CHALMERS
UNIVERSITY OF TECHNOLOGY

Analysis of AEB False Positives in Urban Bus Operations

A Criteria-Based Offline Workflow for Detecting False Positive AEB Braking Events Using MF4 Data and TTC Validation

Master's thesis in Systems Control and Mechatronics

YOULI GONG

JINGYU WANG

DEPARTMENT OF ELECTRICAL ENGINEERING

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's Thesis report 2026

Analysis of AEB False Positives in Urban Bus Operations

A Criteria-Based Offline Workflow for Detecting False Positive AEB Braking Events Using MF4 Data and TTC Validation

YOULI GONG
JINGYU WANG

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
Chalmers University of Technology
Gothenburg, Sweden 2026

Analysis of AEB False Positives in Urban Bus Operations
A Criteria-Based Offline Workflow for Detecting False Positive AEB Braking Events Using MF4 Data and TTC Validation
YOULI GONG & JINGYU WANG

Supervisor:
Jordanka Kovaceva, Department of Mechanical Engineering, Chalmers University of Technology
Zain Zubair, Team Driver Assistance, Volvo Buses

Examiner:
Nikolce Murgovski, Department of Electrical Engineering, Chalmers University of Technology

Master's Thesis report 2026
Department of Electrical Engineering
Chalmers University of Technology
SE-412 96 Gothenburg
Sweden
Telephone +46 31 772 1000

Cover: by Sergio Souza / Pexels, used under the Pexels license.

Typeset in LaTeX
Gothenburg, Sweden 2026

Analysis of AEB False Positives in Urban Bus Operations
A Criteria-Based Offline Workflow for Detecting False Positive AEB Braking Events Using MF4 Data and TTC Validation
YOULI GONG & JINGYU WANG
Department of Electrical Engineering
Chalmers University of Technology

Abstract

False positive (FP) activations of Automatic Emergency Braking (AEB) in urban buses can cause unnecessary harsh braking and reduce system acceptance for situations involving both passengers and vulnerable road users. This report implements a traceable, criteria-based pipeline to identify candidate FP events from raw ASAM MDF 4 (MF4) recordings, structured as file-level screening, qualified brake-event detection, and FP classification. The implementation uses core Controller Area Network (CAN) signals to detect AEB-active braking events and applies collision-risk and driver-response criteria for classification, where time-to-collision (TTC) is computed from recorded signals and used as one decision input. The TTC computation matches a ground-truth TTC reference provided by a colleague within the Group Trucks Technology team on the evaluated data, and the end-to-end detection pipeline runs on an internal dataset to automatically extract candidate FP events for further attribution and performance analysis.

Keywords: Automatic Emergency Braking, false positive detection, urban buses, criteria-based classification, CAN signal analysis

Acknowledgements

Firstly, we would like to acknowledge, with our deepest gratitude, Zain Zubair at Volvo Buses and Jordanka Kovaceva at Chalmers University of Technology. They have supervised this project and provided us with essential guidance, feedback, and support throughout the work.
We would also like to acknowledge and emphasize the significant contributions of our examiner Professor Nikolce Murgovski at the Department of Electrical Engineering at Chalmers University of Technology.
We would also like to thank our colleagues Haiguang Chen, Nikolaos Mardas, Anton Pettersson, Sreelekha Velugotti, Gustav Wikman, in the Team Driver Assistance at Volvo Buses for valuable discussions, for providing access to the dataset, and for practical hands-on support with data access and tooling. Additionally, We would also like to express our sincere gratitude to our manager Louise Jonsson at Volvo Buses for recruiting us to this project and providing us with the opportunity to carry out this work. Furthermore, we thank Martin Claesson for providing the ground-truth TTC reference used for validation.
Lastly, we would like to acknowledge Volvo Buses for hosting the project and enabling this work.

Youli Gong & Jingyu Wang, Gothenburg, June 2026

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis, listed in alphabetical order:

ACC	Adaptive Cruise Control
ADAS	Advanced Driver Assistance Systems
AEB	Autonomous Emergency Braking
ASAM	Association for Standardisation of Automation and Measuring Systems
BLF	Vector Binary Logging Format
CAKF	Constant-Acceleration Kalman Filter
CAN	Controller Area Network
CM	Collision Mitigation
CVKF	constant-velocity Kalman filter
DACU	Driver Assistance Control Unit
DBC	Database Container
ECU	Electronic Control Unit
eTTC	Enhanced Time-to-Collision
Euro NCAP	European New Car Assessment Programme
FCW	Forward Collision Warning
FLC	Forward-Looking Camera
FLR	Forward-Looking Radar
FOT	Field Operational Test
FP	False Positive
FSM	Finite State Machine
GBT	Global Buses Technology
GT	Ground Truth
GTT	Group Trucks Technology
IQR	Interquartile Range
ISO	International Organization for Standardization
LPoB	Last Point of Braking
MF4	ASAM MDF 4
NIS	Normalised Innovation Squared
QBE	Qualified Brake Event
RMSE	Root Mean Square Error
ROC	Receiver Operating Characteristic
RTS	Rauch–Tung–Striebel Smoother

SAE	Society of Automotive Engineers
SOTIF	Safety of the Intended Functionality
TfL	Transport for London
TLC	Time-to-Lane-Crossing
TP	True Positive
TTC	Time to Collision
UTC	Coordinated Universal Time
VRU	Vulnerable Road User
XBR	External Brake Request

Nomenclature

Below is the nomenclature of indices, parameters, and variables that have been used throughout this thesis.

Indices

i	Sample index
k	Time step index (Kalman filter)
N	Total number of measurements in a sequence

Kinematic Variables

d	Longitudinal headway distance (m)
v	Ego vehicle speed (km/h)
v_{ego}	Ego vehicle speed (m/s)
v_{rel}	Relative closing speed (m/s)
a_{rel}	Relative longitudinal acceleration (m/s^2)
a_x	Longitudinal acceleration (m/s^2)
$\dot{a}_x$	Longitudinal jerk, $\approx \Delta a_x / \Delta t$ (m/s^3)
Δt	Time step between samples (s)

TTC Definitions

TTC	Time-to-collision (first-order model): $-d \,/\, v_{\text{rel}}$; assumes both vehicles maintain current speed (s)
eTTC	Enhanced TTC: smallest positive root of $d + v_{\text{rel}}\, t + \frac{1}{2} a_{\text{rel}}\, t^2 = 0$; incorporates relative acceleration (s)
$\text{TTC}_{\text{threshold}}$	Dynamic TTC threshold: $\max\Big(1.4,\ v_{\text{ego}} \,/\, (2\, a_{\text{LPoB}})\Big)$ (s)
a_{LPoB}	Braking deceleration used in the TTC-threshold definition associated with the Last Point of Braking ($3.0\,\text{m/s}^2$)

Kalman Filter & RTS Smoother

z_k	Measurement at time step k
$\hat{x}_{k\|k-1}$	Predicted (a priori) state estimate
$\hat{x}_{k\|k}$	Updated (filtered) state estimate
$\hat{x}_{k\|N}$	RTS smoothed state estimate
$P_{k\|k-1}$	Predicted error covariance
$P_{k\|k}$	Updated error covariance
K_k	Kalman gain
C_k	RTS smoother gain
F	State transition matrix
Q	Process (transition) noise covariance
R	Observation noise covariance

Brake Simulation Parameters

$a_{\text{XBR}}(t)$	Reconstructed XBR deceleration command
a_{init}	Initial deceleration ($-0.5\,\text{m/s}^2$)
$\dot{a}_{\text{down}}$	Deceleration ramp rate / jerk ($-4.5\,\text{m/s}^3$)
$a_{\max}$	Maximum deceleration ($-5.0\,\text{m/s}^2$)
$a_{\text{brake}}[i]$	Realised brake deceleration at sample i
α	Exponential decay factor: $e^{-\Delta t/\tau}$
k_{gain}	Brake realisation gain (0.78)
τ_{rise}	Brake lag rise time constant (0.20 s)
τ_{fall}	Brake lag fall time constant (0.40 s)

Qualified Brake Event Criteria

$q[i]$	Qualified braking event Boolean mask at sample i
$c_{\text{decel}}[i]$	Deceleration gate: $a_{\text{brake}}[i] \leq -1.5\,\text{m/s}^2$
$c_{\text{aeb}}[i]$	AEB active: $\texttt{CM\_Status}[i] \in \{2, 3\}$
$c_{\text{gate}}[i]$	Minimum speed gate: $v[i] > 10\,\text{km/h}$
$c_{\text{range}}[i]$	Scenario-specific speed range gate
$e_{\text{start}}[i]$	Rising edge flag: $q[i] \wedge \neg\, q[i-1]$
Δt_{merge}	Gap merge threshold (1.0 s)

Contents

List of Figures

List of Tables

1
Introduction

1.1 Background

Urban buses operate in dense, mixed traffic environments where interactions with vulnerable road users (VRU), such as pedestrians and cyclists, are frequent. Advanced driver assistance systems (ADAS), including Forward Collision Warning (FCW) and Autonomous Emergency Braking (AEB), are therefore increasingly mandated and deployed to mitigate imminent conflicts and reduce crash severity [1, 2]. The case for requiring AEB on city buses has been articulated by regulatory bodies and researchers who have developed specific test procedures for this vehicle class [3]. Broader estimates suggest that crash avoidance technologies—including AEB—could prevent a substantial share of large-vehicle rear-end collisions [4, 5], and research programmes have sought to accelerate voluntary adoption on commercial fleets [6]. Recent evaluation studies continue to refine how AEB performance is assessed across test protocols and operational contexts [7, 8]. Real-world evidence from heavy-truck fleets indicates that FCW and AEB can reduce police-reportable crashes, with particularly strong effects for rear-end crash types [9], findings corroborated by field studies on light vehicles [10] and Society of Automotive Engineers (SAE) evaluations of commercial forward collision avoidance systems [11]. Complementary naturalistic analyses further show that full mitigation braking events are relatively rare and that drivers often brake before later intervention stages, underlining the complexity of driver–system interaction and the importance of interpreting activation logs in their operational context [12].
However, effectiveness in crash reduction is only one dimension of system performance, especially for buses. In urban bus operation, sudden or unnecessary braking can negatively affect passenger comfort and may cause standing-passenger injuries [13], while also creating secondary risks for surrounding traffic. Large-truck safety reviews identify false activations as a critical adoption barrier [14], and adaptive warning systems that account for differential driving behaviour have been proposed to reduce unnecessary alerts [15]. As a result, false or nuisance activations—hereafter referred to as false positives (FP)—are not only an engineering issue but also an acceptance issue. Naturalistic heavy-vehicle studies have documented false activations in crash avoidance systems, including stationary-object-related false alerts triggered by infrastructure such as overpasses and signage, and have highlighted that FP events can erode driver trust and reduce system acceptance [16]. Similar acceptance challenges have been reported in the transit-bus domain, where operators frequently reported false alarms and distraction in collision avoidance warning

system deployments [17]. These findings motivate systematic approaches that can identify, quantify, and characterise candidate FP events from operational data.
A common approach for such studies is offline, event-based analysis of recorded fleet data, a strategy already established in large naturalistic driving programs such as UDRIVE, SHRP 2, and euroFOT [18–20], where interpretable surrogate measures are used to reason about risk and timing. Time-to-collision (TTC) is one widely used surrogate safety measure, defined as the time remaining until two road users collide if they maintain their current speed and heading [21]. TTC appears in both research and deployed warning logic; for example, naturalistic studies of blind-spot monitoring systems use TTC-based thresholds to trigger collision warnings, illustrating the practical relevance of TTC as a risk indicator in heavy vehicle and bus contexts [22]. At the same time, TTC computed from recorded signals is sensitive to signal definitions, object selection, and timing alignment, which motivates careful implementation and validation when TTC is used as a decision input.
In urban bus AEB deployment, the core challenge is not only assessing system effectiveness but also establishing a traceable and scalable engineering workflow for identifying false positive events from large volumes of operational fleet data. Existing studies provide useful evidence on crash-reduction performance and driver response, but they offer limited guidance on how raw vehicle recordings can be transformed into analysis-ready FP event catalogues for routine engineering investigation. Most existing studies focus on crash-avoidance effectiveness or driver-response behaviour, whereas broader scenario-based safety assessment, counterfactual benefit estimation, and predictive risk modelling are typically addressed in separate streams [23–25]. The practical challenge of processing large volumes of raw CAN recordings, applying traceable FP definitions, and producing analysis-ready event catalogues remains under-addressed.
In this thesis, we address this gap by developing and applying an offline analysis pipeline to raw ASAM MDF 4 (MF4) recordings from an internal Global Buses Technology (GBT) dataset collected in the Transport for London (TfL) operational context. The FP definition is based on whether it satisfies Condition A (TTC-based excessive margin at brake initiation) or Condition B (absence of driver brake response within a defined reaction-time window after AEB activation).

1.2 Purpose

The purpose of this thesis is to develop and demonstrate a traceable, criteria-based workflow that can automatically identify candidate FP AEB-associated braking events in urban bus operational data. The workflow is intended to support systematic engineering analysis by:

- extracting candidate events from large volumes of raw recordings in a scalable and reproducible manner,
- providing an auditable FP definition aligned with the project criterion document (Condition A and Condition B),
- enabling subsequent attribution analysis by outputting event timestamps and supporting signal traces.

The thesis is written following general principles for technical report structure and clarity, including explicit positioning of contributions and clear separation between background, methods, and results [26].

1.3 Research Questions

The following research questions guide the analysis:

RQ1 What is the FP rate of AEB-associated braking events in the TfL urban bus fleet under the adopted Condition A / Condition B definition?

RQ2 How are the identified FP events distributed across rule-based groups defined by radar target presence and TTC-to-threshold relationship (G0–G3), and what do these groups reveal about the dominant physical circumstances of false activations?

RQ3 To what extent do observable driver behaviour signals (brake pedal, accelerator, steering) at the moment of AEB activation correlate with the FP/TP classification, and can these signals support attribution of FP root causes?

1.4 Goals

The goals of the work are:

1. **Implement a traceable offline pipeline on MF4 data:** Build an end-to-end workflow that ingests raw MF4 recordings and extracts the core CAN-based signals required for scalable file screening, event detection, target reconstruction, TTC assessment, and driver-response assessment.
2. **Detect and time-localise qualified brake events (QBE):** Identify AEB-associated braking events using explicit qualification logic based on braking intensity, AEB activation state, and scenario speed gating, and produce a structured event timeline with key anchor timestamps.
3. **Construct a traceable ego–target kinematic timeline and compute TTC as a decision input:** Reconstruct the target trajectory from multi-sensor object signals through temporal alignment, target association, identity selection, and ego–target merging, then compute TTC/Enhanced Time-to-Collision (eTTC) with a clearly defined time reference so that the collision-risk estimate can be used consistently in Condition A evaluation.
4. **Perform a preliminary ground truth (GT)-based assessment of the TTC computation:** Compare the computed TTC against an independent GT TTC reference provided by the Group Trucks Technology (GTT) team to calibrate key timing assumptions and establish an initial level of confidence in TTC as an input to the FP classification logic.
5. **Classify FP candidates with a traceable definition:** Classify each qualified brake event as a candidate FP or non-FP based on two explicit conditions (Condition A and Condition B, Section 2.3.3) that are traceable to specification and regulatory requirements.
6. **Produce analysis-ready outputs:** Export a structured list of candidate FP events with timestamps and supporting signals to enable downstream review,

clustering, and attribution analysis.

All goals follow the project workflow definition.

1.5 Scope and Demarcation

This work has the following demarcations:

- **Retrospective analysis only:** The workflow processes recorded MF4 data retrieved from fleet storage (S3) and do not run on the vehicle in real time nor modify the embedded AEB function. Although the pipeline supports automated incremental processing as new recordings become available, the analysis is inherently retrospective: conclusions are drawn from historical data rather than influencing the system in the field.
- **Criteria-based FP definition:** FPs are defined strictly according to the project criterion document (Condition A and Condition B). Events not captured by these criteria are outside the scope of this thesis, even if they may be perceived as nuisance braking by users.
- **Scope statement:** The analysis is restricted to a single bus platform. The analysis does not claim a fleet-wide safety impact.
- **Perception uncertainty not fully modelled:** The pipeline relies on recorded signals and assumes that the provided TTC and object-related information are sufficiently consistent for offline event reasoning. Residual uncertainty from perception and tracking is not exhaustively modelled.
- **Shadow-mode data context:** The analysed fleet data represent a shadow-mode or non-intervention observation context rather than the full closed-loop behaviour of an AEB system deployed in active control. Consequently, the recorded event traces may not fully reproduce the exact process or outcome of real AEB interventions in service. Driver behaviour signals observed around these events therefore provide correlational evidence for attribution, but do not establish a fully closed-loop behavioural response to an actual AEB system intervention.

2
Data and Problem Definition

This chapter describes the dataset used in this study, the vehicle platform and its onboard sensing architecture, and the formal problem definition for AEB FP detection. Section 2.1 introduces the data sources and their characteristics. Section 2.2 details the CAN bus architecture and signal taxonomy. Section 2.3 formalises the FP detection problem, including the regulatory criteria that underpin the classification logic.

2.1 Data Sources

The data used in this project originate from two complementary sources: a large-scale Field Operational Test (FOT) conducted in real urban traffic, and a set of controlled track tests following European New Car Assessment Programme (Euro NCAP)-style protocols [1]. Both datasets were collected on Volvo city bus platforms. The FOT data are recorded in the MF4 format, while the track test data use the Vector Binary Logging Format (BLF). Both formats are ingested by the pipeline's data loading stage.

2.1.1 Field Operational Test Data

The primary dataset consists of naturalistic driving recordings from a fleet of 20 Volvo city buses operating on regular service routes in London under the TfL programme. Each vehicle is identified by a fleet code (B-2035 through B-2055), and the recordings cover the period from December 2023 to May 2026. The data are stored on an Amazon S3 object store and accessed programmatically via the project pipeline.

Each recording session produces two MF4 files:

- A **CAN-only file** (`*_CAN.mf4`), containing all decoded CAN signals at their native sampling rates.
- A **full file** (`*.mf4`), containing the same CAN data together with synchronised front-facing video.

The CAN-only files are used for automated batch processing due to their smaller size (typically 50–150 MB per session), while the full files are reserved for manual video validation of selected events. Table 2.1 summarises the FOT dataset characteristics. As of June 1, 2026, the analysis dataset comprises 3 160 CAN MF4 files from 20 vehicles. This data snapshot constitutes the basis for all FP classification results presented in Chapter 4.

Table 2.1: Summary of FOT dataset characteristics (data snapshot as of June 1, 2026).

Property	Value
Vehicle platform	Volvo city bus
Fleet size	20 vehicles (B-2035 – B-2055)
Operating environment	Urban routes, London (TfL)
Recording period	December 2023 – May 2026
Total CAN MF4 files	3 160
Data format	ASAM MDF 4 (MF4)
CAN channels	4 (J1939, Backbone, VicinityNet 1/2)
Raw signals per file	$\approx$ 5 500 unique signal names
File pairing	CAN-only + video-embedded
Storage	Amazon S3

2.1.2 Controlled Track Test Data

In addition to the FOT fleet recordings, a set of controlled track tests was conducted on a dedicated test facility using a development vehicle (B-1706, August 2025). These tests follow Euro NCAP evaluation protocols for AEB systems and cover a range of collision scenarios involving cars, pedestrians, and cyclists at various speeds and lighting conditions (Table 3.9 in Section 3.6.2).
The track test recordings use the BLF rather than MF4. The primary purpose of these recordings is to calibrate the brake simulation model described in Section 3.6.2: the seven events from vehicle B-1706 provided GT deceleration profiles against which the External Brake Request (XBR) ramp parameters, brake lag time constants, and realisation gain were fitted (Table 3.8). The track tests also supply known-outcome AEB activations under repeatable conditions, serving as a sanity check for the pipeline's TTC computation. However, all FP analysis results presented in Chapter 4 are based exclusively on the TfL FOT fleet data.

