Beatrix Koltai & András Gazdag

Anomaly detection in CAN with TCN

www.crysys.hu

Background

Modern cars are full of embedded controllers (ECUs), which are connected by internal networks, like CAN (Controller Area Network). ECUs communicate with each other mostly periodically, and signals are encoded in CAN messages. One message type usually contains more signals.

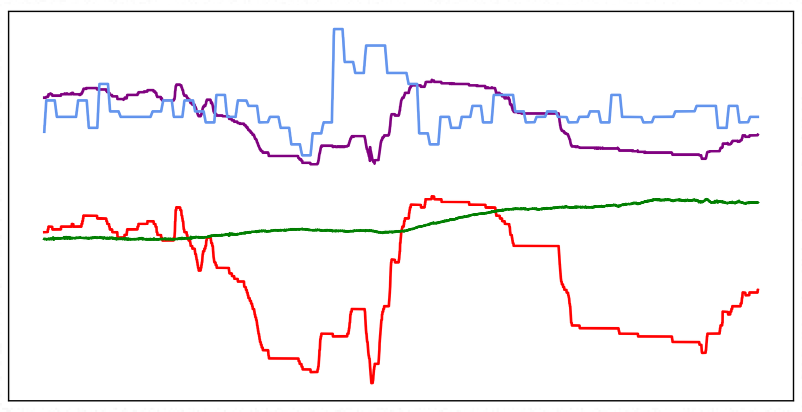

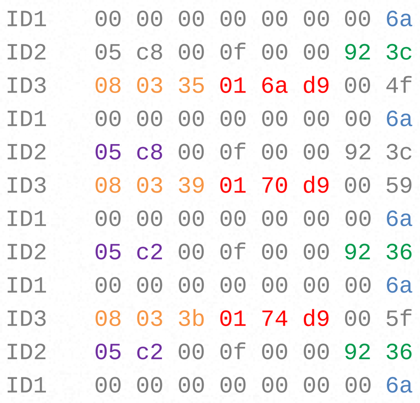

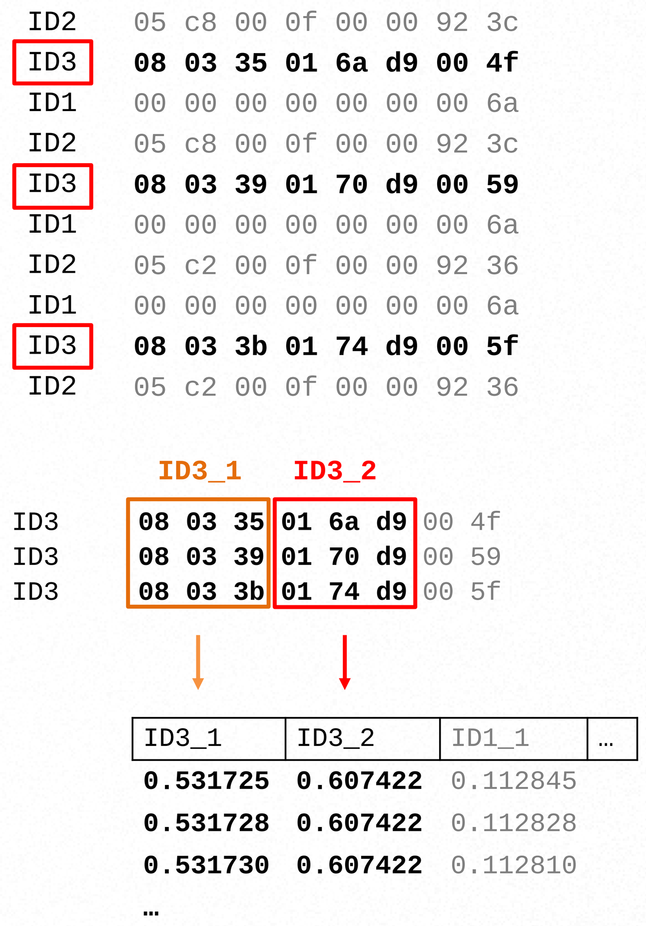

Message preprocessing

Step 1) Filter on message ID. Messages with the same ID will contain multiple signals.

Step 2) Signal extraction from messages. Slicing masks are calculated with the analysis of bit flip rate.

Step 3) Create input dataset. Input data will contain these individual signals from all message types. Missing values are filled with zeros.

