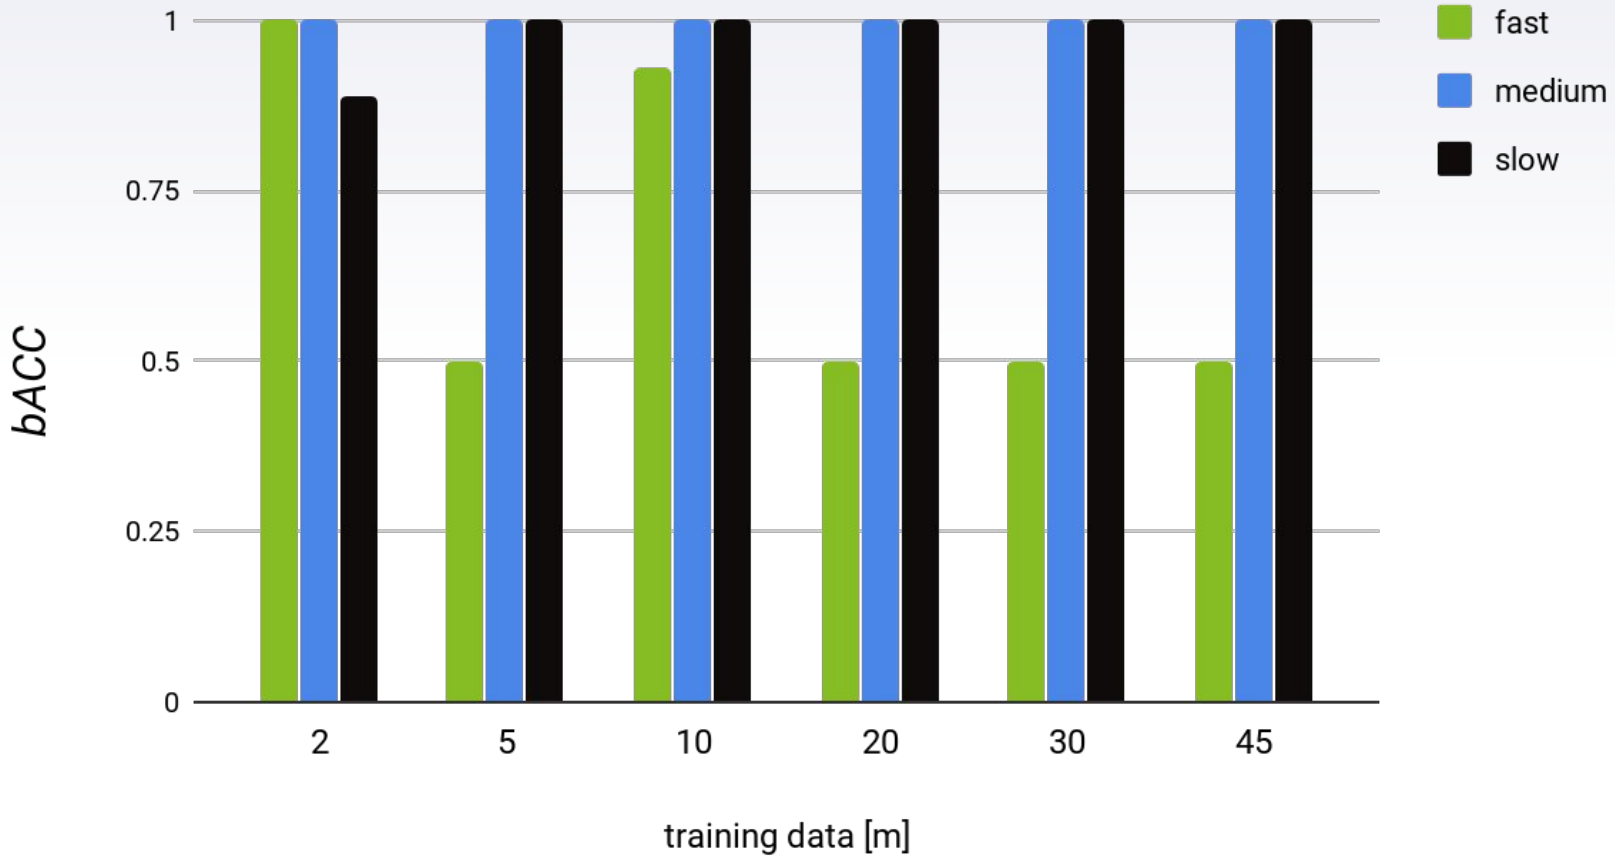

Masquerade attack (LS) - 1 message

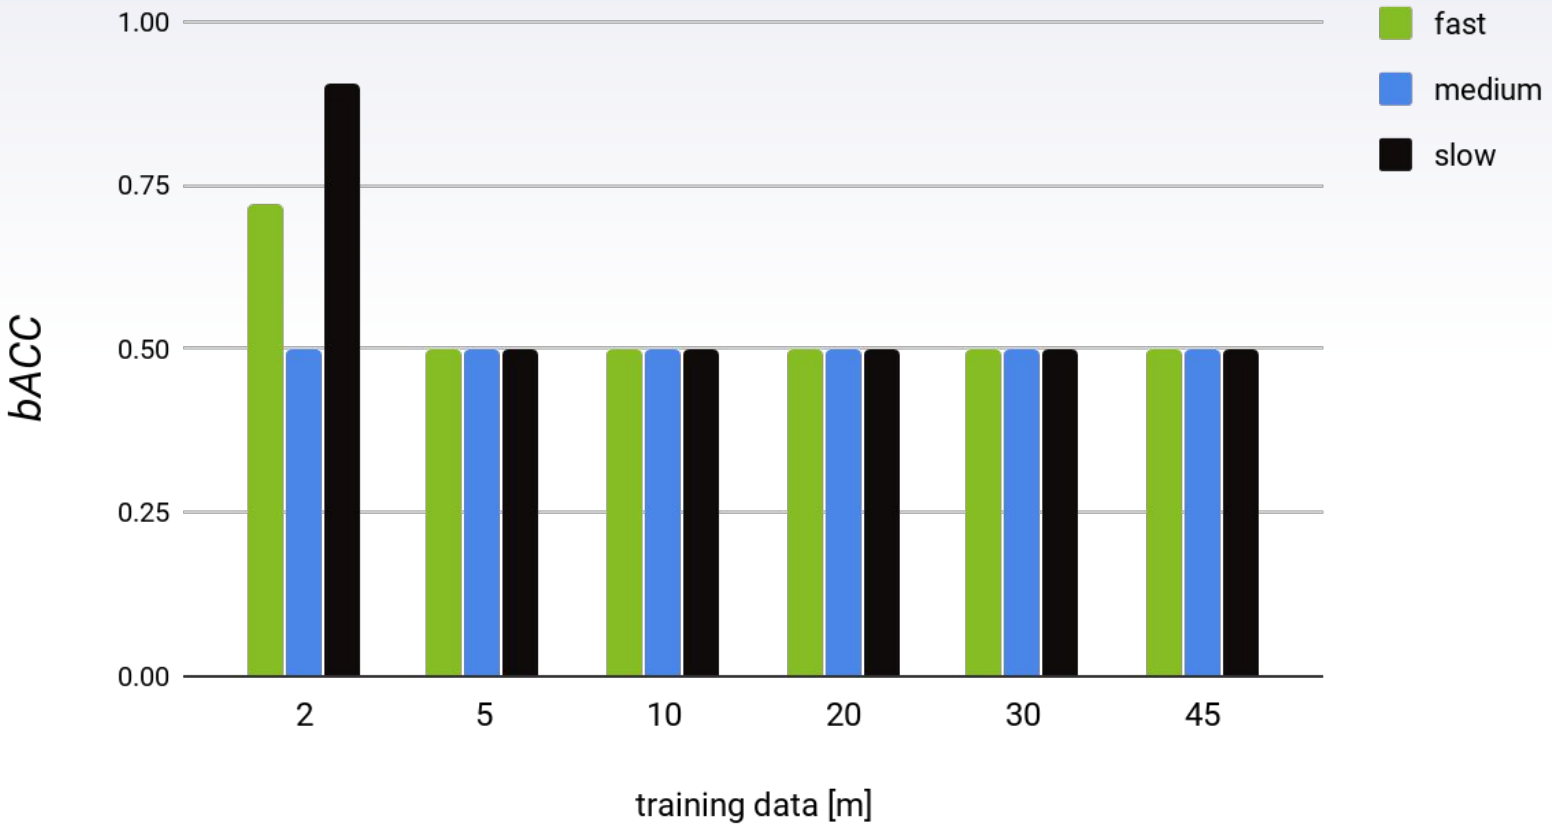

Suspension attack - 1 message

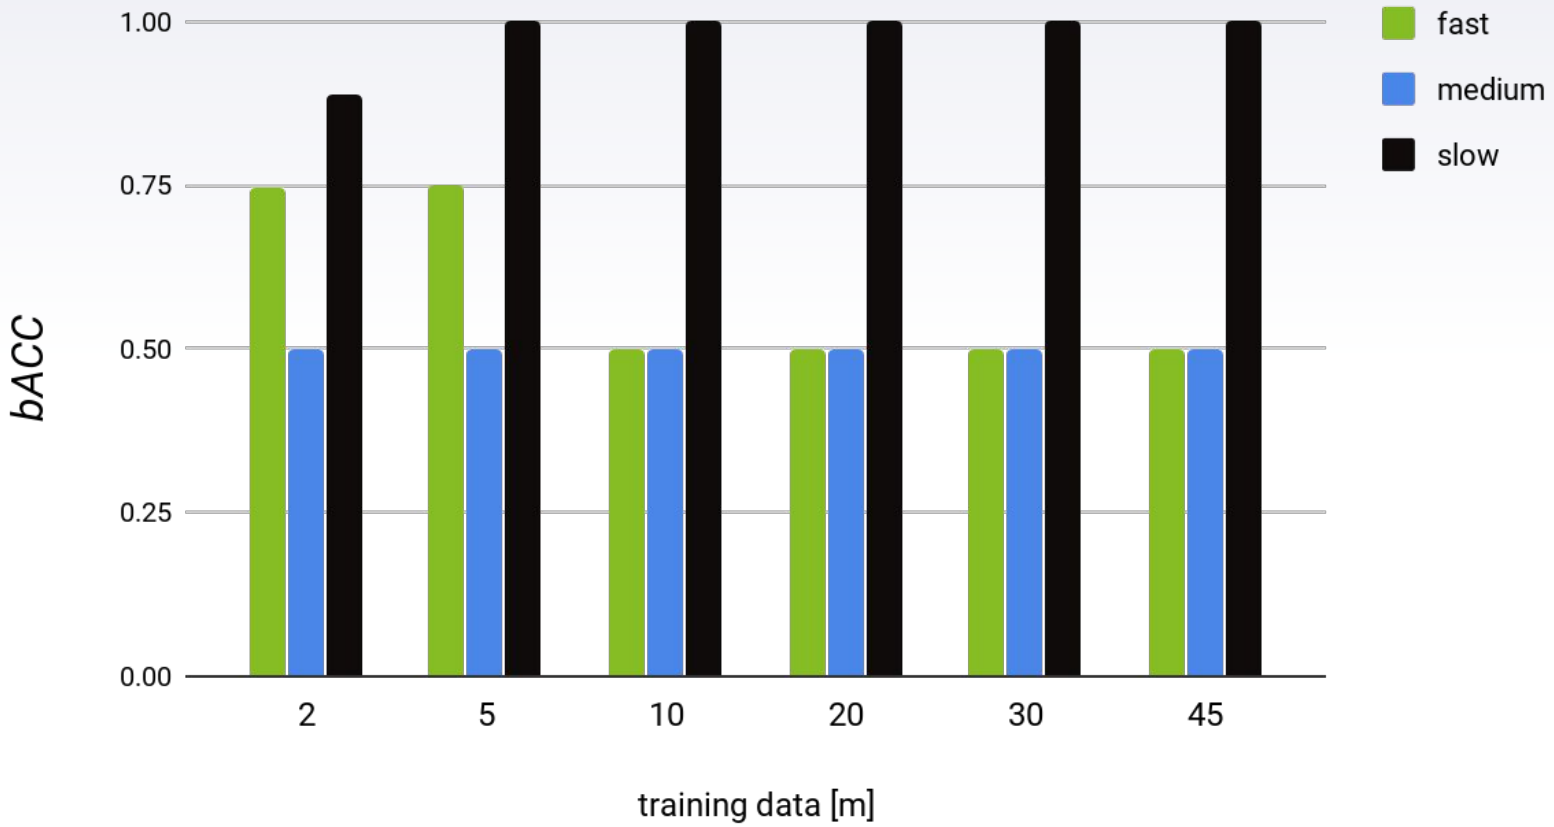