2.2 Vehicle Platform and CAN Bus Architecture

2.2.1 CAN Bus Structure

The Volvo city bus platform transmits onboard signals via a multi-bus CAN architecture. The Driver Assistance Control Unit (DACU) logger captures traffic from four CAN channels, each decoded using a corresponding Database Container (DBC) file[1]. Table 2.2 summarises the channel assignment.

Table 2.2: CAN channel assignment and DBC decoding files.

[1]The DBC file names shown here are abbreviated for readability; the actual filenames include version and variant suffixes (e.g. `VicinityNet2-T2_1.27.0 6.dbc`).

Channel	Bus Name	Domain	DBC File
2	J1939	Heavy-duty powertrain	`J1939_1.dbc`
3	Backbone	Vehicle backbone	`Backbone_2.dbc`
4	VicinityNet 2	Perception (radar/camera)	`VicinityNet2.dbc`
5	VicinityNet 1	Perception (radar/camera)	`VicinityNet1.dbc`

The J1939 bus carries standardised heavy-duty vehicle signals including wheel-based speed, brake pressure, and powertrain status. The Backbone bus provides vehicle-level coordination signals such as steering wheel angle, accelerator pedal position, and driver vigilance indicators. The two VicinityNet buses carry perception-layer outputs from the forward-looking radar (FLR) and the forward-looking camera (FLC), including per-object tracking data (position, velocity, classification) and system-level status flags such as `CM_Status`.

2.2.2 Signal Taxonomy

From the approximately 5 500 raw CAN signals available in each MF4 file, the current pipeline no longer operates on a flat manually curated list of single channels. Instead, it consumes a compact set of ego/control and driver-response signals together with two structured perception families: 24 forward-looking radar slots and 24 forward-looking camera slots, each decoded with a fixed set of per-object attributes. The signal taxonomy is therefore better described by functional groups than by a legacy candidate count. Table 2.3 summarises the functional signal groups used by the pipeline, while the exact signal inventory and family-level field definitions are provided in Appendix A.
The perception input is organised as two slot-indexed object families,`FLRObj00`–`FLRObj23` and `FLCObject00`–`FLCObject23`. For each slot, the pipeline consumes object-level kinematic states together with confidence, classification, and continuity-related fields. These radar and camera object tracks are then processed through gating, matching, and ranking stages to construct the target assignment used in TTC evaluation. The exact signal inventory and family-level field definitions are provided in Appendix A.

2.3 Problem Definition

This section formalises the false positive detection problem addressed in this thesis. It is organised as follows. Section 2.3.1 defines the AEB system behaviour and what constitutes an activation event. Section 2.3.2 establishes the qualification criteria that filter out trivial activations. Section 2.3.3 presents the two-condition classification logic that labels each qualified event as FP or TP. Section 2.3.5 introduces the TTC formulations used in the classification at a conceptual level; the full estimation methodology, including state-space modelling, adaptive filtering, and smoothing—is deferred to Chapter 3. Finally, Section 2.3.6 defines the target-presence categories used for root cause analysis.

Table 2.3: Current signal groups and their roles in the detection pipeline.

Category	Signal Group	Role
Ego-motion and event state	`CM_Status`, speed channels, longitudinal-acceleration channels, `BrakeSwitch`	Event segmentation, qualified-brake-event screening, ego-speed thresholding, and temporal anchoring of AEB activation
Driver-response and system context	`BrakePedalPosition`, accelerator-pedal channel, `SteeringWheelAngle1`, `EBSBrakeSwitch`, and system-availability flags	Driver-response assessment, brake-authenticity cross-checks, and scenario-bucket attribution diagnostics
Absolute timebase	`YearUTC`, `MonthUTC`, `DayUTC`, `HoursUTC`, `MinutesUTC`, `SecondsUTC`	Reconstruction of absolute CAN UTC timestamps for cross-file time alignment and auditability
FLR object family	`FLRObj`*XX*`_`{kinematics, confidence, variance, shape, slot continuity, class}	Radar hard/soft gating, FLR–FLC association, target ranking, and the longitudinal component of TTC estimation
FLC object family	`FLCObject`*XX*`_`{kinematics, confidence, variance, corner geometry, slot continuity, class}	Camera-side association geometry, target identity continuity, lateral state estimation, and width reconstruction for fused targets

2.3.1 AEB System Behaviour

The Collision Mitigation (CM) system monitors the forward scene using radar and camera sensors. When an imminent collision is detected, the system issues an autonomous braking command. The system state is reported via the `CM_Status` signal, where values 2 and 3 indicate active AEB intervention (partial and full braking, respectively), as defined in internal files.

An AEB activation event is defined as a contiguous time interval during which `CM_Status` $\in \{2, 3\}$. Consecutive active samples separated by gaps shorter than a configurable merge threshold are merged into a single event. The timestamp of the first sample where `CM_Status` transitions to an active state is referred to as the AEB anchor point, which serves as the temporal reference for all subsequent analyses.

Figure 2.1 illustrates the sensor layout and AEB activation concept. The FLR and FLC are mounted at the front of the bus and project overlapping fields of view to detect potential collision targets. When a target is confirmed within the critical zone and TTC drops below the intervention threshold, the CM system triggers autonomous braking.

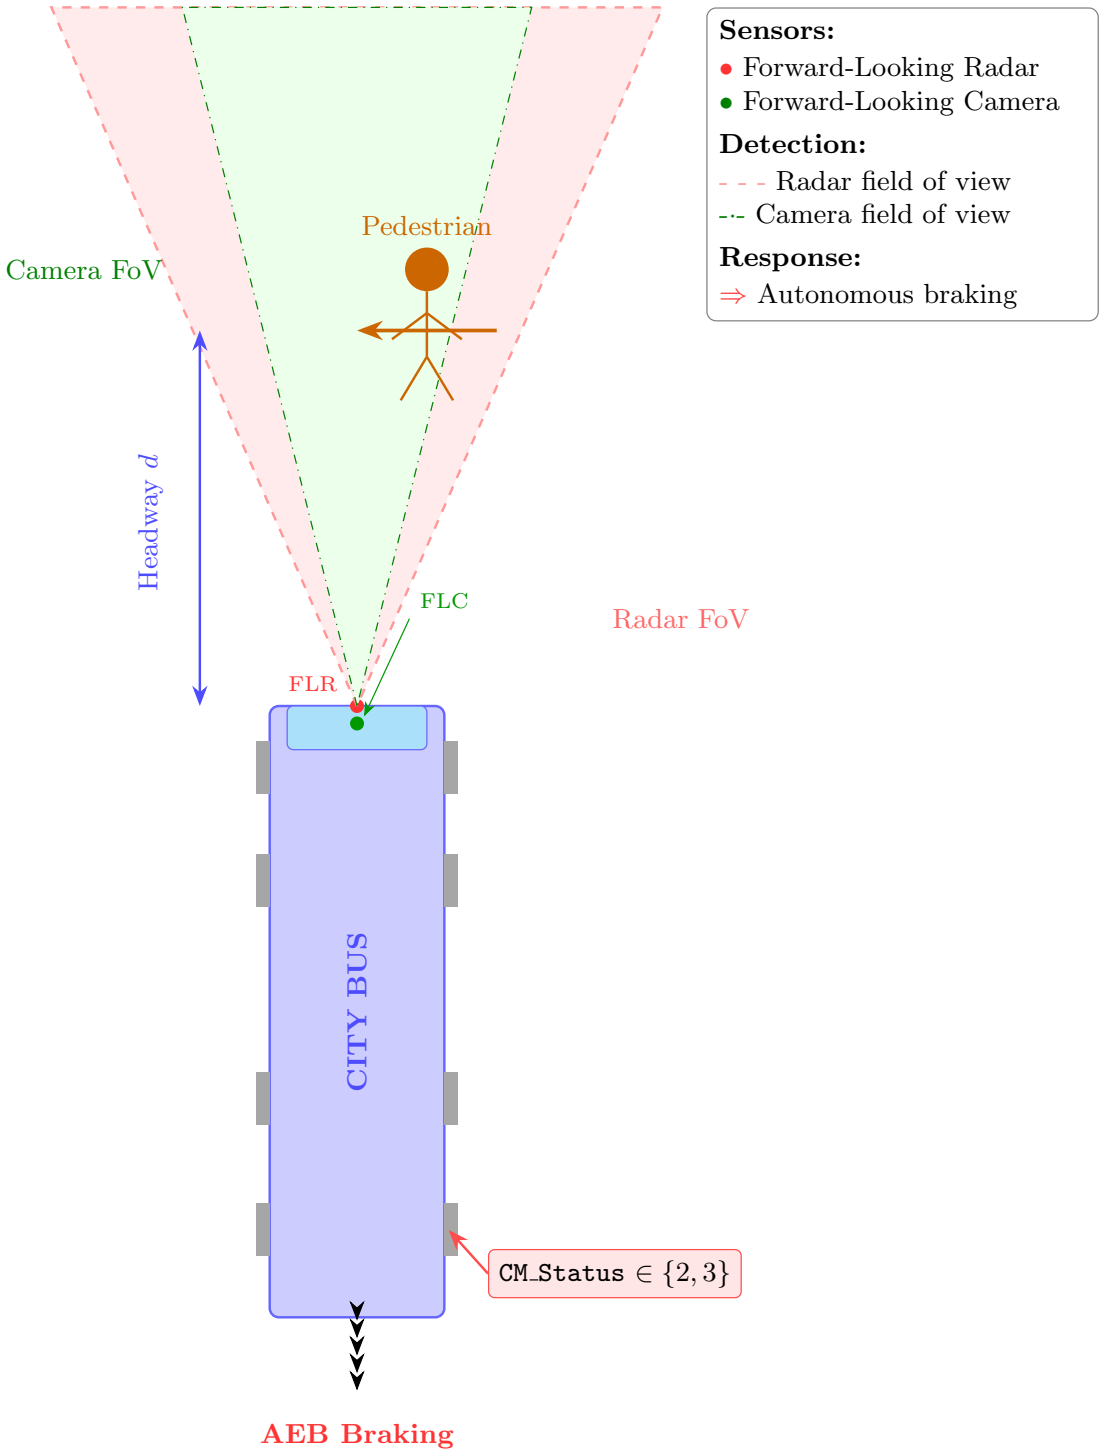

Figure 2.1: Top-down schematic of the bus AEB sensor layout and activation concept. The FLR and FLC provide overlapping detection coverage. When a collision threat is confirmed, the CM system issues an autonomous braking command (`CM_Status` $\in \{2, 3\}$).

2.3.2 Qualified Brake Event Criteria

Not every AEB activation constitutes a safety-relevant event. To filter out trivial or spurious activations, the pipeline applies a set of qualified brake event (QBE)

criteria derived from regulatory references. An AEB activation is promoted to a QBE only if all of the following conditions are satisfied simultaneously:

1. **AEB system active:** `CM_Status` $\in \{2, 3\}$
2. **Deceleration threshold:** $a_{\text{long}} \leq -1.5\ \text{m/s}^2$
3. **Speed range:** Speed gate: $v > 10\ \text{km/h}$

The deceleration criterion currently uses the raw `LongitudinalAcceleration` signal (with `EgoMotionXYplane_LongAcc` as a fallback). A simulated brake response module, which derives the commanded braking demand from `CM_Status` activation timing, has been developed as an alternative but is not yet integrated as the primary input. The raw acceleration approach is conservative: it reflects the physical vehicle response rather than the commanded demand and is therefore affected by road surface, vehicle mass, and brake system dynamics.

2.3.3 False Positive Classification

A QBE is classified as FP if either of the following conditions holds:

Condition A — Excessive TTC: The TTC at the AEB anchor exceeds a safety threshold, indicating that no emergency intervention was necessary:

$$\text{TTC} > \text{TTC}_{\text{threshold}} \tag{2.1}$$

where the threshold is defined as:

$$\text{TTC}_{\text{threshold}} = \max\left(1.4\ \text{s},\ \frac{v_{\text{ego}}}{2 \cdot a_{\text{LPoB}}}\right) \tag{2.2}$$

where LPoB denotes the Last Point of Braking and $a_{\text{LPoB}} = 3.0\ \text{m/s}^2$ is the assumed braking deceleration used in the adopted TfL threshold definition. The speed-dependent term $v_{\text{ego}}/(2\,a_{\text{LPoB}})$ provides a kinematically motivated time margin derived from this assumed braking capability: for a vehicle travelling at speed v_{ego}, the corresponding full-stop time under constant deceleration a_{LPoB} is $t_{\text{stop}} = v_{\text{ego}}/a_{\text{LPoB}}$, so the adopted threshold uses half of that time as a conservative intervention boundary. In practical terms, this means that higher ego speed implies a larger acceptable TTC at intervention onset, reflecting the longer time needed for the vehicle to reduce collision risk. The 1.4 s floor prevents the threshold from becoming unrealistically small at low speeds; this lower bound is specified in TfL's internal AEB system classification document and is therefore treated here as a requirement of the adopted FP criterion rather than a parameter tuned in this thesis.

Condition B — Absent driver brake response: No driver brake input is detected within 1.2 s of the AEB activation:

$$t_{\text{driver\_brake}} > 1.2\ \text{s} \quad \text{or no brake detected} \tag{2.3}$$

In the implementation, driver braking is detected by whether `BrakeSwitch` > 0.5 or `BrakePedalPosition` exceeds a configurable threshold within the 1.2 s window following the AEB anchor. The 1.2 s threshold is specified in TfL's internal AEB

system classification specification, reflecting operational requirements for urban bus driving conditions. The absence of any brake input within this window supports the FP classification. Note that the current implementation does not assess driver attentiveness directly—it relies solely on the presence or absence of a physical brake signal.
Events satisfying neither condition are classified as TP, representing genuine collision-avoidance interventions.

2.3.4 Signal Smoothing Strategy

Radar and camera measurements of relative kinematics—particularly relative velocity and acceleration—exhibit substantial high-frequency noise and discontinuities caused by sensor slot reassignment. These artefacts can produce spurious TTC spikes that confound the FP classification logic. To obtain reliable TTC estimates, the pipeline applies a Rauch–Tung–Striebel (RTS) smoother [27]—a two-pass optimal state estimator that leverages both past and future measurements.

2.3.4.1 Kalman Filter Background

The Kalman filter [28] is a recursive state estimator that processes noisy measurements sequentially to produce an optimal estimate of a system's internal state. At each time step, the filter performs two operations:

1. **Prediction:** The current state estimate is propagated forward in time using a dynamic model (the state transition matrix $\mathbf{F}$), yielding a predicted state and an associated prediction uncertainty.
2. **Update:** A new sensor measurement is compared against the predicted observation. The discrepancy—termed the *innovation*—is then incorporated into the state estimate, weighted by the Kalman gain $\mathbf{K}$.

The Kalman gain balances trust between the model and the measurement: when measurement noise is large (high observation covariance $\mathbf{R}$), the gain is small and the filter relies more on the predicted state; when the model is uncertain (high process noise covariance $\mathbf{Q}$), the gain increases to track the measurements more closely. The result at each step is a minimum-variance estimate of the true state given all observations received so far.
Different choices of the state vector and transition matrix $\mathbf{F}$ yield filter variants suited to different physical models. Three variants are used in this thesis:

- **Random-walk model** ($F = 1$, scalar state): assumes the signal evolves slowly between samples. Used in the RTS smoother described below.
- **Constant-velocity (CV) model**: state $\mathbf{x} = [p,\, v]^\top$; assumes position changes at a constant rate. Used for inter-sensor pair smoothing in the cross-sensor target construction stage (Section 3.3.1.2).
- **Constant-acceleration (CA) model**: state $\mathbf{x} = [d,\, v_{\text{rel}},\, a_{\text{rel}}]^\top$; assumes acceleration persists between samples. Used for the adaptive TTC pre-smoothing filter (Section 3.5).

In all cases, the underlying principle is the same: the filter fuses a dynamic prediction with noisy sensor data, and the Q/R balance determines how aggressively the output is smoothed.

2.3.4.2 RTS Smoother Formulation

The RTS smoother extends the Kalman filter for offline analysis. A standard Kalman filter is causal: the estimate at time k uses only measurements up to time k. When all data are available (as in this offline pipeline), the RTS backward pass revisits each estimate in reverse chronological order, incorporating information from future measurements to produce a globally optimal smoothed trajectory.
The general formulation is given below in matrix notation so that it applies to any of the three filter variants introduced above. For the scalar random-walk smoother, all matrices reduce to scalars.

Forward pass (Kalman filter).

$$\hat{\mathbf{x}}_{k|k-1} = \mathbf{F}\,\hat{\mathbf{x}}_{k-1|k-1}, \tag{2.4}$$

$$\mathbf{P}_{k|k-1} = \mathbf{F}\,\mathbf{P}_{k-1|k-1}\,\mathbf{F}^{\top} + \mathbf{Q}, \tag{2.5}$$

$$\mathbf{K}_k = \mathbf{P}_{k|k-1}\,\mathbf{H}^{\top}\Big(\mathbf{H}\,\mathbf{P}_{k|k-1}\,\mathbf{H}^{\top} + \mathbf{R}\Big)^{-1}, \tag{2.6}$$

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k\Big(\mathbf{z}_k - \mathbf{H}\,\hat{\mathbf{x}}_{k|k-1}\Big), \tag{2.7}$$

$$\mathbf{P}_{k|k} = \Big(\mathbf{I} - \mathbf{K}_k\,\mathbf{H}\Big)\,\mathbf{P}_{k|k-1}. \tag{2.8}$$

The term $(\mathbf{z}_k - \mathbf{H}\,\hat{\mathbf{x}}_{k|k-1})$ in Equation (2.7) is the *innovation*: the discrepancy between the sensor measurement and the model's prediction. The Kalman gain $\mathbf{K}_k$ determines how much of this discrepancy is incorporated into the revised state estimate. When observation noise dominates ($\mathbf{R}$ large relative to $\mathbf{Q}$), the gain is small and the filter relies more on the dynamic model, producing stronger smoothing.

Backward pass (RTS smoothing).

$$\mathbf{C}_k = \mathbf{P}_{k|k}\,\mathbf{F}^{\top}\,\mathbf{P}_{k+1|k}^{-1}, \tag{2.9}$$

$$\hat{\mathbf{x}}_{k|N} = \hat{\mathbf{x}}_{k|k} + \mathbf{C}_k\Big(\hat{\mathbf{x}}_{k+1|N} - \hat{\mathbf{x}}_{k+1|k}\Big), \tag{2.10}$$

$$\mathbf{P}_{k|N} = \mathbf{P}_{k|k} + \mathbf{C}_k\Big(\mathbf{P}_{k+1|N} - \mathbf{P}_{k+1|k}\Big)\,\mathbf{C}_k^{\top}. \tag{2.11}$$

The result $\hat{\mathbf{x}}_{k|N}$ is the smoothed state estimate at time k given *all* N measurements. This eliminates the phase delay inherent in causal filters: the smoothed signal remains temporally aligned with the true collision geometry rather than being shifted to the right.

Scalar random-walk special case. For the signal-level RTS smoother used in this pipeline, the state is a single scalar (x_k) with $F = 1$ and $H = 1$. The matrix equations above simplify to:

$$K_k = \frac{P_{k|k-1}}{P_{k|k-1} + R}, \tag{2.12}$$

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k\,(z_k - \hat{x}_{k|k-1}), \tag{2.13}$$

$$C_k = \frac{P_{k|k}}{P_{k+1|k}}, \tag{2.14}$$

$$\hat{x}_{k|N} = \hat{x}_{k|k} + C_k\,(\hat{x}_{k+1|N} - \hat{x}_{k+1|k}). \tag{2.15}$$

Here, R/Q is the single tuning parameter: a higher ratio yields stronger smoothing because the filter trusts temporal continuity more than individual measurements.