Problem

Malicious modification of these messages must be detected for the safety of the vehicle.

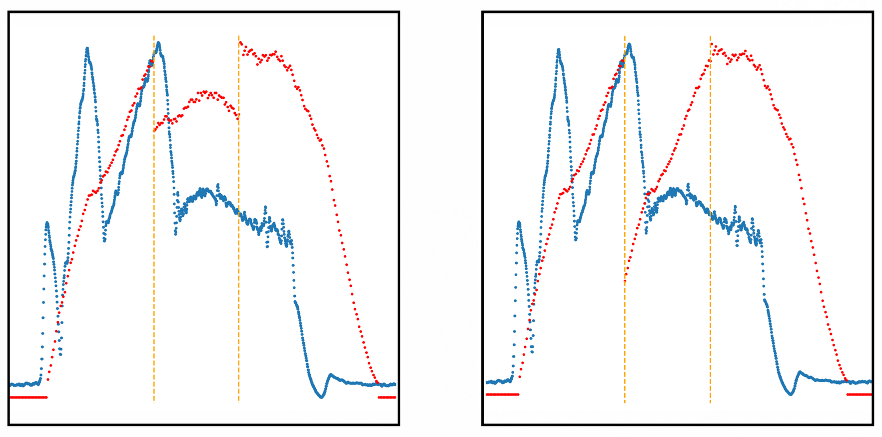

Dataset

The used dataset contains ~2,5 hours of traffic (~5,5 million messages) from different driving scenarios. Six different modification attacks were performed on the traces. During each attack only one signal is modified.

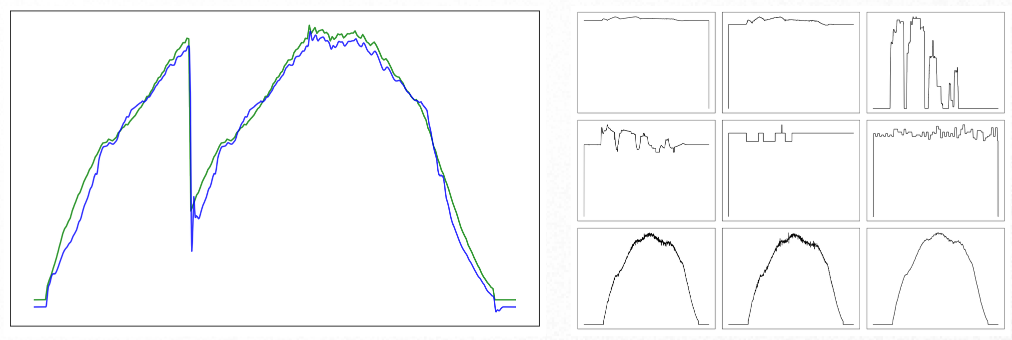

Correlation

Signals are grouped by their correlation. A multi-channel model is trained on each of these groups, where one channel corresponds to one signal.