Suspension attack - 2500/500/50

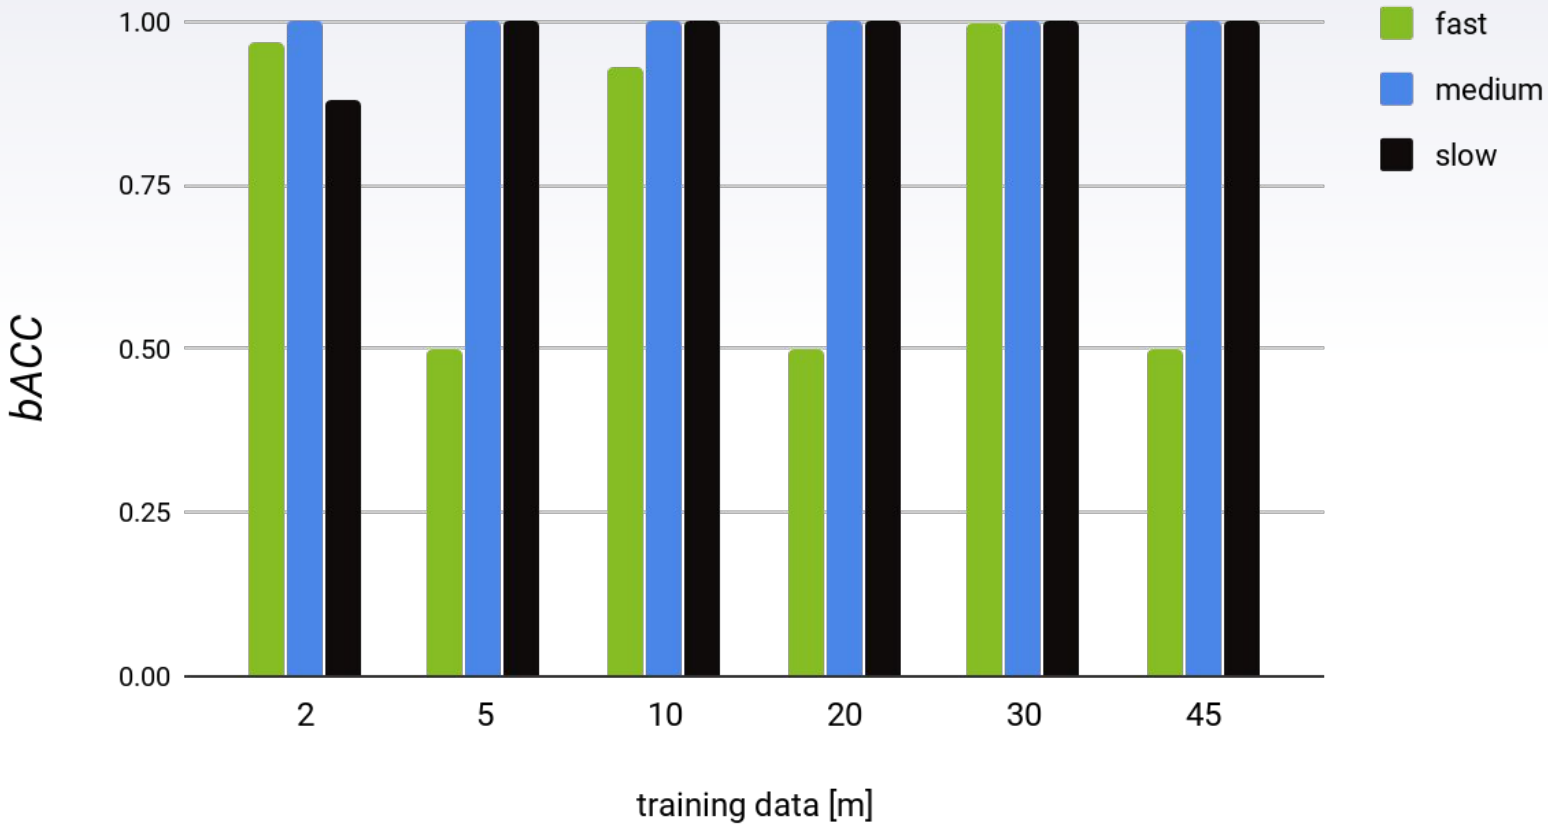

Conclusion

- Able to obtain CAN traffic with microseconds timestamps
- Different data for different vehicle models
- Amount of training data does not have significant influence
  - Depends on attack and CAN ID

Discussion

Limitations & future work

- Not all CAN IDs tested
- Only attack information is a time frame
- Non-recurring CAN frames
- Vehicle model specific
- Algorithm does not utilize CAN data field
- Proof of concept needs to work on input stream of data