2.3.4.3 Why RTS Over Simpler Alternatives

The choice of the RTS smoother over conventional alternatives (e.g. moving average or Butterworth low-pass filters) is motivated by two properties:

1. It is non-causal, using both past and future data to eliminate phase delay—critical for an offline FP criterion that evaluates TTC at a precise anchor timestamp.
2. It naturally handles the irregular time steps arising from CAN bus scheduling without requiring prior resampling to a fixed grid.

The ratio R/Q is the single parameter controlling smoothing aggressiveness: a higher ratio yields a smoother output that trusts temporal continuity more than individual measurements. The specific parameter values used in the pipeline and the two-stage application strategy are presented alongside the TTC estimation methodology in Section 3.5.

2.3.5 TTC Computation

TTC is the central metric in the FP classification. Two formulations are used because each captures a different aspect of collision risk. TTC (Eq. 2.16) computes a first-order collision-time estimate from the instantaneous headway and relative velocity at each timestamp; it is well-defined whenever the gap is closing, but because it does not account for relative acceleration it can overestimate collision risk when the target is already decelerating away from the ego vehicle. The enhanced metric eTTC (Eq. 2.17) extends the projection to second order by additionally incorporating the instantaneous relative acceleration, thus better capturing situations where braking or acceleration is ongoing. The limitation of this second-order projection is that the acceleration term is taken as a snapshot at the evaluation instant; in real urban traffic, target acceleration can change rapidly as vehicles brake, release, or change manoeuvre, so the quadratic extrapolation is only valid over a short horizon. eTTC should therefore be interpreted as a locally refined collision-time estimate rather than as a long-range trajectory prediction. In addition, because relative acceleration is derived from noisier signals than position or velocity, eTTC is more sensitive to measurement noise and may become undefined when the estimated relative deceleration is already sufficient to prevent collision entirely.

TTC: Assumes both ego and target maintain their current velocities:

$$\mathrm{TTC} = -\frac{d}{v_{\mathrm{rel}}}, \quad v_{\mathrm{rel}} < 0,\ d > 0 \tag{2.16}$$

where d is the headway distance (from radar `LongPos`) and v_{rel} is the relative closing speed (from radar `LongVel`).

Enhanced TTC (eTTC): Incorporates relative longitudinal acceleration:

$$d + v_{\text{rel}} \cdot t + \tfrac{1}{2}\, a_{\text{rel}} \cdot t^2 = 0 \tag{2.17}$$

The smallest positive root of this quadratic gives eTTC. When the discriminant $v_{\text{rel}}^2 - 2\, a_{\text{rel}}\, d < 0$, the relative deceleration is sufficient to prevent collision and $\text{eTTC} = \infty$. In such cases, TTC is used as a fallback.
Both TTC variants are evaluated on smoothed kinematic inputs produced by the RTS procedure introduced in Section 2.3.4; the specific smoothing parameters and the two-stage application strategy (pre-smoothing of radar inputs followed by post-smoothing of the derived TTC series) are detailed in Section 3.5.

2.3.6 Target Presence Classification

To support root cause analysis, QBEs are further categorised based on whether a valid forward target was tracked by the Adaptive Cruise Control (ACC)/AEB perception system at the AEB anchor. Three categories are defined using the longitudinal position of the primary tracked target:

Table 2.4: Target presence classification at the AEB anchor.

Category	Condition	Interpretation
PRESENT	$d_{\text{long}} > 0$	A valid forward target is tracked ahead of the ego vehicle; TTC computation is meaningful.
ABSENT	$d_{\text{long}} \leq 0$	No forward target or target behind ego; indicates a potential ghost trigger or sensor artefact.
UNKNOWN	d_{long} unavailable	Target distance data missing or invalid; inconclusive.

This three-way classification is combined with the FP/TP label and, where available, unsupervised clustering results to produce a multi-dimensional event grouping. The ABSENT category is of particular interest for FP analysis, as it corresponds to AEB activations where no physical obstacle was present in the forward path.

2.3.7 Summary of Problem Formulation

The FP detection problem addressed in this thesis can be summarised as follows: given a continuous stream of CAN bus signals from an urban bus equipped with an AEB system, the task is to:

1. Detect all AEB activation events by monitoring `CM_Status`.
2. Filter these to QBEs using deceleration and speed criteria.
3. Classify each qualified event as FP or TP using TTC-based collision risk assessment (Condition A) and driver brake response analysis (Condition B).
4. Categorise events by target presence and FP/TP status to support root cause analysis.

This formulation transforms the open-ended question of "why did the AEB brake?" into a structured, criteria-based classification that is both reproducible and traceable to regulatory references.

3
Methodology

This chapter presents the end-to-end offline analysis pipeline developed to detect, classify, and characterise candidate FP AEB braking events from raw fleet recordings. The discussion in this chapter focuses on the conceptual decision logic underlying these stages. Detailed signal-selection policies, imputation rules, calibration procedures, and tuning parameters are collected in Appendix B and are cross-referenced where relevant. The pipeline is organised as a sequence of eleven self-contained processing stages, each consuming the output of the preceding stage. Figure 3.1 provides a high-level block diagram of the complete workflow.

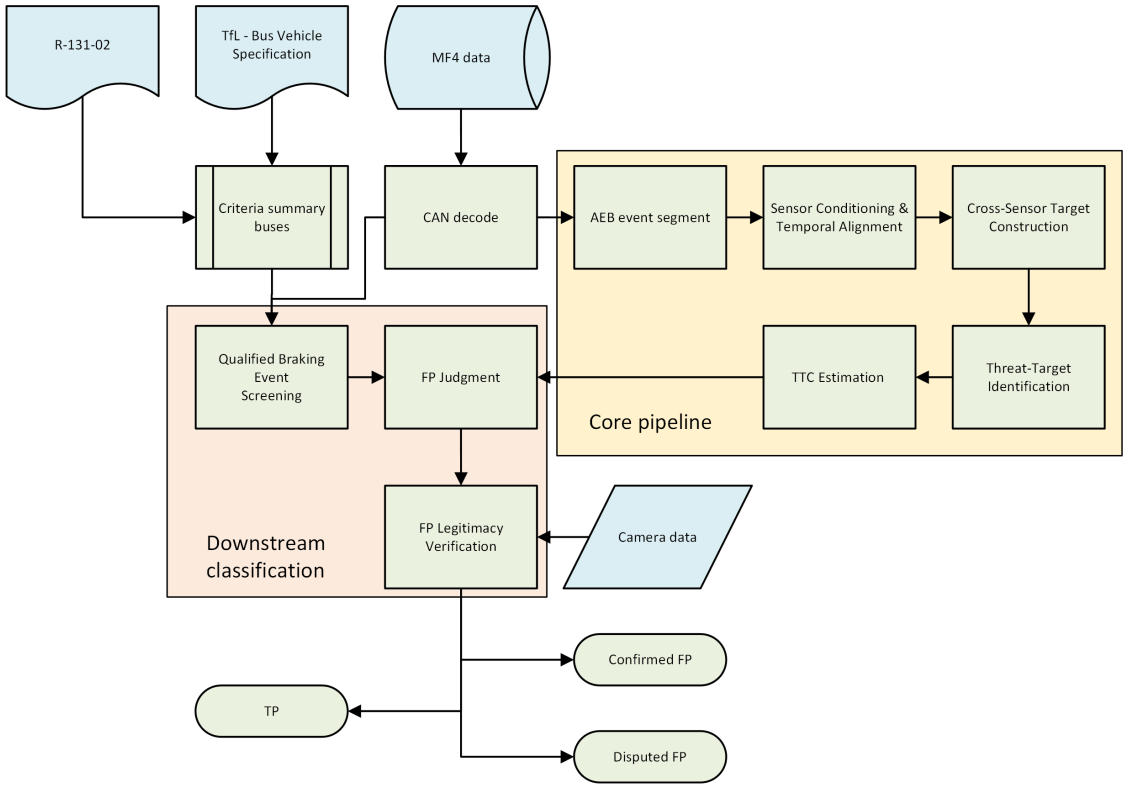

Figure 3.1: Overview of the offline analysis pipeline. Raw MF4 recordings enter at the top; the final output is a classified event catalogue with per-event FP/TP labels and supporting signal traces.

The stages are summarised in Table 3.1, described in the following sections.

Table 3.1: Core pipeline stages and Downstream classification stages

Stage	Name	Purpose
C1	Data Ingest & Event Windowing	Decode MF4/DBC recordings into ego and perception time series, reconstruct the UTC-referenced analysis timeline, and identify candidate AEB-related event windows from `CM_Status` transitions.
C2	Sensor Conditioning & Temporal Alignment	Apply radar/camera frame-level validity control and align camera observations onto the radar-centred processing axis so that downstream cross-sensor target construction is performed on a consistent event-local timeline.
C3	Cross-Sensor Target Construction	Establish radar–camera target correspondences through gated association, conflict resolution, and post-association refinement, producing a unified candidate target pool for each event.
C4	Threat-Target Identification & Timeline Construction	Rank competing target candidates, maintain target identity over time, and construct the ego–target timeline used for event-level risk evaluation.
C5	TTC Estimation	Compute TTC/eTTC from the selected ego–target kinematic timeline and generate the collision-risk quantities used by the downstream false-positive analysis logic.
D1	Qualified Braking Event Screening	Segment `CM_Status` activations into discrete events, simulate the AEB brake request as a piecewise-linear XBR ramp filtered through first-order actuator dynamics, and apply four qualification criteria (deceleration $\leq -1.5\,\mathrm{m/s^2}$, `CM_Status`$\in \{2,3\}$, $v > 10\,\mathrm{km/h}$) to select events eligible for FP evaluation.
D2	FP Judgment	Classify each qualified event as FP or TP via two disjunctive conditions: Condition A (measured TTC exceeds the speed-dependent dynamic threshold $\max(1.4\,\mathrm{s},\ v_{\mathrm{ego}}/2a_{\mathrm{LPoB}})$) and Condition B (no driver brake response within 1.2 s of the AEB anchor), with FP $= A \vee B$.
D3	FP Legitimacy Verification	Extract per-event scenario-bucket features (ego kinematics, target state, TTC context, radar attributes, driver state) at the AEB anchor, partition events into four groups (G0–G3) by target presence and TTC relative to the dynamic threshold, and perform within-group distributional analysis plus cross-group consistency checks to validate FP labels.

3.1 Data Ingest and Event Windowing

3.1.1 Data Ingestion and UTC Timebase

Raw MF4 files are downloaded from the S3 object store and decoded using the `asammdf` library in conjunction with DBC signal definitions parsed by `cantools`. The decoding step produces three wide-format DataFrames:

- **Ego signals** (6 channels): longitudinal speed, longitudinal acceleration, steering wheel angle, `CM_Status`, brake switch, and accelerator pedal position.
- **FLR signals** (24 targets · 19 attributes = 456 channels): per-object longitudinal/lateral position and velocity, longitudinal acceleration, existence confidence, track status, object class, dynamic class, lifetime counter, coast index, and lateral/longitudinal extent fields.
- **FLC signals** (24 targets · 19 attributes = 456 channels): analogous camera-based detections using corner-point representations (distance and angle to far/left/right corners) and reliability confidence.

MF4 files use relative timestamps measured from the file start. To align recordings with external references such as video, the pipeline reconstructs an absolute Coordinated Universal Time (UTC) timebase from CAN-broadcast UTC fields (`YearUTC`, `MonthUTC`, ..., `SecondsUTC`) by linear interpolation between integer-second boundaries. When no valid CAN UTC is available, the MF4 file header `start_time` is used as a degraded fallback that preserves relative timing but does not provide verified absolute UTC accuracy.

3.1.2 AEB Event Segmentation

The event segmentation algorithm identifies candidate analysis windows from the ego signal stream. The procedure is:

1. Forward-fill and backward-fill the `CM_Status` column to eliminate intermediate NaN samples.
2. Identify all timestamps where `CM_Status` $\neq 1$ (`Enabled`).
3. Group consecutive non-`Enabled` timestamps into episodes; merge episodes separated by gaps $\leq \Delta t_{\text{merge}} = 1.0\,\text{s}$ into a single event.
4. For each merged event, cut an analysis window extending $\pm\, w$ seconds beyond the first and last non-`Enabled` timestamps, where $w = 5.0\,\text{s}$ (batch) or $w = 10.0\,\text{s}$ (interactive).

The output is a dictionary mapping 1-based event identifiers to independent DataFrame slices. These slices serve as the unit of analysis for the subsequent processing stages. The $\pm 5\,\text{s}$ window captures both the scenario development preceding the intervention and the immediate post-braking recovery period, while the $1.0\,\text{s}$ merge threshold prevents brief signal jitter from fragmenting one physical event into several artificial segments.

3.2 Sensor Conditioning and Temporal Alignment

3.2.1 Sensor Data Quality Control

Raw radar and camera target frames contain numerous invalid or degraded entries caused by sensor errors, untracked states, or reserved encodings. A two-level gating architecture filters these observations before target association and fusion.

3.2.1.1 Hard and Soft Gating

A target frame is declared unconditionally invalid (`slot_valid` $= 0$) and all associated measurement values are set to NaN if any of the hard gate conditions listed in Table 3.2 is met.

Table 3.2: Hard gate conditions (measurement invalidated on match).

Signal	Invalid encoding	Interpretation
`TrackStatus`	$\in \{6, 7\}$	Error / Not Available
`ObjClass`	$= 14$	Error class
`LongPos`	$\geq 200\,\mathrm{m}$	Idle/placeholder value

Frames that pass the hard gate but still exhibit degraded tracking quality receive a multiplicative confidence penalty:

$$\texttt{slot\_confidence} = \prod_{i \in \text{triggered soft gates}} p_i, \qquad p_i = 0.5. \tag{3.1}$$

Table 3.3: Soft gate conditions (confidence reduced by a factor of 0.5 per trigger).

Signal	Degraded encoding
`TrackStatus`	$\in \{2, 3, 4, 5\}$ (Predicted / Coasting / etc.)
`ObjDynClass`	$= 3$ (Unknown)
`ObjClass`	$\in \{12, 13, 15\}$ (Reserved)

A fully healthy frame has `slot_confidence` $= 1.0$; one soft-gate trigger reduces it to 0.5; two triggers to 0.25. This confidence score is propagated to the downstream association step, where it penalises uncertain pairings in the matching cost function.

3.2.1.2 Slot Switch Detection and Fill Strategy

Within a single radar object slot (e.g. `FLRObj05`), the physical target may change when the sensor reassigns the slot to a new detection. A set of kinematic and metadata discontinuity criteria identifies such reassignment events between consecutive frames. At each detected switch point, forward-fill propagation is interrupted to prevent the kinematic state of the previous physical target from leaking into the new one.

After gating, remaining missing values are filled using a signal-type-aware strategy: continuous kinematic signals receive segment-bounded forward fill (interrupted at slot switches), lateral position uses a sentinel value beyond the lane width to trigger downstream ranking demotion, and categorical signals are set to out-of-range sentinels. The complete slot-switch detection criteria and per-signal fill rules are given in Appendix B.5 (Tables B.4 and B.5).

3.2.2 Multi-Sensor Temporal Alignment

The FLR and FLC broadcast target data at approximately 10 Hz each, but their clocks are not perfectly synchronised, resulting in inter-sensor time offsets of $\pm$ 10–50 ms. To enable subsequent spatial association, camera timestamps must be aligned onto the radar time axis.

3.2.2.1 Alignment Procedure

The alignment proceeds in two rounds:

1. **First-pass matching**: a tolerance-free nearest-neighbour join (`merge_asof`, direction = nearest) between FLR and FLC timestamps yields a raw residual distribution $\{\delta_i = t_{\mathrm{FLR},i} - t_{\mathrm{FLC},i}\}$.
2. **Systematic offset removal**: the median residual $\tilde{\delta} = \mathrm{median}(\delta_i)$ is subtracted from all FLC timestamps to correct for the bulk clock bias.
3. **Outlier exclusion**: residuals exceeding $|\delta_i| > 0.2\,\mathrm{s}$ (indicative of data gaps rather than clock drift) are excluded from tolerance estimation.
4. **Adaptive tolerance**: the effective matching tolerance is set as

$$\tau_{\mathrm{eff}} = \min\Big(\tau_{\mathrm{input}},\ P_{95}(\delta),\ 0.15\,\mathrm{s}\Big). \tag{3.2}$$

5. **Final join**: a second `merge_asof` with tolerance τ_{eff} produces the aligned DataFrame `df_flr_flc_ts`, indexed on FLR timestamps and containing all FLR and FLC signal columns.

The output retains a diagnostic column `flc_attach_dt_abs` recording the absolute time difference for each matched pair, enabling downstream quality monitoring.
To complement the procedural description above, Figure 3.2 provides an event-level view of the residual temporal mismatch after the camera stream has been attached to the radar-centred processing axis. Rather than introducing another transformation step, the figure serves to verify that the remaining FLR–FLC timestamp discrepancy is sufficiently bounded for the subsequent association stage.

3.3 Cross-Sensor Target Construction

3.3.1 Radar–Camera Target Association

Multi-sensor target association establishes a per-frame one-to-one correspondence between FLR and FLC detections of the same physical object. A cascaded three-gate architecture filters candidate pairings before global optimisation.

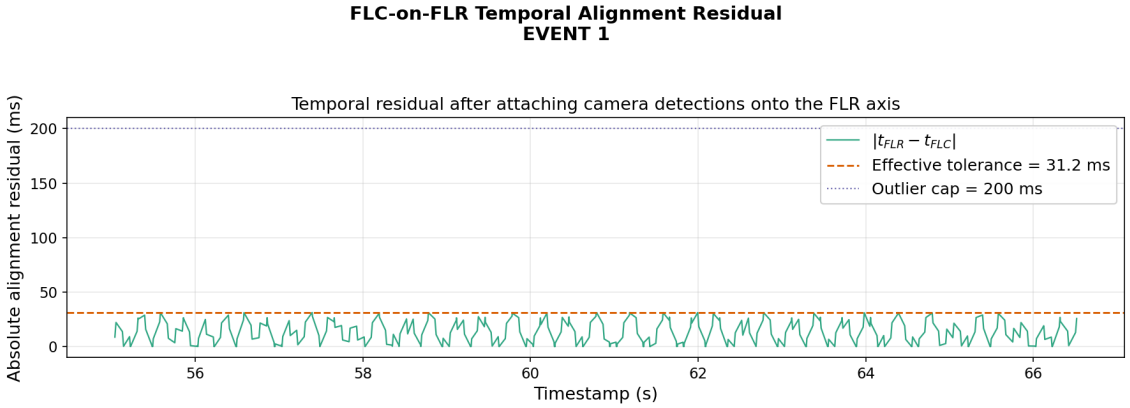

Figure 3.2: FLC-on-FLR temporal alignment residual after bulk offset removal and adaptive tolerance truncation

3.3.1.1 Association Gates

Gate 1: Mahalanobis distance. The squared Mahalanobis distance between an FLR target at position (x_r, y_r) and an FLC target at (x_c, y_c) is:

$$d_M^2 = \frac{(\Delta x)^2}{\sigma_{x,\text{FLR}}^2 + \sigma_{x,\text{FLC}}^2} + \frac{(\Delta y)^2}{\sigma_{y,\text{FLR}}^2 + \sigma_{y,\text{FLC}}^2} \leq \chi_{2,0.99}^2 = 9.21, \tag{3.3}$$

where the threshold corresponds to the 99th percentile of the χ^2 distribution with two degrees of freedom. This threshold means that only pairs that are statistically consistent with the combined radar–camera position uncertainty are admitted; in other words, it rejects geometrically implausible pairings while retaining almost all genuine matches expected under the assumed noise model. Variance bounds are clamped to $\sigma_x^2 \in [0.09, 100.0]$ m^2 and $\sigma_y^2 \in [0.04, 25.0]$ m^2 to prevent numerical instability.