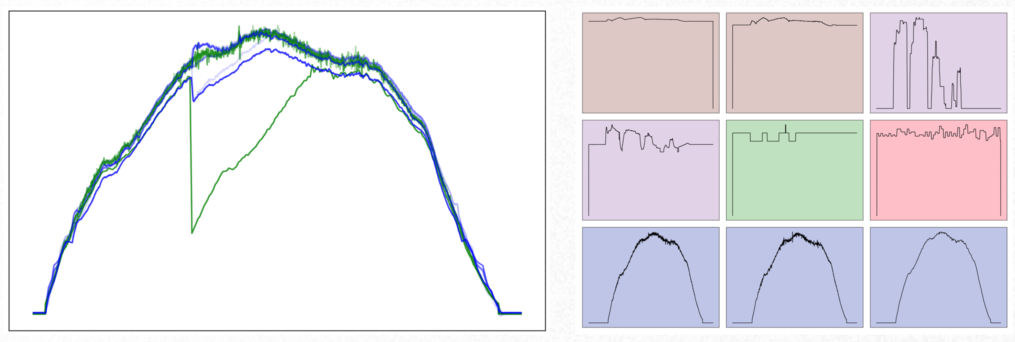

Solution

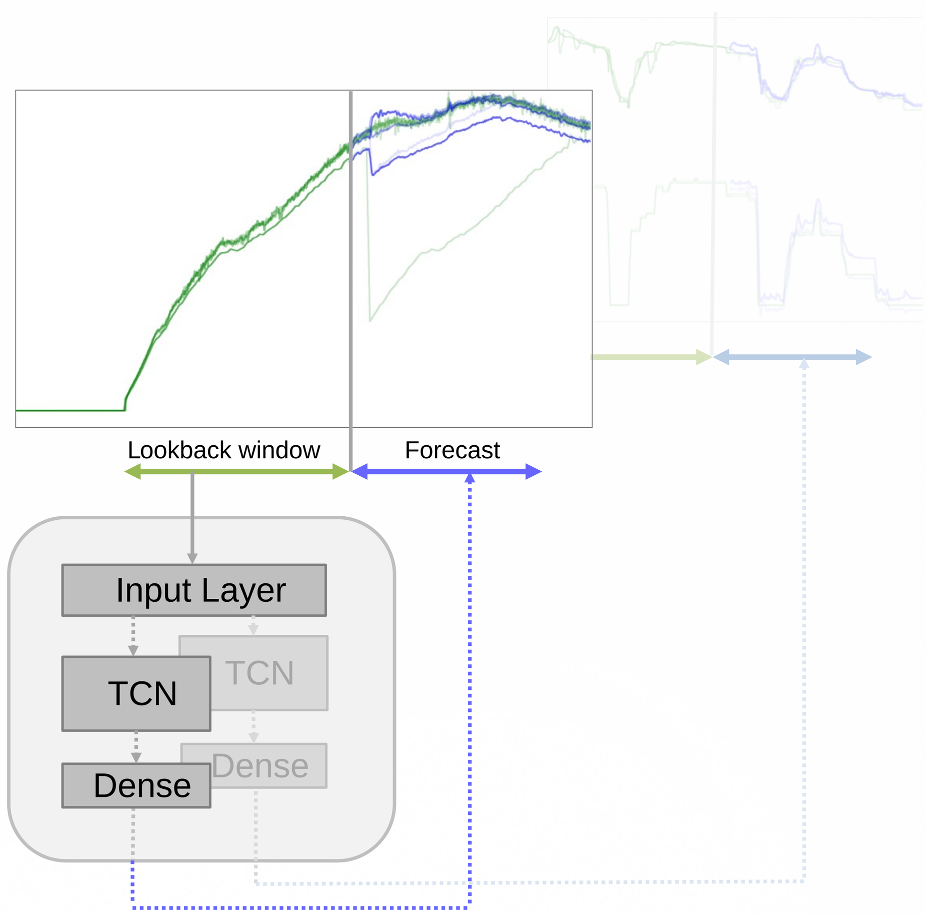

Predicting one signal at a time would not necessarily result in a detectable anomaly. However, the prediction of a group of signals will depend on their correlation, thus an attack in only one signal will alter the prediction of the whole group. This causes a more significant anomaly.

Model

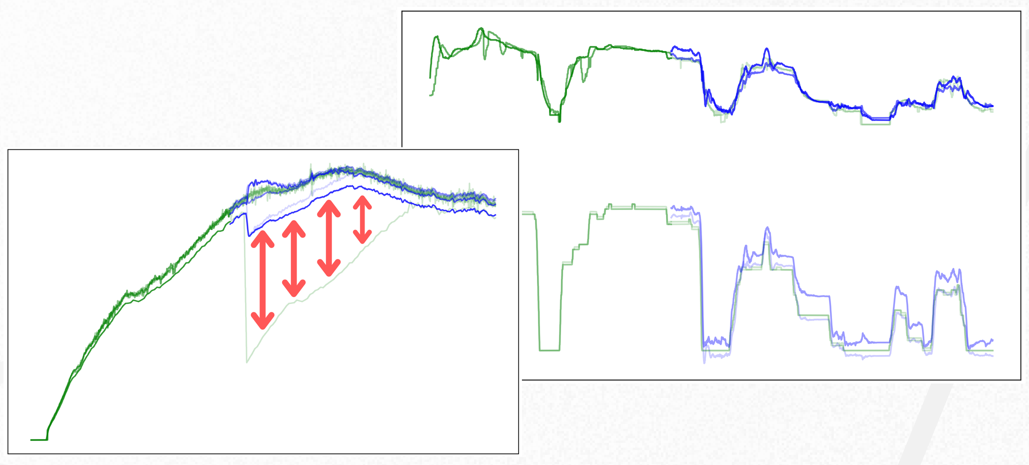

During training, the model receives the relevant signals in a sliding window through an input layer, which passes the correlated signals to a TCN layer. The TCN is composed of four dilatation layers, and predicts next value for each signal in the group by training on the values in the sliding window. The TCN layer is followed by a Dense layer, which produces the output corresponding to the group size.