Gate 2: Bounding-box overlap. Axis-aligned bounding boxes are constructed for each sensor target:

- FLR: derived from `LatExtLeft/Right` and `LongExtFront/Back`, clamped to $[0.25, 5.0]$ m laterally and $[0.25, 15.0]$ m longitudinally; small-object templates use 0.5×0.5 m.
- FLC: computed by projecting corner-point distance/angle measurements into Cartesian coordinates to form an enclosing rectangle.

Gate 2 passes when the intersection area $A_{\text{overlap}} > 0$.

Gate 3: Kinematic consistency. The relative velocity residual between paired targets must satisfy:

$$|\Delta v_x| \leq 3.0 \text{ m/s}, \qquad |\Delta v_y| \leq 2.0 \text{ m/s}. \tag{3.4}$$

3.3.1.2 Assignment Cost and Temporal Smoothing

For all (i, j) pairs that pass all three gates, a composite cost is computed:

$$c_{ij} = d^2_{M,ij} + \frac{|\Delta v_{x,ij}|}{3} + \frac{1}{1 + A_{\text{overlap},ij}} + \lambda\,(2 - c_{\text{FLR},i} - c_{\text{FLC},j}), \tag{3.5}$$

where $\lambda = 2.0$ weights the confidence penalty and $c_{\text{FLR}}, c_{\text{FLC}} \in (0, 1]$ are the slot confidence values produced by the gating step. The cost combines spatial consistency, kinematic consistency, geometric overlap, and measurement reliability in a single score, with the Mahalanobis term acting as the primary discriminator and the confidence penalty serving mainly to break ties among otherwise plausible pairings. The per-frame cost matrix $\mathbf{C} \in \mathbb{R}^{n_{\text{FLR}} \times n_{\text{FLC}}}$ is solved via the Hungarian algorithm to yield the globally optimal one-to-one assignment [29].
Raw per-frame associations exhibit jitter at gate boundaries. Two complementary channels stabilise the pairing decisions over time.

Channel A: Rule-based pair hold. A four-state automaton (UNCONFIRMED → CONFIRMED → HOLD → DROPPED) requires $n = 3$ consecutive successful frames to confirm a pair and tolerates up to $m = 5$ consecutive gate failures (with $d^2_M < 15$) before dropping.
Since pair-hold operates on the 10 Hz resampled timeline, the frame-count thresholds $n = 3$ and $m = 5$ correspond to fixed physical durations of 0.3 s and 0.5 s, respectively, independent of the underlying CAN bus message rate. This frame-counting scheme is distinct from the time-based hold budget ($T_{\text{hold}} = 1.5\,\text{s}$) used in the downstream target-selection FSM, which tracks a different entity — physical target identity persistence rather than inter-sensor pair validity.

Channel B: Kalman-filter pair smoothing. A constant-velocity Kalman filter (CVKF) tracks the inter-sensor residual vector:

$$\mathbf{x} = [\Delta x,\ \Delta y,\ \dot{\Delta x},\ \dot{\Delta y}]^\top. \tag{3.6}$$

After forward filtering and RTS backwards smoothing, the smoothed Mahalanobis distance is evaluated against fixed reference variances ($\sigma_x = 2\,\text{m}$, $\sigma_y = 1\,\text{m}$) rather than the filter covariance $\mathbf{P}$. This prevents overconfident $\mathbf{P}$ from inflating normalised distances for moderate true residuals. A pair is retained when:

$$d^2_s = \frac{\hat{\Delta x}^2}{4} + \frac{\hat{\Delta y}^2}{1} < 9.21. \tag{3.7}$$

The final pairing decision combines both channels: `pass_all` $\vee$ `kf_pair_valid`.
After matching, all valid sensor targets are classified per timestamp into three groups (Table 3.4), forming the target pool used by the fusion stage.

Table 3.4: Target pool groups produced by radar–camera association.

Group	Definition	Minimum lifetime
Fused	Matched FLR+FLC pair	1 frame (0.1 s)
Camera-only	Valid FLC, no FLR match	2 frames (0.2 s)
Radar-only	Valid FLR, no FLC match	5 frames (0.5 s)

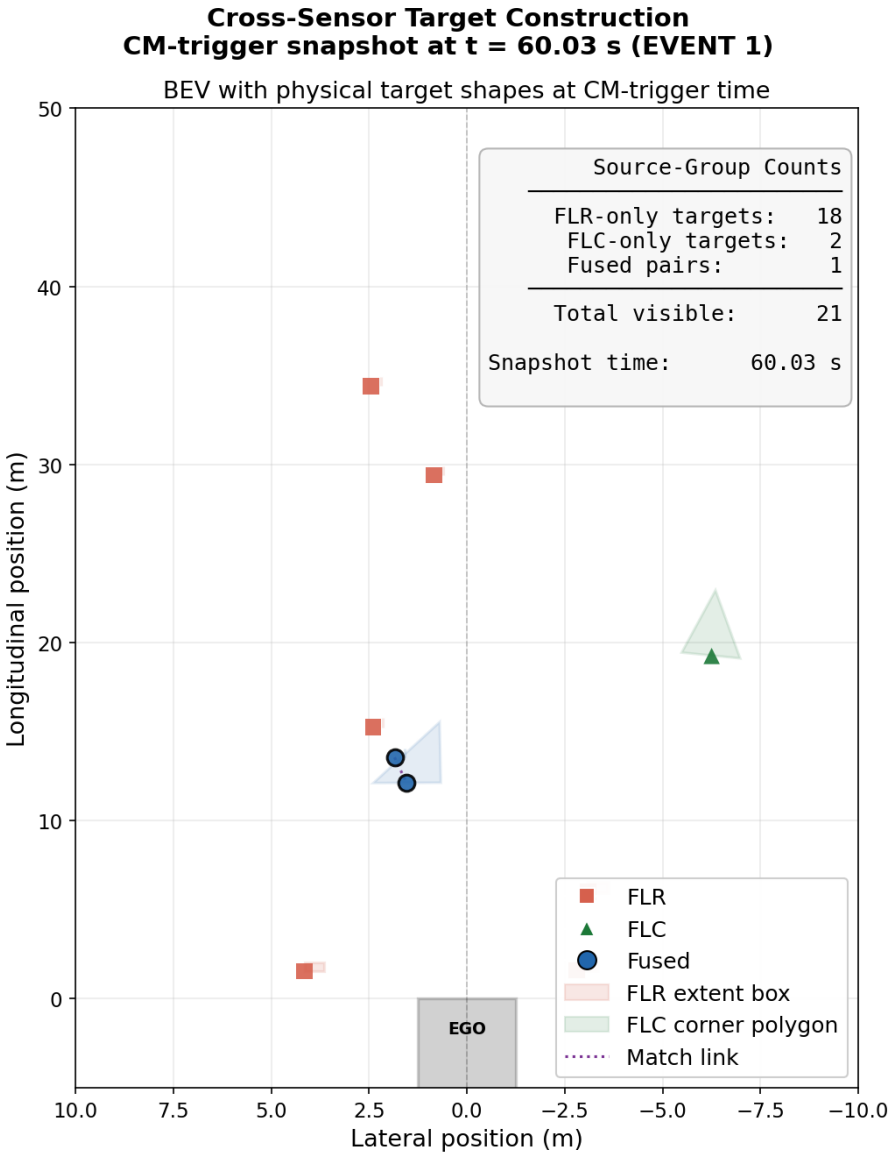

Figure 3.3: Cross-sensor target construction at the CM-trigger snapshot, showing radar-only, camera-only, and fused targets in a common BEV frame.

The outcome of the association stage is not merely a set of accepted radar–camera pairs, but a structured target pool that preserves fused and single-sensor hypotheses in parallel. Figure 3.3 visualises this representation at a representative CM-trigger snapshot, making explicit how radar-only, camera-only, and fused targets coexist before any downstream ranking or temporal smoothing is applied.

3.3.2 Post-Fusion Kalman Filtering

3.3.2.1 State-Space Model

A six-state constant-acceleration Kalman filter (CAKF) refines the fused target kinematics:

$$\mathbf{x} = [x,\ y,\ v_x,\ v_y,\ a_x,\ a_y]^\top. \tag{3.8}$$

The state transition matrix for time step Δt is:

$$\mathbf{F}(\Delta t) = \begin{bmatrix} 1 & 0 & \Delta t & 0 & \frac{\Delta t^2}{2} & 0 \\ 0 & 1 & 0 & \Delta t & 0 & \frac{\Delta t^2}{2} \\ 0 & 0 & 1 & 0 & \Delta t & 0 \\ 0 & 0 & 0 & 1 & 0 & \Delta t \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \tag{3.9}$$

A piecewise-constant jerk model yields the per-axis process noise sub-block:

$$\mathbf{Q}_{\text{blk}}(q, \Delta t) = q^2 \begin{bmatrix} \frac{\Delta t^5}{20} & \frac{\Delta t^4}{8} & \frac{\Delta t^3}{6} \\ \frac{\Delta t^4}{8} & \frac{\Delta t^3}{3} & \frac{\Delta t^2}{2} \\ \frac{\Delta t^3}{6} & \frac{\Delta t^2}{2} & \Delta t \end{bmatrix}, \tag{3.10}$$

with the full $\mathbf{Q} \in \mathbb{R}^{6\times 6}$ constructed as a block-diagonal of longitudinal and lateral sub-blocks.

3.3.2.2 Observation Model and RTS Smoothing

The observation structure reflects the complementary sensing modalities:

- **FLR update** (longitudinal): $\mathbf{H}_{\text{FLR}} \in \mathbb{R}^{3\times 6}$ observes $[x, v_x, a_x]$ (state indices 0, 2, 4).
- **FLC update** (lateral): $\mathbf{H}_{\text{FLC}} \in \mathbb{R}^{2\times 6}$ observes $[y, v_y]$ (state indices 1, 3).

Each frame applies a sequential update: FLR first (correcting the longitudinal sub-space), then FLC (correcting the lateral sub-space). When only one sensor contributes a valid measurement at a given timestamp, the other axis receives prediction only.

An optional velocity-dependent scaling of the process noise spectral density is applied:

$$q_{\text{eff}} = q \cdot \min\Big(\max\Big(1,\ \frac{|v_x|}{v_{\text{ref}}}\Big),\ \alpha_{\max}\Big), \tag{3.11}$$

where $v_{\text{ref}} = 5\,\text{m/s}$ and $\alpha_{\max} = 3$. This enlarges the process noise envelope for fast-moving targets, accommodating potential manoeuvres that a constant-q model would under-represent.

After the forward Kalman filter pass, the Rauch–Tung–Striebel backward recursion [27] refines the estimates using future observations:

$$\mathbf{G}_k = \mathbf{P}_{k|k}\,\mathbf{F}^\top\,\mathbf{P}_{k+1|k}^{-1}, \tag{3.12}$$

$$\hat{\mathbf{x}}_{k|N} = \hat{\mathbf{x}}_{k|k} + \mathbf{G}_k\Big(\hat{\mathbf{x}}_{k+1|N} - \hat{\mathbf{x}}_{k+1|k}\Big), \tag{3.13}$$

$$\mathbf{P}_{k|N} = \mathbf{P}_{k|k} + \mathbf{G}_k\Big(\mathbf{P}_{k+1|N} - \mathbf{P}_{k+1|k}\Big)\mathbf{G}_k^\top. \tag{3.14}$$

The smoothed acceleration estimates exhibit significantly reduced noise compared to the forward-only filter, which is critical for reliable eTTC computation.

3.4 Threat-Target Identification and Timeline Construction

3.4.1 Object Threat Ranking

The fused target pool may contain tens of simultaneously tracked objects per frame. A lexicographic multi-criterion scoring system identifies the most collision-relevant candidate at each timestamp.

3.4.1.1 Admission and Priority Scoring

A candidate target must satisfy:

- `slot_valid` $= 1$ (passed the hard gate),
- `LongPos_used` > 0 (target must be longitudinally ahead of the ego vehicle, eliminating rearward sensor artefacts).

Non-admitted targets receive a sentinel priority score of 99 999 and are excluded from winner selection.
Each admitted candidate is assigned a composite integer score:

$$P = \sum_{i=1}^{7} r_i \cdot w_i, \tag{3.15}$$

where a *lower* score indicates higher collision relevance. The exponentially separated weights enforce a strict lexicographic ordering, so that a target with higher lane relevance cannot be outranked by another target merely because it has better values in lower-priority dimensions such as confidence or source type. Table 3.5 summarises the seven scoring dimensions, ordered by weight.

Table 3.5: Lexicographic ranking dimensions (lower rank = higher threat).

w_i	Dimension	Levels	Criterion
10^6	Lane relevance	0–2	In-path / TLC-encroaching / Out-of-path
10^5	Closing rate	0–2	$v_{\text{rel}} < -0.5$ / < 0 / $\geq 0\,\text{m/s}$
10^4	Range bin	0–4	0–10 / 10–25 / 25–50 / 50–100 / $>100\,\text{m}$
10^3	Track health	0–2	Tracked / Predicted / Other
10^2	Maturity	0–1	Lifetime $\geq$ threshold / below
10^1	Obstacle conf.	0–3	Confidence quartile
10^0	Source type	0–2	Fused (0) / Radar (1) / Camera (2)

3.4.1.2 Lane Relevance and Time-to-Lane-Crossing

The effective lateral clearance is computed as:

$$d_{\text{lat,eff}} = \max\Big(|\texttt{LatPos}| - w_{\text{target}}/2,\ 0\Big). \tag{3.16}$$

A target is classified as in-path (rank 0) when $d_{\text{lat,eff}} \leq w_{\text{lane}}/2 = 1.75\,\text{m}$. Otherwise, the time-to-lane-crossing (TLC) is evaluated:

$$\text{TLC} = \frac{\max(d_{\text{lat,eff}} - w_{\text{lane}}/2,\ 0)}{v_{\text{toward}}}, \tag{3.17}$$

where $v_{\text{toward}} = -\text{sign}(\texttt{LatPos}) \cdot \texttt{LatVel}$. A target with TLC $\leq 2.0\,\text{s}$ is marked as encroaching (rank 1); all others receive rank 2.
Prioritises targets approaching the ego vehicle. This dimension outweighs range because a receding near-field target poses no collision risk, whereas a distant but rapidly closing target does.
Acts only as a tiebreaker at the lowest weight, which reflects the generally superior kinematic accuracy of fused targets over single-sensor detections.

3.4.2 Target State Machine

The ranking output selects the instantaneously best candidate per frame, but this selection may oscillate between targets due to measurement noise. A five-state finite-state machine (FSM) provides temporal persistence and hysteresis.

3.4.2.1 State Definitions

The FSM distinguishes five operational states, corresponding to target acquisition, confirmation, stable tracking, temporary loss, and challenge-based switching.
Search No target is tracked; awaiting a viable candidate.
Candidate A prospective target has appeared but is not yet confirmed.
Locked The system is committed to a confirmed target for TTC evaluation.
Hold The locked target is temporarily lost; the system retains its identity.
Override Pending A challenger target satisfies the override condition and awaits confirmation before switching.

3.4.2.2 Transition Logic

The hold state monitors all admitted candidates, not only the frame-best, for the locked track key. Recovery to LOCKED occurs as soon as the physical track reappears in any priority position, provided no qualifying challenger has accumulated during the absence.
A challenger target may replace the currently locked target when it demonstrates substantially greater collision urgency:

$$(x_{\text{locked}} - x_{\text{challenger}}) \geq 5\ \text{m} \quad \text{or} \quad (\text{TTC}_{\text{locked}} - \text{TTC}_{\text{challenger}}) \geq 0.7\ \text{s}. \tag{3.18}$$

A canonical track key of the form `PAIR:FLRObj`*nn*`|FLCObj`*mm* is used to identify physical targets across their fused, radar-only, and camera-only manifestations. To prevent slot reassignment on the radar side from being interpreted as continuity of

Table 3.6: FSM state transitions.

From	To	Condition
SEARCH	CANDIDATE	Best-ranked candidate appears
CANDIDATE	LOCKED	Candidate sustained for confirmation window
LOCKED	HOLD	Locked target absent for $\leq T_{\text{hold}} = 1.5\,\text{s}$
HOLD	SEARCH	Absence exceeds T_{hold} (time-based timeout)
LOCKED	OVERRIDE PENDING	Challenger satisfies override criterion
OVERRIDE PENDING	LOCKED(new)	Challenger confirmed

the same physical target, the FSM additionally applies a continuity guard before re-accepting a returning track. The primary criterion is preservation of the upstream FLR segment identifier; when this metadata is unavailable, the guard falls back to raw LifeTime reset evidence and, only as a final symptom-level fallback, a large longitudinal position jump. If any of these checks indicates an identity discontinuity, the state machine resets to SEARCH and requires the new slot occupant to pass the full confirmation path before being locked again. This prevents spurious state transitions when a target temporarily loses one sensor modality. Camera-only detections are likewise mapped to the dominant radar–camera pair key through the event-level partner map, preventing identity fragmentation when a target temporarily loses radar visibility.

Input Rate Normalisation Raw CAN bus messages arrive at ∼1.8 kHz due to multiplexed J1939 frames, far exceeding the actual sensor update cycle of ∼10 Hz. Sub-millisecond inter-frame gaps would trigger spurious hold timeouts under any frame-count logic. The FSM therefore resamples the candidate pool to a uniform 10 Hz grid (100 ms time bins, retaining the latest observation per bin per object) prior to state-machine processing. This ensures that the time-based parameters (T_{hold}, T_{confirm}) operate on physically meaningful intervals.

3.4.3 Ego–Target Timeline Merge

The FSM-selected target timeline (at ∼10 Hz sensor rate) is aligned onto the ego signal timeline via nearest-neighbour temporal matching. The merged DataFrame carries:

- Target kinematics: `LongPos_used`, `LongVel_used`, `LongAcc_used`, `LatPos_used`, `LatVel_used`.
- Metadata: `candidate_id`, `source_group`, `flr_obj`, `flc_obj`.
- Quality flag: `merged_is_ffill` (true when no valid target exists at that timestamp and the previous measurement is forward-filled).

When the state machine produces no locked target for an event (e.g. because the

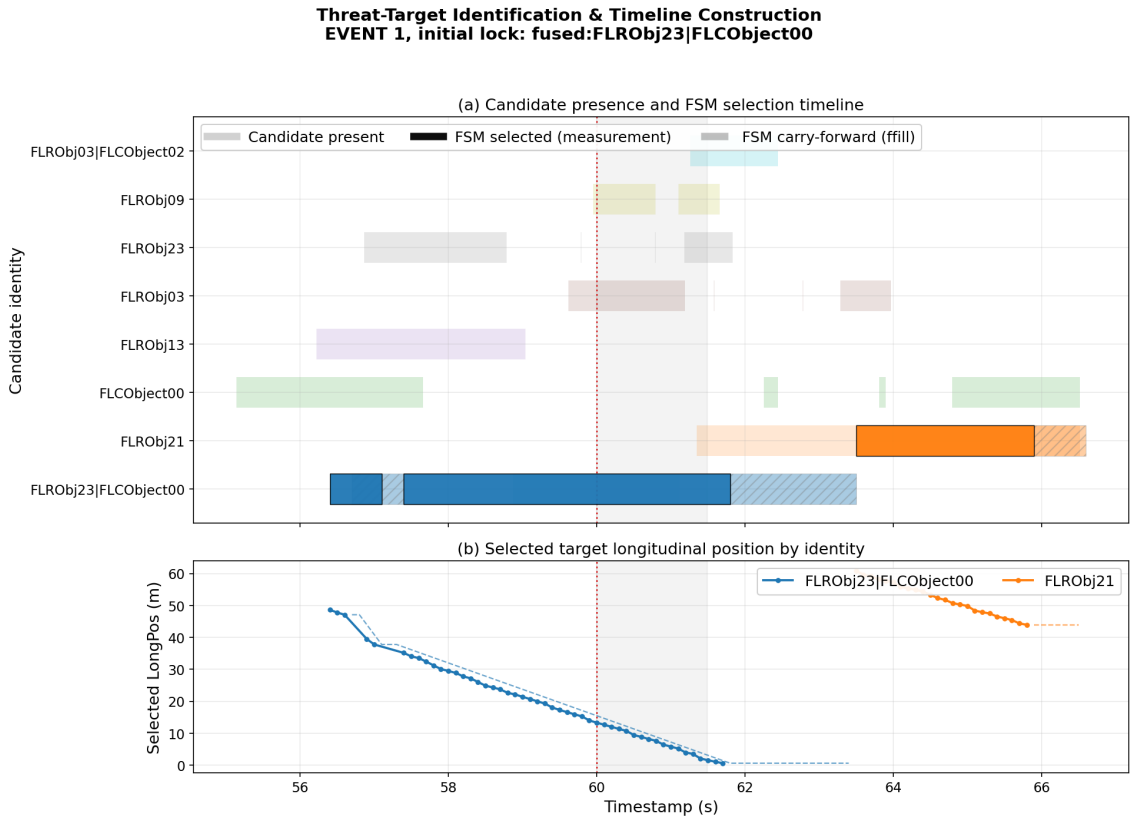

Figure 3.4: Threat-target identification and timeline construction, showing candidate identity persistence, FSM-selected intervals, and the resulting selected longitudinal trace.

target pool remains empty), the ego event timeline is retained with all target columns set to NaN, preserving the event for diagnostic review while explicitly indicating an unresolved target state. Accordingly, forward-fill is not applied unconditionally during HOLD. It is permitted only while the continuity guard indicates that the locked target remains on the same physical track; once a segment change, `LifeTime` reset, or equivalent discontinuity is detected, the carry-forward buffer is cleared and the merged timeline reverts to an unresolved target state until a new target has been confirmed.

Taken together, the ranking logic, finite-state hysteresis, and ego-aligned merge define not only which candidate is preferred at an isolated frame, but also how a persistent threat-target identity is maintained over the duration of an event. Figure 3.4 summarises this combined effect by showing both the temporal presence of competing candidate identities and the longitudinal trace ultimately carried forward for downstream risk evaluation.

3.5 TTC Estimation

3.5.1 RTS Smoothing Application

The scalar RTS smoother introduced in Section 2.3.4 is applied at two stages in the TTC estimation pipeline:

1. **Pre-smoothing of radar inputs** — applied independently to each object's headway d, relative speed v_{rel}, and relative acceleration a_{rel} before TTC computation, suppressing measurement noise at the source.
2. **Post-smoothing of computed TTC/eTTC** — applied to the resulting collision-time series to remove residual discontinuities caused by sensor slot swaps, gate transitions, and brief target occlusions.

Table 3.7 lists the tuning parameters used at each stage.

Table 3.7: RTS smoother parameters for TTC-related signals.

Signal	R **(obs. cov.)**	Q **(trans. cov.)**	R/Q
Pre-smoothing (per-object radar inputs)			
Headway d	30.0	0.05	600
Rel. speed v	30.0	0.05	600
Rel. accel. a	50.0	0.02	2500
Post-smoothing (computed TTC series)			
TTC / eTTC	80.0	0.05	1600

The same R and Q values are assigned to headway and relative speed because both signals are comparatively stable and directly measured by radar. A higher R/Q ratio (2500) is used for relative acceleration because this quantity is derived from differentiation of noisy velocity measurements, amplifying high-frequency noise. The TTC/eTTC post-smoothing ratio (1600) is chosen between these extremes to suppress residual spikes without removing the overall decay trend of the collision-time curve near the AEB anchor. These values were selected empirically to balance noise suppression against temporal fidelity, informed by visual inspection of representative events across the fleet.

The adaptive CAKF described in the following subsection operates after this scalar RTS pre-smoothing, providing a second layer of kinematic refinement with a richer state model.

3.5.2 Kalman-Based Pre-Smoothing

Raw target kinematics inherited from the merged ego–target timeline retain residual measurement noise that can produce spurious TTC spikes. A coupled CAKF is applied to the longitudinal state vector before TTC evaluation.

The three-state constant-acceleration model tracks the relative longitudinal kinematics between ego and the locked target:

$$\mathbf{x} = [d,\ v_{\text{rel}},\ a_{\text{rel}}]^{\top}, \tag{3.19}$$

with state transition:

$$\mathbf{F}(\Delta t) = \begin{bmatrix} 1 & \Delta t & \frac{\Delta t^2}{2} \\ 0 & 1 & \Delta t \\ 0 & 0 & 1 \end{bmatrix}. \tag{3.20}$$

The full observation vector $\mathbf{z} = [d, v_{\text{rel}}, a_{\text{rel}}]$ is used when all three measurements are available; otherwise a position–velocity sub-model ($\mathbf{H} = [\mathbf{I}_2 \mid \mathbf{0}]$) is employed.

3.5.3 Adaptive Process Noise and Gap Handling

The filter monitors model–measurement consistency via the Normalised Innovation Squared (NIS):

$$\text{NIS}_k = \mathbf{y}_k^\top \mathbf{S}_k^{-1} \mathbf{y}_k, \tag{3.21}$$

where $\mathbf{y}_k = \mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_{k|k-1}$ is the innovation and $\mathbf{S}_k = \mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^\top + \mathbf{R}$ is the innovation covariance. An exponentially weighted moving average tracks the NIS history:

$$\overline{\text{NIS}}_k = \alpha\, \overline{\text{NIS}}_{k-1} + (1-\alpha)\, \text{NIS}_k, \qquad \alpha = 0.98. \tag{3.22}$$

The adaptive scaling factor is then:

$$q_{\text{adapt}} = \text{clip}\left(\frac{\overline{\text{NIS}}_k}{n_z},\ \frac{1}{c},\ c \right), \qquad c = 10, \tag{3.23}$$

where n_z is the observation dimension and c is the ceiling parameter. When the innovations are consistent with the model ($\overline{\text{NIS}} \approx n_z$), the factor remains near unity; persistent model mismatch drives q_{adapt} upward, widening the process noise to track rapid kinematic changes.
An additional multiplicative term ensures responsiveness during braking manoeuvres:

$$q_{\text{eff}} = q_{\text{base}} \cdot (1 + \beta\, |\hat{a}_{\text{rel}}|) \cdot q_{\text{adapt}}, \qquad \beta = 0.5. \tag{3.24}$$

This guarantees that the filter widens its uncertainty envelope proportionally to the estimated deceleration magnitude, preventing lag during the critical AEB braking phase.
Time gaps exceeding 2.0 s (e.g. caused by target occlusion or forward-fill interruption) trigger an independent filter re-initialisation to avoid bridging across physically disconnected observation segments. In addition to temporal gaps, a longitudinal position discontinuity exceeding 10 m between consecutive frames also triggers segment re-initialisation. This accommodates radar slot re-use scenarios in which the same hardware object slot is reassigned to a physically distinct target after collision passage. Segments shorter than three frames fall back to unsmoothed raw values.
After forward filtering, an RTS backward recursion (Equations 3.12–3.14) refines the smoothed estimates $[\hat{d}, \hat{v}_{\text{rel}}, \hat{a}_{\text{rel}}]$ using both past and future observations, producing the final kinematic inputs to the TTC formulae below.

3.5.4 Collision-Time Measures

The smoothed relative kinematics $[\hat{d}, \hat{v}_{\text{rel}}, \hat{a}_{\text{rel}}]$ are converted into the TTC and eTTC measures defined in Section 2.3.5. In the present pipeline, both quantities are evaluated on the fixed locked-target ego–target timeline described below, using the smoothed longitudinal states produced by the adaptive CAKF–RTS stage. TTC is retained as a conservative baseline and grouping variable. When no valid locked target is available, the TTC outputs remain undefined (NaN) rather than being interpreted as infinite collision time.

3.5.5 Latency Compensation

The computed TTC curve exhibits a systematic temporal lag relative to the true physical collision geometry. Three sources contribute:

1. **CAN transmission delay** (∼40–70 ms): transport latency from the sensor Electronic Control Unit (ECU) through the CAN bus to the data logger.
2. **Kalman filter group delay** (∼50–120 ms): The forward-pass filter inherently lags the rapid kinematic changes; the RTS backward pass partially recovers this for offline analysis, but the residual phase shift persists for step-like transitions.
3. **Sensor processing latency** (∼20–60 ms): Internal radar cycle time and camera frame-to-output delay.

A configurable temporal offset Δt_{shift} is applied to the final TTC timeline:

$$t_{\text{corrected}} = t_{\text{computed}} - \Delta t_{\text{shift}}, \tag{3.25}$$

shifting the TTC curve earlier by Δt_{shift} seconds. The offset Δt_{shift} is determined by calibration against an independent ground-truth TTC reference provided by the GTT team; video-based review, when available, serves only as a secondary cross-check rather than the primary calibration source. Engineering analysis indicates a total system latency of 110–250 ms, which defines the expected calibration range.

3.5.6 CM-Active Target Locking

For each event, the AEB-locked target is identified as follows:

1. Define the CM-active window: all frames where `CM_Status` $\in \{2, 3, 4\}$.
2. Identify the lock anchor as the first frame with `CM_Status` $= 2$; if absent, use the first active frame.
3. Query the FSM state timeline for the nearest measurement frame—defined as `merged_is_ffill` $=$ `False` and `sm_state` $\in \{$`Candidate`, `Locked`$\}$—to the lock anchor. The `track_key` of that frame identifies the locked target.
4. If the FSM remains in SEARCH at the lock anchor (no eligible measurement frame available), the system falls back to the instantaneous priority ranking.

The single-target TTC series is then computed by filtering the candidate pool on `track_key` rather than `obj`, ensuring that all sensor manifestations (fused, radar-only, camera-only) of the same physical target contribute to a continuous kinematic history.

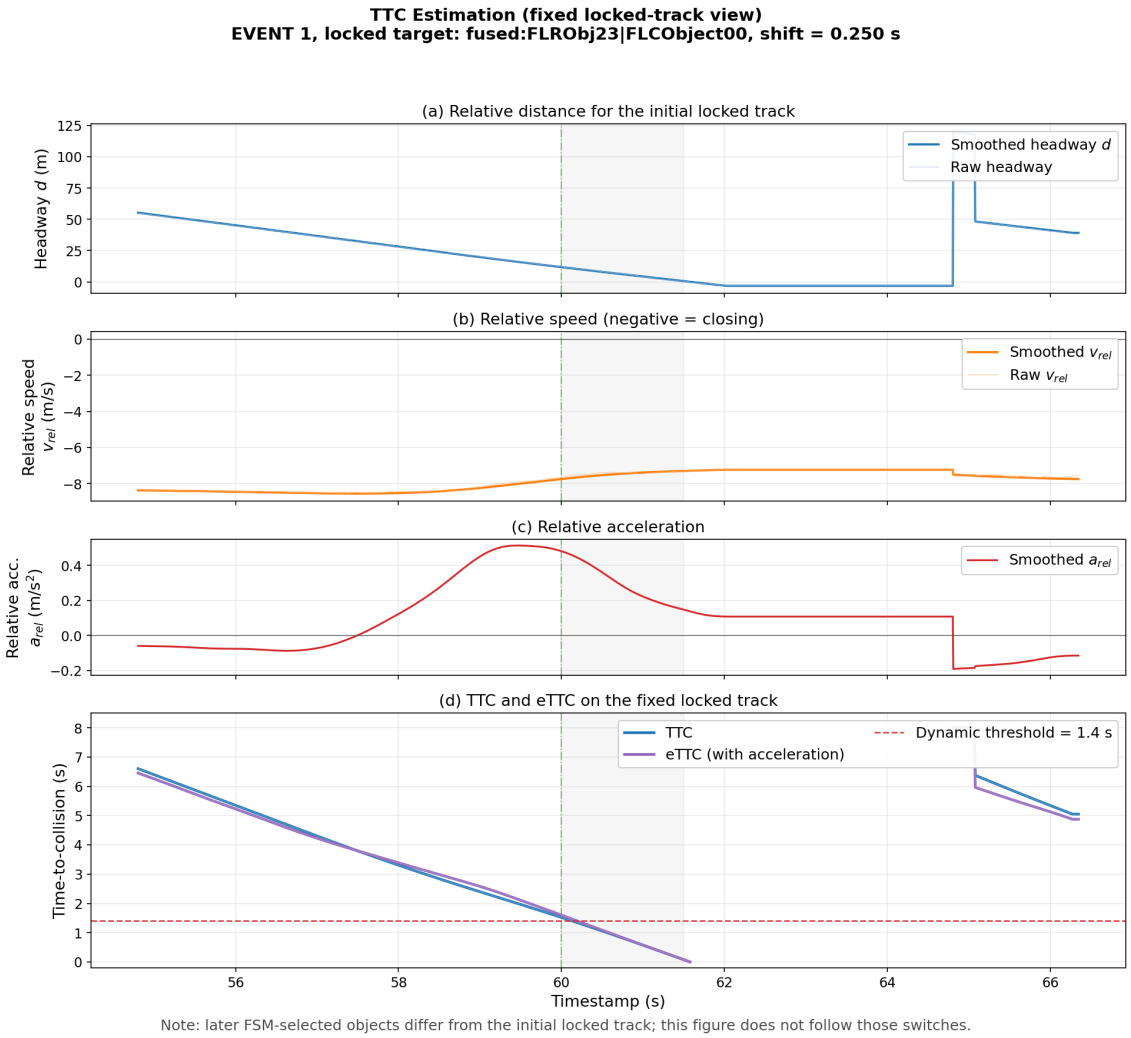

Figure 3.5: Fixed locked-track TTC estimation, including headway, relative speed, relative acceleration, and the resulting TTC/eTTC curves.

When the fused track commences later than the lock anchor (e.g. because the contour-overlap gate delays fusion confirmation), the single-target TTC computation extends its kinematic history backwards by incorporating the constituent radar-only or camera-only tracks sharing the same FLR anchor. This provides pre-trigger observability without sacrificing post-fusion accuracy.
Once the CM-active lock anchor and the corresponding track-based history have been defined, the TTC computation can be interpreted as operating on a fixed target trajectory rather than on a frame-wise reselected object. Figure 3.5 therefore presents the complete fixed-track estimation chain, from relative kinematics to TTC and eTTC, under the temporal-shift convention adopted for the offline analysis.

3.5.7 Unresolved Target Events

When no target can be locked (e.g. due to an empty target pool throughout the CM-active window), the event is classified as target-unresolved. In this case:

- TTC and eTTC are not computed (left as NaN rather than defaulting to ∞).
- The event is excluded from threshold-based false positive evaluation but retained for aggregate statistics.

This distinction prevents conflating "no measurable collision geometry" with "verified absence of collision risk."

3.6 Qualified Braking Event Screening

The QBE screening stage determines, from continuous CAN time-series recordings, which samples correspond to a genuine AEB-initiated emergency brake intervention that satisfies the regulatory and specification-level requirements. The output is a time-indexed Boolean mask indicating qualified rows and a set of discrete event IDs derived from it.

3.6.1 CM-Status Event Segmentation

The QBE screening stage operates on the candidate event windows produced in Stage C1. The interpolation of `CM_Status`, gap merging, and window construction are therefore not repeated here; only the qualification logic applied within each event window is described below. The corresponding segmentation parameters are summarised in Appendix B.2.

3.6.2 Simulated Brake Response Model

Because the vehicle CAN log does not expose the internal XBR command issued by the AEB system controller, a physics-based brake simulation model is used to reconstruct the deceleration that the AEB system demanded. The model consists of three cascaded stages: XBR profile generation, first-order brake dynamics, and speed integration.

3.6.2.1 XBR Profile Generation

For each activation window where `CM_Status` transitions into $\{2, 3\}$, the XBR command is reconstructed as a piecewise-linear ramp:

$$a_{\mathrm{XBR}}(t) = \begin{cases} 0, & t < t_0 + t_{\mathrm{offset}}, \\ a_{\mathrm{init}} + \dot{a}_{\mathrm{down}}\,(t - t_0 - t_{\mathrm{offset}}), & t_0 + t_{\mathrm{offset}} \le t < t_{\mathrm{sat}}, \\ a_{\mathrm{max}}, & t = t_{\mathrm{sat}}, \\ a_{\mathrm{max}} - \dot{a}_{\mathrm{down}}\,(t - t_{\mathrm{sat}}), & t_{\mathrm{sat}} < t \le t_{\mathrm{end}}, \text{ clamped to } \le 0, \end{cases} \tag{3.26}$$

where t_0 is the first timestamp at which `CM_Status` enters $\{2, 3\}$, and $t_{\mathrm{sat}} = t_0 + t_{\mathrm{offset}} + (a_{\mathrm{max}} - a_{\mathrm{init}})/\dot{a}_{\mathrm{down}}$ is the instant at which the ramp saturates. After saturation, the XBR recovers toward zero along a symmetric up-ramp at rate $|\dot{a}_{\mathrm{down}}|$ and is clamped to ≤ 0. The parameters are listed in Table 3.8.

Table 3.8: Brake simulation model parameters (calibrated from 7 test-track AEB events, vehicle B-1706, August 2025).

Parameter	Symbol	Value	Unit
XBR onset offset	t_{offset}	0.0	s
Initial deceleration	a_{init}	-0.5	m/s^2
Deceleration ramp rate (jerk)	$\dot{a}_{\text{down}}$	-4.5	m/s^3
Maximum deceleration	$a_{\max}$	-5.0	m/s^2
Brake realization gain	k_{gain}	0.78	—
Brake lag (rise)	τ_{rise}	0.20	s
Brake lag (fall)	τ_{fall}	0.40	s
Speed response delay	t_{speed}	0.1	s

3.6.2.2 First-Order Brake Dynamics

The realized brake deceleration is obtained by passing the scaled XBR command through a first-order exponential lag filter:

$$a_{\text{brake}}[i] = \alpha\, a_{\text{brake}}[i-1] + (1-\alpha)\, k_{\text{gain}}\, a_{\text{XBR}}[i], \tag{3.27}$$

with $\alpha = e^{-\Delta t/\tau}$, where the time constant switches between two values depending on whether the XBR demand is releasing or building:

$$\tau = \begin{cases} \tau_{\text{fall}} = 0.40\ \text{s}, & \text{if } a_{\text{XBR}}[i] > a_{\text{XBR}}[i-1]\ \text{(releasing)}, \\ \tau_{\text{rise}} = 0.20\ \text{s}, & \text{otherwise (building)}. \end{cases} \tag{3.28}$$

The asymmetric time constants model the physical observation that hydraulic brake pressure builds up faster (active pump pressurization) than it decays (passive valve release and spring return) [30].

3.6.2.3 Speed Integration

The simulated vehicle speed (in km/h) is obtained by Euler integration of the delayed brake response:

$$v[i] = \max\Big(0,\ v[i-1] + a_{\text{brake}}[i - n_{\text{delay}}]\, \Delta t \cdot 3.6\Big), \tag{3.29}$$

where $n_{\text{delay}} = \lceil t_{\text{speed}}/\Delta t \rceil$ and the factor 3.6 converts the acceleration increment from m/s to km/h. When $v \le 0.1\,\text{km/h}$, the vehicle is considered stopped, and all subsequent XBR values are zeroed.