Anomaly detection is done by comparing the prediction with the actual value.

Results

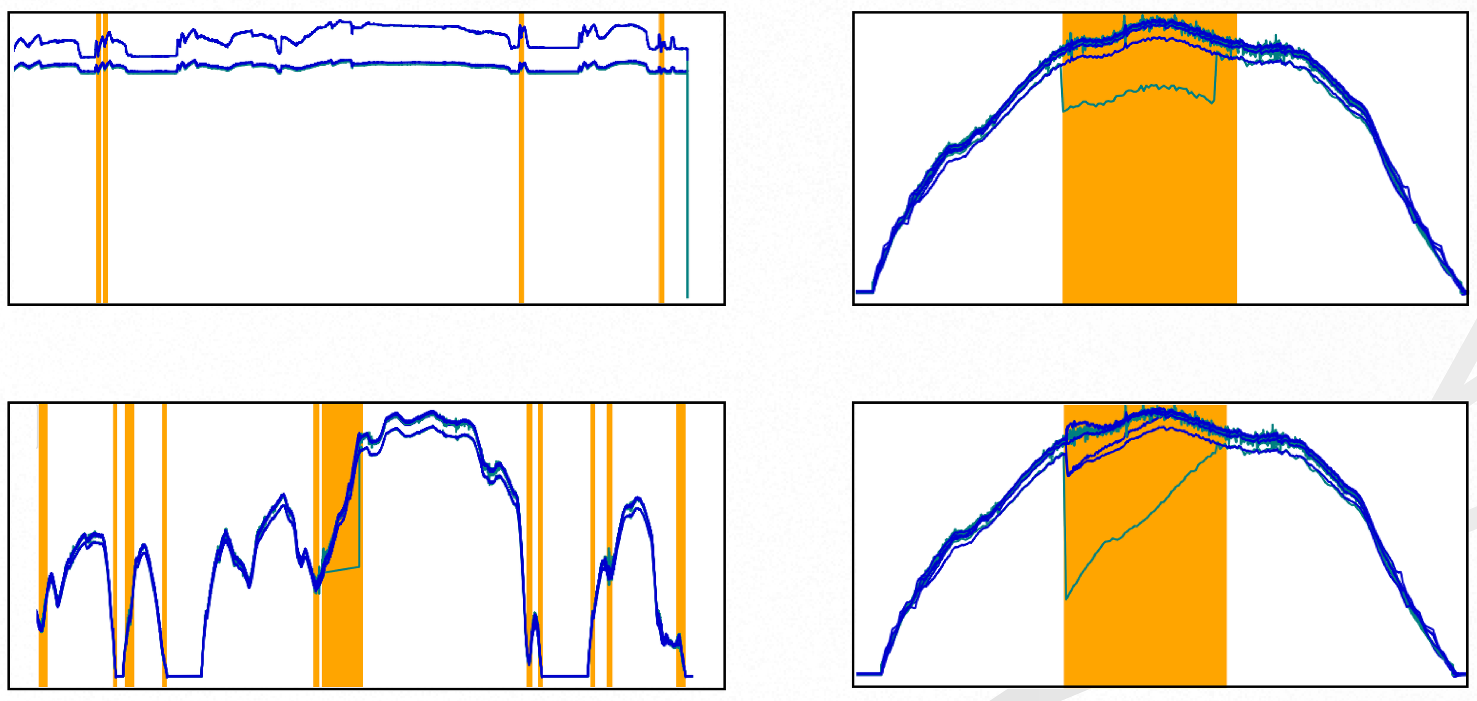

Predicting a group of signals allows us to detect attacks that modify the signal in a way that would otherwise be normal. Using this method, we were able to detect the majority of attacks, but attacks of different lengths were problematic. Choosing the right evaluation window size is an important task in the detection process.

Conclusion

- CAN anomalies could be efficiently detected with TCN networks.
- Grouping signals based on correlation improves anomaly detection and reduces the required resources.
- Finding the proper window and threshold values is essential for high accuracy and low FPR results.

Laboratory of Cryptography and System Security
Budapest University of Technology and Economics