3.6.2.4 Model Calibration

The primary goal of the calibration is to ensure that the simulated deceleration profile closely tracks the real longitudinal acceleration during the *onset ramp*—i.e. from the moment the brake command is issued until the deceleration reaches the $-1.5\,\text{m/s}^2$ qualification threshold (Criterion 1, Section 3.6.3). Accurate matching in this region is critical because the instant at which the simulation crosses $-1.5\,\text{m/s}^2$

Table 3.9: Track test scenario categories (Euro NCAP-style protocols, vehicle B-1706).

Scenario	Target	Speed	Lighting
Car-to-Car Rear Stationary	Stationary vehicle	–	Day
Bicycle Longitudinal Adult	Crossing cyclist	25 km/h	Day
Bicycle Longitudinal Adult	Crossing cyclist	50 km/h	Day
Pedestrian Nearside Adult	Crossing pedestrian	25 km/h	Day
Pedestrian Nearside Adult	Crossing pedestrian	75 km/h	Day
Pedestrian Nearside Adult	Crossing pedestrian	25 km/h	Night
Pedestrian Nearside Adult	Crossing pedestrian	75 km/h	Night
Pedestrian Farside Adult	Crossing pedestrian	50 km/h	Day
Pedestrian Nearside Child	Crossing child	50 km/h	Day
Bus stop departure	Moving/stationary	–	Day
False positive / aborted crossing	Pedestrian abort	–	Day

determines both the start of the QBE window and the corresponding TTC estimate used for false-positive judgment; a timing error in the onset translates directly into a TTC bias.

The model was calibrated against 13 controlled AEB braking events recorded on a dedicated test track using development vehicle B-1706 in August 2025 (see Section 2.1.2). The events span seven Euro NCAP scenario types—BBLA50 (car-to-bicycle lateral), BCRS (car-to-car rear stationary), BPFA-50 (car-to-pedestrian farside), BPNA-75 and BPNA25 (car-to-pedestrian nearside adult), BPNA_NIGHT variants, and BPNC50 (car-to-child)—at speeds ranging from 20 to 51 km/h. Table 3.9 lists the full set of track test scenario categories. For each event, the measured longitudinal acceleration from the CAN bus served as the ground-truth reference.

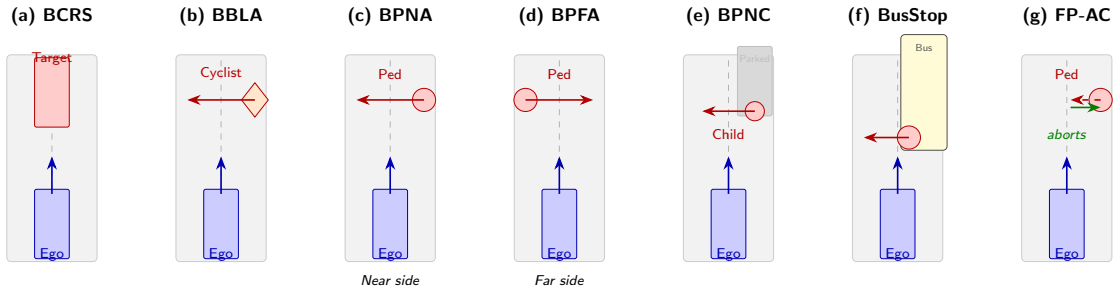

Figure 3.6: Bird's-eye view sketches of the track test scenario types. Blue rectangles represent the ego vehicle (driving upward); red symbols denote targets. Arrows indicate direction of motion. In (g), the dashed arrow shows the pedestrian's initial crossing path and the solid green arrow indicates the abort (return to kerb).

From the 13 track test events, seven remained after excluding scenarios without a full AEB braking intervention (FCW-only, bus-stop, aborted crossings), events with driver-initiated braking preceding the AEB command (onset delay ≤ 0), and events

where the driver's brake pedal signal was active. The calibration set characterises two systematic discrepancies:

- **Onset delay**: median 310 ms (IQR [276, 327] ms), indicating a consistent ∼0.3 s lag between the commanded and realised deceleration.
- **Peak gain ratio**: median $G = 0.78$ (range 0.59–1.29), reflecting that the physical brake generally undershoots the command at low-to-moderate speeds.

The detailed per-event statistics and exclusion criteria are provided in Appendix C.3.

Correction model. To bridge the gap between the idealised simulation and the measured response, a four-parameter cascaded correction model was fitted:

$$a_{\text{corr}}(t) = \text{LPF}_{\tau}\Big(\text{RateLim}_{R}\Big(G \cdot a_{\text{brake}}(t - \Delta t)\Big)\Big), \tag{3.30}$$

where Δt is a time shift (system reaction time), G is an amplitude gain, R is a slew-rate limiter modelling the hydraulic valve opening rate, and τ is a first-order lag time constant capturing calliper compliance.

The uncorrected simulated brake request overlaid on the real longitudinal acceleration for all 13 analysed events is shown in Appendix D (Figures D.1–D.4), with CM-status transitions indicated by background shading.

Figure 3.7 compares the original and corrected simulation against the real measurement after applying the fitted cascade correction for two representative calibration events (BBLA50 at 35 km/h and BCRS at 30 km/h). The remaining calibration events, grouped by scenario type, are shown in Appendix D (Figures D.5–D.7).

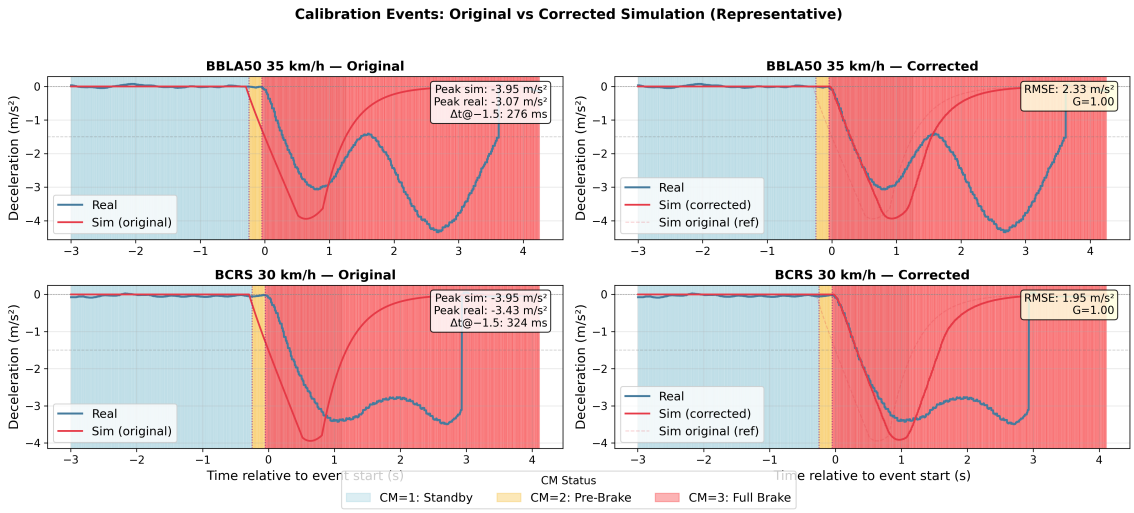

Figure 3.7: Original simulation (left column) versus corrected simulation (right column) for two representative calibration events. The dashed grey curve in the right column shows the original simulation for reference; the solid red curve shows the corrected output after applying the cascade correction (Eq. 3.30).

The calibrated parameters are specific to the Volvo city bus brake hardware. For the FOT fleet (B-2035 through B-2055), which shares the same brake system architecture as B-1706, the parameters are assumed to transfer without adjustment.

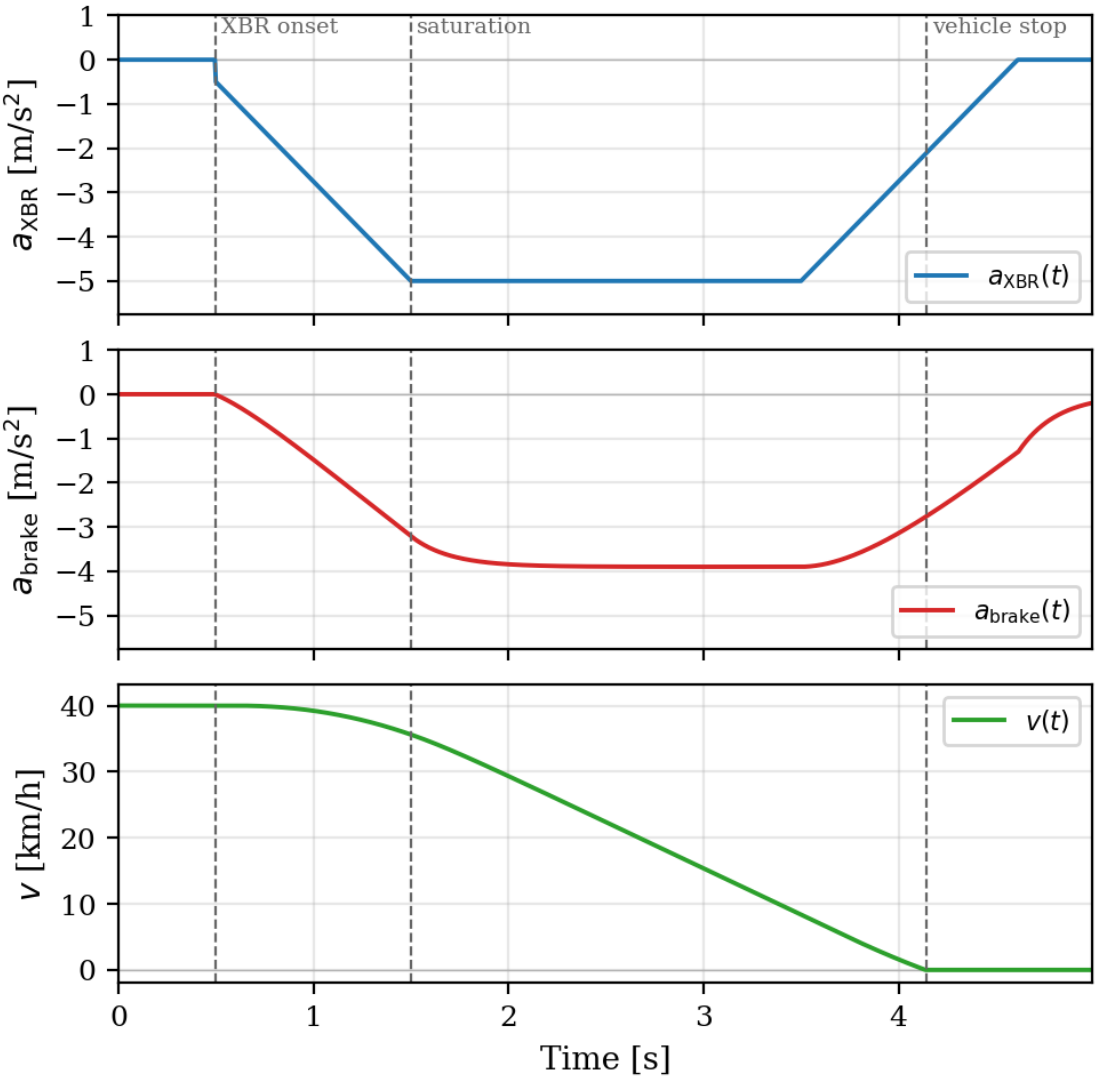

Figure 3.8: Three-stage brake simulation model. Top panel: piecewise-linear XBR command profile $a_{\mathrm{XBR}}(t)$ (Eq. 3.26). Middle panel: realised brake deceleration $a_{\mathrm{brake}}(t)$ after first-order lag filtering (Eq. 3.27) with asymmetric time constants. Bottom panel: integrated vehicle speed $v(t)$ (Eq. 3.29). Dashed vertical lines mark the XBR onset, saturation, and vehicle-stop instants. Parameters from Table 3.8.

In the downstream false-positive detection pipeline, only the onset delay compensation (∼300 ms, corresponding to an effective TTC correction of ∼0.15 s) is applied; the full amplitude correction is not used because the qualification criterion (Section 3.6.3) depends on the sign rather than the magnitude of the simulated deceleration.

3.6.3 Qualification Criteria

A time sample is marked as a QBE if and only if four criteria are satisfied simultaneously:

$$q[i] = \underbrace{c_{\text{decel}}[i]}_{\text{Criterion 1}} \ \wedge \ \underbrace{c_{\text{aeb}}[i]}_{\text{Criterion 2}} \ \wedge \ \underbrace{c_{\text{gate}}[i]}_{\text{Criterion 3}} \ \wedge \ \underbrace{c_{\text{range}}[i]}_{\text{Criterion 4}} . \tag{3.31}$$

Criterion 1 — Deceleration gate. The simulated brake response must equal or exceed $-1.5\,\mathrm{m/s^2}$ in magnitude:

$$c_{\text{decel}}[i] \iff a_{\text{brake}}[i] \leq -1.5\ \mathrm{m/s^2}. \tag{3.32}$$

This threshold originates from the TfL Bus Vehicle Specification v2.6, which defines it as the lower bound for an emergency braking phase. When the simulated brake response is unavailable (e.g. due to missing input signals), the pipeline falls back to the raw `LongitudinalAcceleration` signal (or `EgoMotionXYplane_LongAcc`) as an alternative deceleration source.

Criterion 2 — AEB active status. The AEBS must be in an active braking phase:

$$c_{\text{aeb}}[i] \iff \texttt{CM\_Status}[i] \in \{2, 3\}, \tag{3.33}$$

where status 2 denotes active partial braking (level 1) and status 3 denotes active full braking (level 2).

Criterion 3 — Minimum speed gate. A global minimum speed threshold is applied regardless of scenario type, as defined in UN R-131 [31] §5.2.1.3:

$$c_{\text{gate}}[i] \iff v[i] > 10\ \mathrm{km/h}. \tag{3.34}$$

This criterion reflects the minimum speed at which the AEB system function is required to be active and is therefore applied to all events. When no additional scenario-specific speed restriction is imposed, the resulting event set is referred to as the speed-gate set.

3.6.4 Event Edge Detection

Individual QBEs are extracted from the continuous Boolean mask $q[i]$ by detecting rising edges:

$$e_{\text{start}}[i] = q[i] \ \wedge \ \neg\, q[i-1]. \tag{3.35}$$

Each event is assigned a monotonically increasing integer identifier via a cumulative sum of the start flags:

$$\text{EventID}[i] = \begin{cases} \displaystyle\sum_{k=1}^{i} e_{\text{start}}[k], & \text{if } q[i] = \text{true}, \\ \varnothing, & \text{otherwise.} \end{cases} \tag{3.36}$$

No minimum duration or secondary gap-merge is applied at this stage; every contiguous block of qualified samples, regardless of its length, constitutes a separate event.

3.6.5 Signal Selection and Fallback Strategy

Across the FOT fleet, different vehicles and firmware versions expose the same physical quantity under varying CAN signal names. To handle this heterogeneity the pipeline employs a scored-fallback strategy: for each required signal category (e.g., vehicle speed, longitudinal acceleration) a ranked list of candidate signal names is defined. When multiple candidates are present in the decoded data, the pipeline scores each by data quality (number of unique values, non-null count, and non-zero count) and selects the highest-scoring signal; if only one candidate exists it is selected directly.
The candidate signal lists and channel-retry logic are documented in Appendix B.6; the unit-detection heuristic is described in Appendix B.3.

3.7 False Positive Judgment

This section describes the implementation of the false positive classification logic defined in Section 2.3.3. The classification operates on the TTC estimation output and produces a per-event binary label (FP or TP).

3.7.1 TTC Resolution Chain

Not all input event records contain a pre-computed TTC column. When a direct TTC value is unavailable, the pipeline derives it from lower-level signals using the following resolution chain:

1. **Direct TTC column:** If a `TTC` or `TTC_smooth` column exists in the input record, it is used directly.
2. **Derived from headway and relative speed:** The headway (`LongPos_used`), relative speed (`LongVel_used`), and relative acceleration (`LongAcc_used`) are smoothed according to the TTC estimation procedure described in Section 3.5 and then used to compute eTTC together with TTC.
3. **Unavailable:** If neither direct nor derived TTC can be computed (e.g. no radar target present), Condition A cannot be evaluated and defaults to `False`.

3.7.2 Condition A: Excessive TTC

At the AEB anchor point, the TTC estimate selected by the resolution chain is compared against the dynamic threshold $\mathrm{TTC}_{\mathrm{threshold}}$ (Eq. 2.2, Section 2.3.3). The TTC estimate is sampled within a search window of ± 0.5 s around the AEB anchor; within that window, the sample closest in time to the anchor is selected. If the selected TTC estimate exceeds $\mathrm{TTC}_{\mathrm{threshold}}$, Condition A is satisfied.

3.7.3 Condition B: Absent Driver Brake Response

The pipeline searches for evidence of driver braking within a configurable window after the AEB anchor (default: 1.2 s). The search scope extends over the full recording for the same source file, not only the qualified event slice, so that driver braking occurring after the qualified interval ends is still captured. The driver is considered to have braked if either of the following is detected:

- `BrakeSwitch` > 0.5 (discrete brake switch activated, also covering `BrakePedalSwitch_BB1_X_V` and `EBSBrakeSwitch` as fallbacks), or
- `BrakePedalPosition` $> 1.0\%$ (continuous pedal position above noise threshold).

If no driver braking is detected within the window, Condition B is satisfied.

3.7.4 FP Determination

The final classification follows the disjunctive rule FP = Condition A $\vee$ Condition B (Eq. 2.1–2.3, Section 2.3.3).

Each event record includes the individual condition flags, the measured TTC, the computed threshold, and the driver brake delay (if detected), ensuring full traceability of the classification decision.

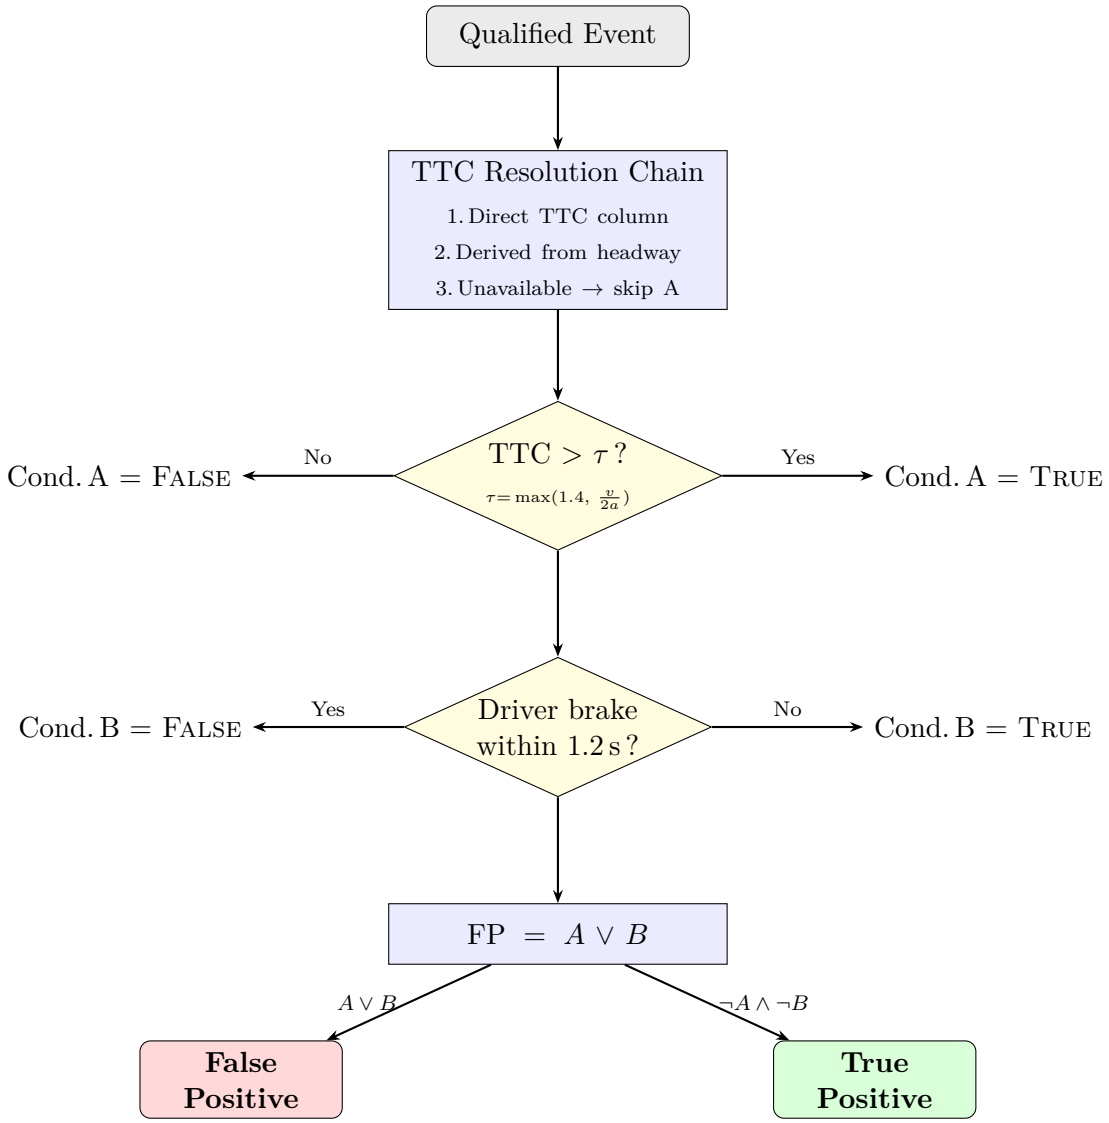

Figure 3.9: False positive judgment decision flowchart. For each qualified event, the TTC resolution chain selects the best available TTC estimate. Condition A compares the measured TTC against the speed-dependent dynamic threshold $\max(1.4\text{ s},\ v_{\text{ego}}/2a_{\text{LPoB}})$. Condition B checks for driver brake response within 1.2 s of the AEB anchor. The event is classified as FP if either condition is satisfied ($\text{FP} = A \vee B$).

3.8 False Positive Legitimacy Verification

The binary FP/TP label produced by the FP judgment stage (Section 3.7) is a rule-based decision grounded in TTC and driver-response criteria defined by the TfL Bus Vehicle Specification. To assess whether those labels are physically plausible and to identify the dominant failure modes, a post-hoc verification layer is applied. It comprises three steps: scenario-bucket extraction, rule-based event grouping, and within-group distributional analysis.

3.8.1 Scenario-Bucket Feature Extraction

For each qualified braking event, a feature vector is extracted at the AEB anchor — the first time sample at which `CM_Status` $\in \{2, 3\}$ within the event. The extracted features are organised into six categories:

1. **Ego kinematics:** vehicle speed v (km/h), longitudinal acceleration a_x (m/s^2), a finite-difference jerk proxy $\dot{a}_x \approx \Delta a_x / \Delta t$ computed over ± 2 samples around the anchor, and the braking duration (time span of contiguous qualified rows).
2. **Forward-target state:** longitudinal distance (`LongPos_used`), relative speed (`LongVel_used`), and a ternary target-presence flag:

$$\text{target} = \begin{cases} \text{PRESENT}, & d > 0, \\ \text{ABSENT}, & d \leq 0, \\ \text{UNKNOWN}, & d = \varnothing, \end{cases} \tag{3.37}$$

   where d denotes the decoded longitudinal position.
3. **TTC context:** Kalman-smoothed TTC (`TTC_smooth`), enhanced TTC (`eTTC_smooth`), headway, and relative speed at the anchor.
4. **Radar target attributes:** object classification (`ObjClass`), dynamic classification (`ObjDynClass`), track status (`TrackStatus`), existence confidence, and lateral position state of the primary tracked target.
5. **Driver state:** brake-switch status, brake-pedal position, and three binary driver-action flags evaluated over a 1.2 s window after the anchor — accelerator pedal ($> 5\%$), kickdown switch, and steering-wheel angle change ($|\Delta\theta| > 0.05$ rad ≈ 3ř) [12].
6. **FP judgment decision:** the FP/TP label, measured TTC, computed threshold, and the individual Condition A / B flags, joined from the FP judgment output by event ID.

In addition, a set of Boolean data-availability flags records which signals were successfully decoded for each event, supporting downstream quality filtering.
The result is a single-row-per-event table (`scenario_buckets.csv`) that enables cross-event statistical analysis without re-reading the full time-series data.

3.8.2 Rule-Based Event Grouping

Events are partitioned into four mutually exclusive groups using two physically motivated axes: target presence and TTC value at the AEB anchor relative to the

per-event dynamic threshold at the AEB anchor relative to the per-event dynamic threshold (Eq. 2.2). If the per-event threshold is unavailable (e.g. speed data missing), the floor value of 1.4 s is used as a fallback. The grouping rules are summarised in Table 3.10.

Table 3.10: Rule-based event grouping.

Group	Condition	Interpretation
G0	target = PRESENT ∧ TTC ≤ $\text{TTC}_{\text{threshold}}$	Near-target AEB activation
G1	target = PRESENT ∧ TTC > $\text{TTC}_{\text{threshold}}$	Premature trigger (target far from collision)
G2	target = PRESENT ∧ TTC undefined (NaN)	Target receding (TTC undefined)
G3	target ∈ {ABSENT, UNKNOWN}	No forward radar target (ghost / sensor artefact)

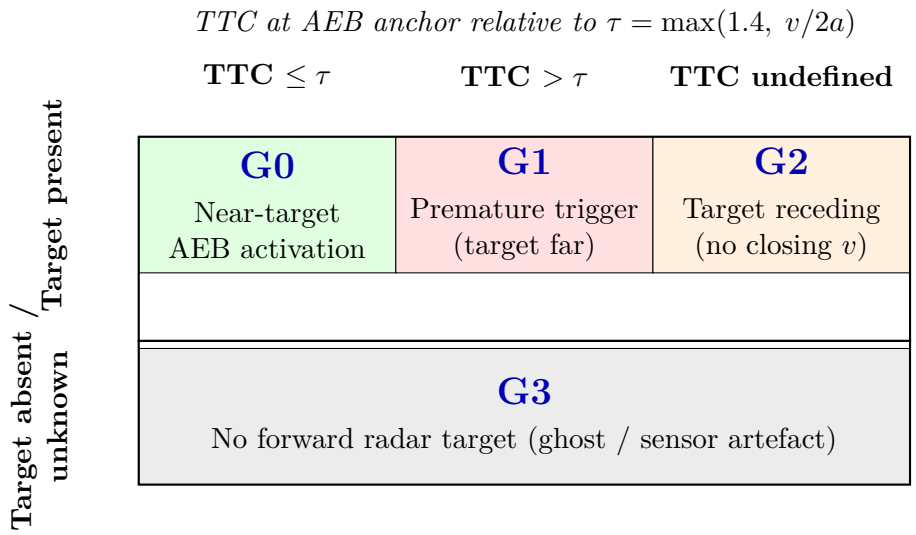

Figure 3.10: Rule-based event grouping matrix. Events are partitioned along two axes: target presence (present vs. absent/unknown) and TTC relative to the per-event dynamic threshold (TTC ≤ τ vs. TTC > τ vs. undefined). The four resulting groups (G0–G3) correspond to the conditions in Table 3.10.

Note that the TTC boundary between G0 and G1 is *not* a fixed value but varies per event with the ego speed; typical values in the FOT dataset ranges from 1.40 s to approximately 2.2 s. The TTC value used for grouping is the `ttc_cv_smooth_s`, whereas the FP judgment Condition A evaluates the eTTC, which incorporates the rate of change of the closing speed (relative acceleration). When relative acceleration is negative (i.e. the gap is closing at an increasing rate), eTTC yields lower values than TTC; consequently, some events whose TTC exceeds the threshold (placing them in G1) have an eTTC that falls below the threshold, causing Condition A to

evaluate as false and the event to be classified as TP. G1 is therefore *not* expected to be 100% FP.
This grouping serves two purposes: (i) it highlights the divergence between the two TTC models — G1 events classified as TP are precisely those where the eTTC correction is decisive; and (ii) it isolates distinct physical failure modes (ghost triggers in G3, premature activations in G1) to guide targeted mitigation.

3.8.3 Within-Group Distributional Analysis

Within G0, the FP and TP subpopulations are compared across the scenario-bucket features to identify which signals carry the most discriminative power. The analysis includes:

- **TTC distribution:** overlaid histograms and boxplots of TTC values, stratified by FP/TP label. A clear separation between the two distributions validates that TTC is the primary discriminating variable.
- **Target distance and kinematics:** scatter plots of TTC versus forward target distance, coloured by FP/TP. Physical consistency requires that TP events cluster at shorter distances and lower TTC values [30].
- **Target contingency tables:** cross-tabulations of `ObjClass`, `ObjDynClass`, and `TrackStatus` against the FP/TP label, identifying whether certain object types or tracking states are over-represented among false positives.
- **Driver brake response:** contingency tables of `BrakeSwitch` status at the AEB anchor. A higher proportion of active braking among TP events corroborates Condition B of the false positive definition.
- **Driver action features:** frequency tables for the three 1.2 s window indicators (accelerator pedal, kickdown, steering) stratified by FP/TP, presented as annotated heatmaps.
- **Speed and braking duration:** compared to confirm that these features do *not* differ systematically between FP and TP, ruling out confounding by operating-point bias.
- **System availability:** cross-tabulations of
  - `CMReducedAbilityStatus`
  - `EmergencyBrakingInhibition`
  - `VDCFullyOperational`

  against the FP/TP label, to whether system degradation correlates with FP occurrence.

Summary statistics (mean, standard deviation, and the 5th, 25th, 50th, 75th, and 95th percentiles) are computed for each feature within each subgroup and reported alongside the visual analysis. Per-scenario feature–quantile heatmaps (row-normalised) provide a compact overview of how median feature values vary across G0–G3. Per-vehicle FP rate breakdowns are included to detect fleet outliers.

3.8.4 Cross-Group Consistency Check

As a final consistency check, the FP rate is computed per group and compared against the expected pattern:

- G1 (high TTC) should exhibit $\approx 100\%$ FP rate, since the target is beyond the collision boundary defined by the per-event TTC threshold.
- G3 (no target) should exhibit a high FP rate, since the absence of a radar target implies a ghost trigger or sensor false alarm — a scenario recognised in International Organization for Standardization (ISO) 21448 (SOTIF) [32] as a known limitation of perception-based safety systems.
- G0 (near-target) should contain the majority of TP events; any TP events appearing in G1–G3 would indicate an error in the upstream target-matching or TTC computation stages.

Deviations from these expectations trigger a review of the upstream pipeline stages (target matching, TTC computation, or signal decoding) rather than the FP classification logic itself.

4

Results

This chapter presents the outputs of the 15-stage detection and classification pipeline applied to the full fleet dataset. Section 4.1 summarises the data scope and processing yield. Section 4.2 characterises the qualified braking events in terms of ego kinematics and system state. Section 4.3 reports the FP/TP classification outcome and analyses the TTC distribution. Section 4.4 introduces a rule-based grouping of events by target presence and TTC level, followed by per-group deep dives. Finally, Sections 4.5 and 4.6 examine the G0 borderline cases and driver behaviour patterns that underlie the classification results.

4.1 Processing Summary

The batch processing pipeline was executed on data retrieved from the fleet S3 storage on June 1, 2026. A total of 3 160 CAN MF4 files from 20 vehicles (B-2035 through B-2055) were processed, covering logger recordings from December 2023 to May 2026. Of these, 3 151 files were decoded successfully (9 failed due to decoding errors or corrupted recordings).

Each file corresponds to one continuous recording session—typically a single driver shift or service block, spanning from logger power-on to power-off. A session therefore covers an entire sequence of route services, depot manoeuvres, and layovers during which the vehicle was operating. Session durations range from under one hour to over eight hours, depending on the driver's shift length and logger configuration.

Each file was evaluated under the speed-gate scenario definition, which requires ego speed $v > 10\,\mathrm{km/h}$ in addition to the qualified-braking criteria (deceleration threshold and AEB status).

Applying the speed-gate QBE criteria defined in Section 3.6, 117 files (3.7%) yielded at least one qualified event, producing 123 events in total. The FP judgment step (Section 3.7) identified 63 events as FPs (51.2%) and 60 as TPs.

Table 4.1 provides a per-vehicle breakdown for the speed-gate scenario. Fourteen of the 20 vehicles produced at least one qualified event; the remaining six recorded zero qualifying activations during the observation period.

Table 4.1: Per-vehicle processing summary (speed-gate scenario).

Vehicle	Files	OK	Errors	Events	FP	TP
B-2035	110	110	0	9	0	9
B-2036	92	92	0	8	4	4
B-2037	119	119	0	6	0	6
B-2038	141	141	0	10	4	6
B-2039	122	122	0	4	3	1
B-2040	128	128	0	9	3	6
B-2041	71	71	0	1	0	1
B-2042	184	184	0	3	0	3
B-2043	338	336	2	0	0	0
B-2044	62	62	0	2	0	2
B-2046	263	263	0	0	0	0
B-2047	199	198	1	0	0	0
B-2048	304	299	5	0	0	0
B-2049	32	32	0	2	1	1
B-2050	184	184	0	7	4	3
B-2051	115	114	1	0	0	0
B-2052	200	200	0	9	6	3
B-2053	138	138	0	0	0	0
B-2054	97	97	0	16	11	5
B-2055	261	261	0	37	27	10
Total	**3 160**	**3 151**	**9**	**123**	**63**	**60**

Note on data freshness. The fleet logger data is continuously uploaded to S3 as vehicles operate. The results presented in this chapter reflect a snapshot as of June 1, 2026. The pipeline supports incremental processing, so subsequent runs extend the dataset without reprocessing previously analysed files. Proportional trends (e.g. FP rates per group) are expected to remain stable as the dataset grows, while absolute counts will increase.

Video review. Among the fleet recordings, only one event had a corresponding video-embedded MF4 file available. This event was classified as TP by the pipeline and captured a complete AEB intervention sequence, with `CM_Status` progressing from 1 (Enabled) through 2 (PreMitigation) and 3 (Mitigation) to 4 (Finished). Visual review of the forward-facing video confirmed the presence of a genuine collision threat, consistent with the pipeline's TP classification. The CAN-based FP/TP classification (Conditions A and B) therefore remains the primary basis for all results presented in this chapter.

4.2 Qualified Brake Event Detection

Of the 3 160 CAN MF4 files processed, 117 files (3.7%) contained at least one QBE satisfying all speed-gate qualified-braking criteria simultaneously: deceleration $a \leq -1.5\,\mathrm{m/s^2}$, AEB status `CM_Status` $\in \{2, 3\}$, and ego speed $v > 10\,\mathrm{km/h}$. A total of 123 events were identified.

4.2.1 Event Yield per Vehicle

The event yield varied considerably across the fleet. Fourteen out of 20 vehicles produced speed-gate-qualified events, with per-vehicle rates ranging from 1.4 to 16.5 events per 100 files (Table 4.1). Vehicle B-2055 contributed the most events (37), while ten vehicles recorded zero qualifying activations throughout the observation period.

4.2.2 Ego Kinematics at AEB Activation

Table 4.2 summarises the ego-vehicle kinematics at the AEB anchor point (first sample where `CM_Status` $\in \{2, 3\}$) for all 123 speed-gate events.

Table 4.2: Ego kinematics at the AEB anchor point ($n = 123$ speed-gate events).

Parameter	Mean	Std	Min	Max
Speed (km/h)	23.4	8.0	10.5	44.7
Braking duration (s)	0.64	0.48	0.01	1.53

The speed distribution is concentrated in the 20–35 km/h range (64 of 123 events), consistent with urban bus operation. Forty-eight events (39.0%) fell in the 10–20 km/h band—these are the low-speed activations captured only by the speed-gate criteria. Only three events exceeded 40 km/h.

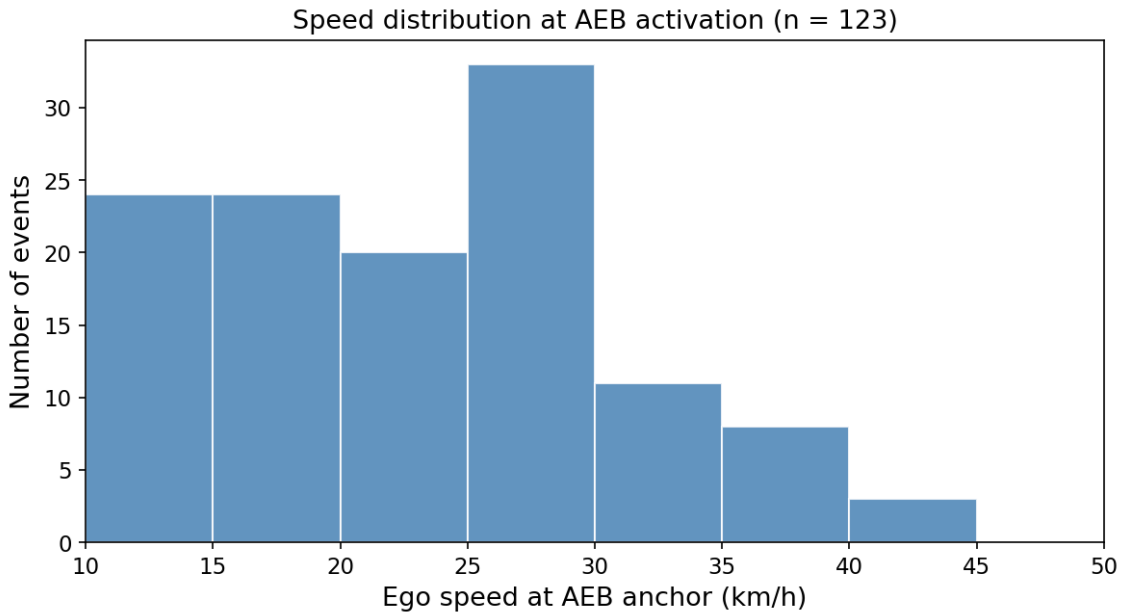

Figure 4.1: Distribution of ego speed at the AEB anchor point for all 123 speed-gate events (5 km/h bins).

4.2.3 AEB Status and Forward Target

Of the 123 speed-gate events, 60 (48.8%) had `CM_Status` = 2 (warning with partial braking) and 63 (51.2%) had `CM_Status` = 3 (full emergency braking) at their latest activation level.
A forward radar target was present in 63 events (51.2%), while 60 events (48.8%) had no identifiable target at the moment of activation, suggesting possible ghost detections or transient sensor artefacts.

4.2.4 Driver State at Activation

At the AEB anchor, 74 of 123 events (60.2%) showed the brake switch already active (`BrakeSwitch` = 1), indicating the driver was already braking when the system intervened.
All 123 events showed the AEB system operating without degradation:

- `CMReducedAbilityStatus` = 0
- `EmergencyBrakingInhibition` = 0
- `VDCFullyOperational` = 1

across the board.

4.3 False Positive Classification

The FP judgment step classified each of the 123 speed-gate-qualified events as either FP or TP according to the two-condition disjunction defined in TfL internal specification (Section 3.7). An event is classified as FP if either Condition A or Condition B is satisfied.

4.3.1 Overall Classification

Of the 123 events, 63 (51.2%) were classified as FPs and 60 (48.8%) as TPs. Table 4.3 shows how the two conditions contributed to the FP count.

Table 4.3: FP condition breakdown ($n = 123$ speed-gate events). Conditions are not mutually exclusive; an event may satisfy both.

Condition combination	Events	Share
Condition A only (excessive TTC)	21	17.1%
Condition B only (absent driver response)	21	17.1%
Both A and B	21	17.1%
Neither (True Positive)	60	48.8%
Total	123	100%

Condition A and Condition B contributed equally to the FP classification, each triggering in 42 of the 63 FP events (66.7%). In 21 of those, both conditions were satisfied simultaneously, reinforcing the FP indication.

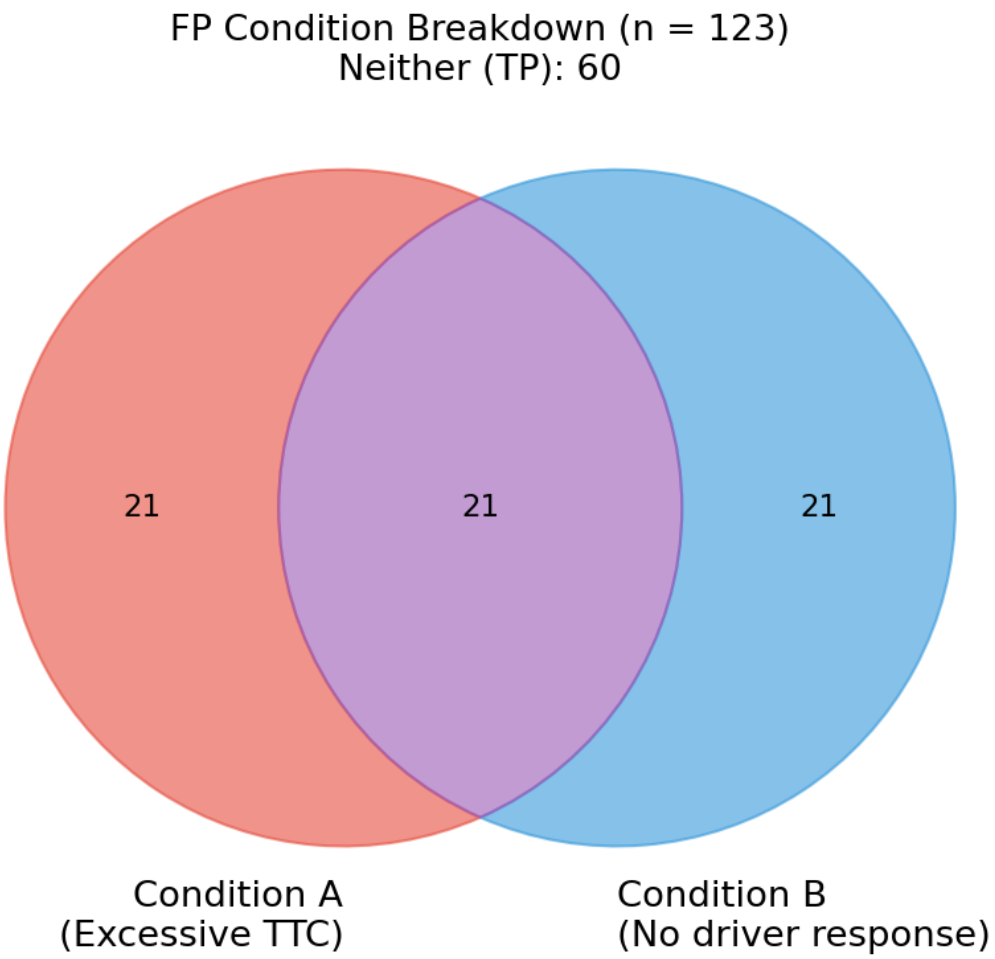

Figure 4.2: Overlap of FP Condition A and Condition B among the 123 speed-gate events. A-only: 21, B-only: 21, Both: 21, Neither (TP): 60.

Twenty-one events were classified as FP solely through Condition B—in these cases the measured TTC was below the per-event threshold, yet the driver did not apply the brake pedal within 1.2 s.

4.3.2 TTC Distribution

A valid TTC measurement was available for 57 of 123 speed-gate events; the remaining 66 lacked a computable TTC (e.g. no converging radar target). Table 4.4 compares the TTC distribution between FP and TP events.

Table 4.4: TTC at the moment of AEB braking command, split by classification ($n = 57$ speed-gate events with valid TTC).

	n	Mean (s)	Std (s)	Min (s)	Max (s)
FP	45	1.89	0.59	1.05	5.12
TP	12	1.23	0.22	0.92	1.65
All	57	1.75	0.60	0.92	5.12

The TTC distributions for FP and TP events show clear separation (mean 1.89 s vs 1.23 s). FP events had substantially higher TTC at the moment of AEB activation, reflecting the Condition A criterion: events where the system intervened despite a large time margin to collision are classified as FP. All 12 TP events with valid TTC had values below 1.65 s, consistent with genuinely close encounters.
The per-event TTC threshold ranged from 1.40 s to 2.07 s (mean 1.45 s), calculated as $\max(1.4,\ v/(2 \cdot 3.0))$ per Bus Spec §4.3.2.3.

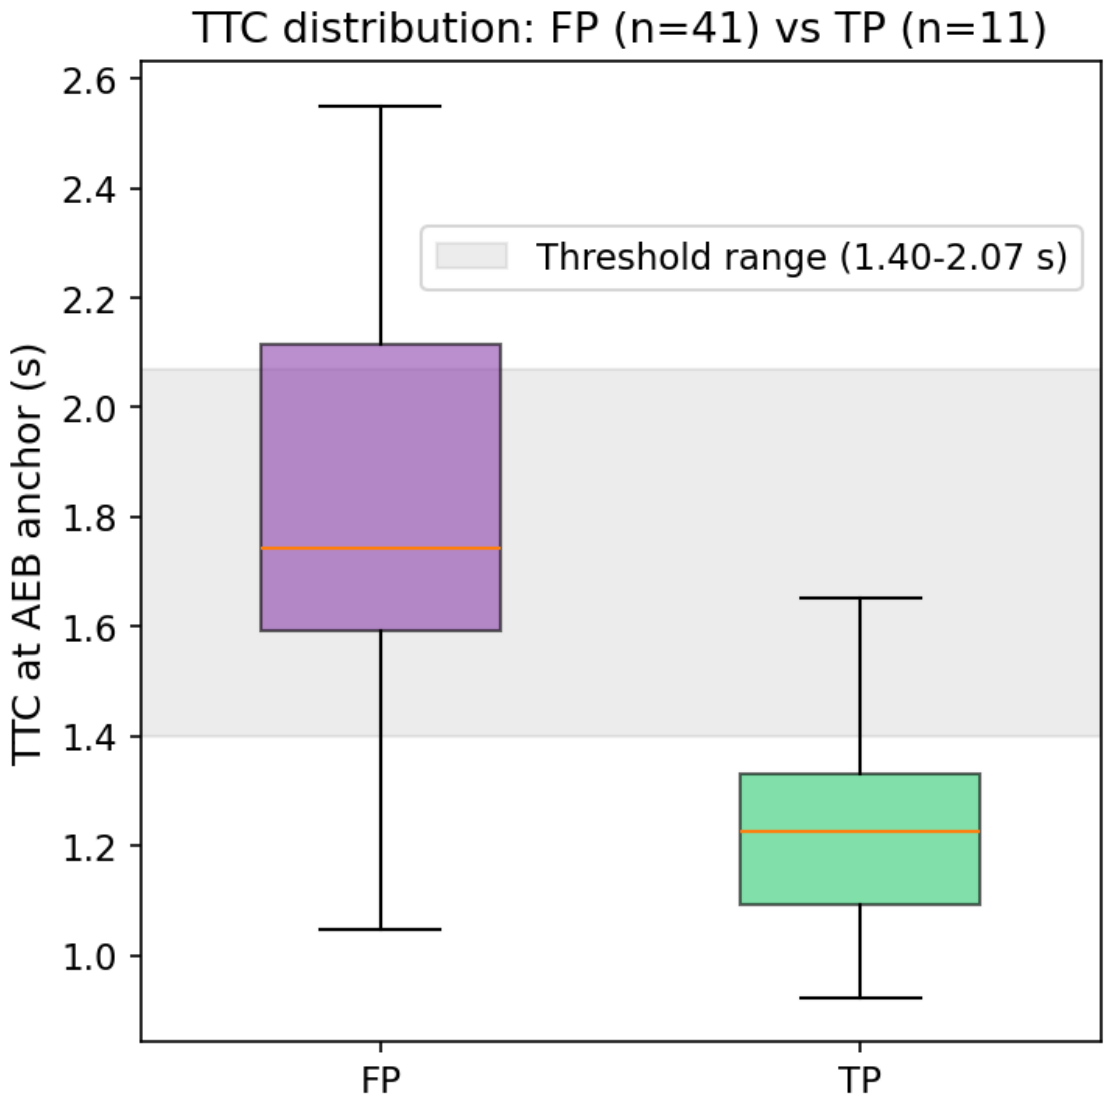

Figure 4.3: Box plot of TTC at the braking command for FP and TP events ($n = 57$ speed-gate events with valid TTC). The grey band marks the per-event threshold range (1.40–2.07 s).

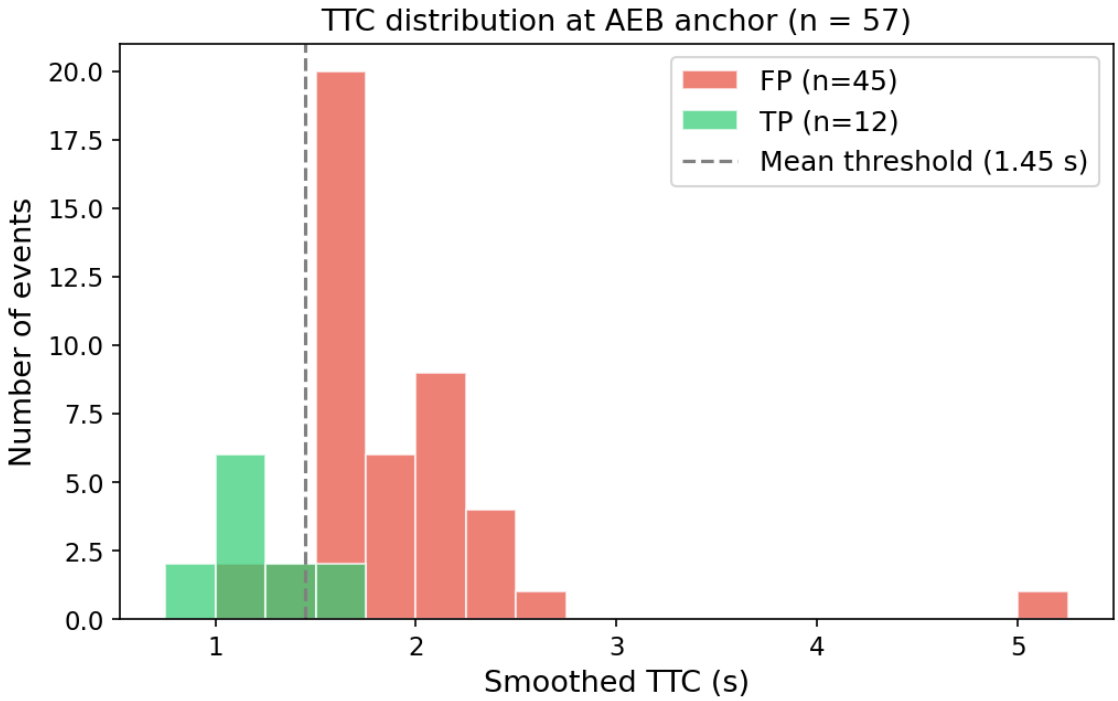

Figure 4.4: Histogram of smoothed TTC at the AEB anchor for all 57 speed-gate events with valid TTC (0.25 s bins), coloured by FP/TP classification. The dashed vertical line marks the mean per-event threshold (1.45 s).

4.3.3 Events Without TTC

Sixty-six events had no computable TTC at the braking anchor. Of these, 48 were classified as TP (neither condition triggered) and 18 as FP (via Condition B). These no-TTC events are further examined in Section 4.4 under groups G2 and G3.

4.3.4 Per-Vehicle FP Rate

The FP rate varied across vehicles as shown in Table 4.5.

Table 4.5: FP rate per vehicle, speed-gate scenario (only vehicles with events).

Vehicle	Events	FP	TP	FP rate (%)
B-2035	9	0	9	0.0
B-2036	8	4	4	50.0
B-2037	6	0	6	0.0
B-2038	10	4	6	40.0
B-2039	4	3	1	75.0
B-2040	9	3	6	33.3
B-2041	1	0	1	0.0
B-2042	3	0	3	0.0
B-2044	2	0	2	0.0
B-2049	2	1	1	50.0
B-2050	7	4	3	57.1
B-2052	9	6	3	66.7
B-2054	16	11	5	68.8
B-2055	37	27	10	73.0
Total	**123**	**63**	**60**	**51.2**

B-2055 (73.0%) and B-2039 (75.0%) had the highest FP rates, while several vehicles (B-2035, B-2037, B-2041, B-2042, B-2044) recorded zero FP events.

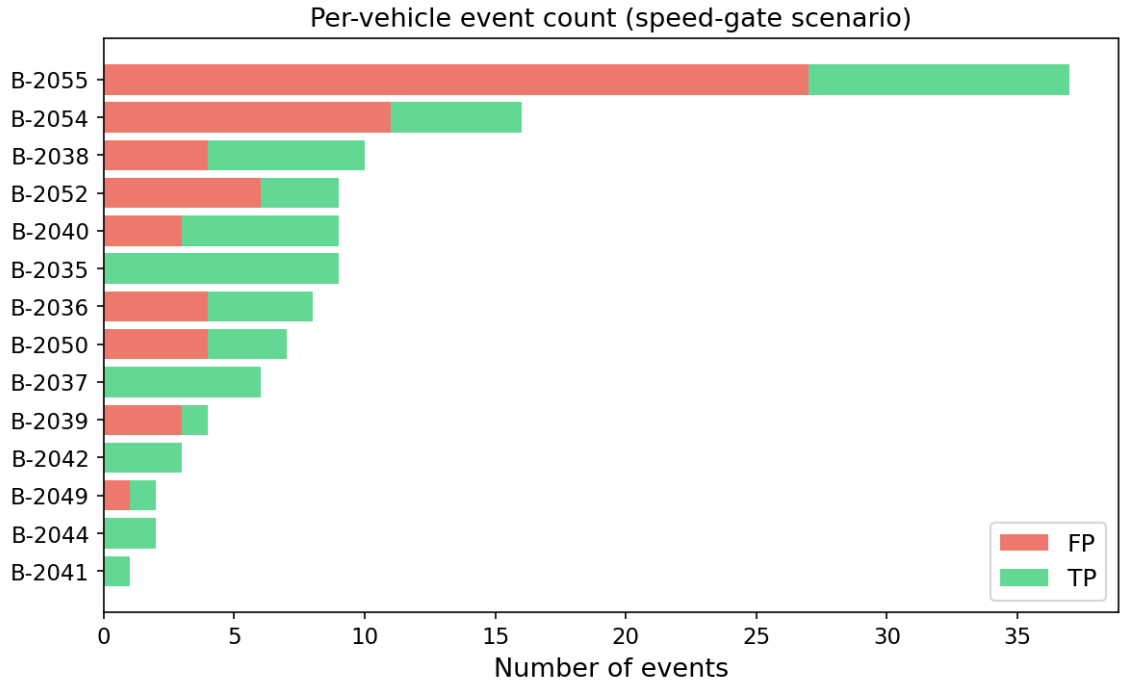

Figure 4.5: Per-vehicle event count (speed-gate scenario), sorted by total events descending. Red segments indicate FP events, green segments indicate TP events. Only vehicles with at least one event are shown.

4.4 Rule-Based Event Grouping

To gain insight into the physical circumstances of each AEB activation, the 123 speed-gate events were assigned to four mutually exclusive groups based on two observable dimensions: whether a forward radar target was present at the AEB anchor point, and how the smoothed TTC compared to the per-event threshold $\tau = \max(1.4,\ v/(2 \cdot 3.0))$. Table 4.6 defines the grouping rules.

Table 4.6: Rule-based event group definitions.

Group	Rule	Physical interpretation
G0	Target present, TTC $\leq \tau$	Typical near-target AEB activation
G1	Target present, TTC $> \tau$	Target far from collision risk
G2	Target present, TTC undefined	Target moving away or no closing velocity
G3	No radar target	Ghost trigger / sensor false alarm

4.4.1 Group Distribution and FP Rates

Table 4.7 presents the cross-tabulation of group assignment against FP/TP classification for the speed-gate scenario.

Table 4.7: Event count and FP rate per group (speed-gate scenario).

Group	Events	FP	TP	FP rate (%)
G0 (normal)	14	3	11	21.4
G1 (high TTC)	38	38	0	100.0
G2 (target, no TTC)	11	5	6	45.5
G3 (no target)	60	17	43	28.3
Total	**123**	**63**	**60**	**51.2**

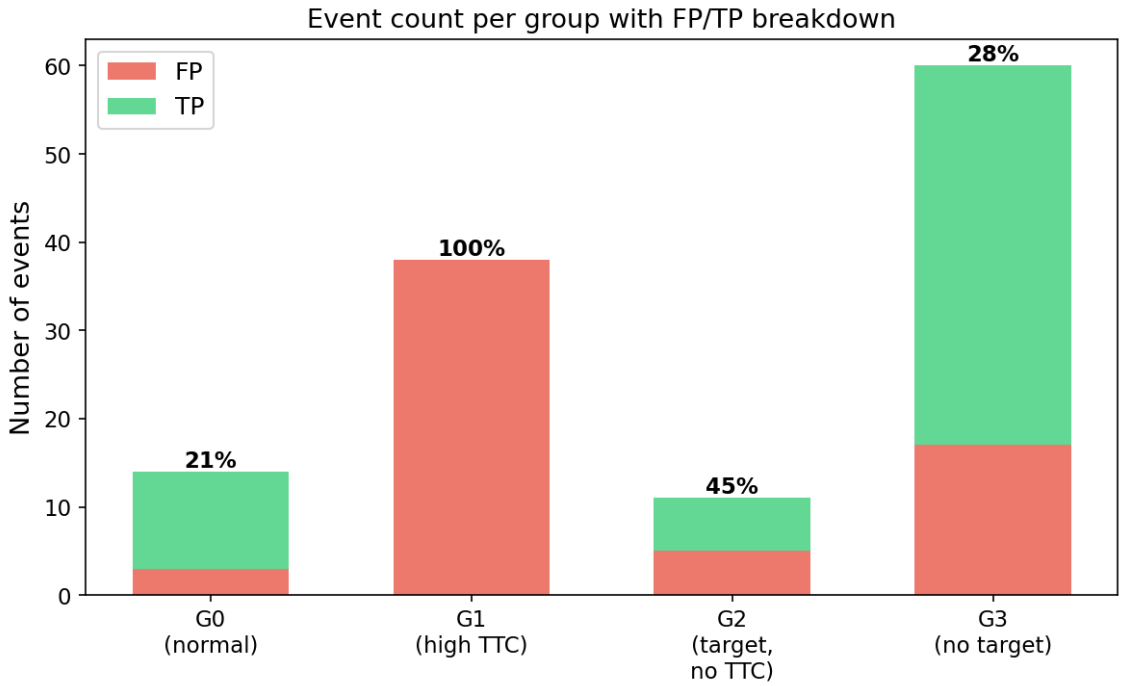

Figure 4.6: Event count per rule-based group (G0–G3) with FP/TP breakdown (speed-gate scenario). Percentage labels indicate the FP rate within each group.

Two dominant patterns emerge:

1. **G1 accounts for the majority of Condition A FP events** (38 of 63, 60.3%) with a 100% FP rate. Every event in this group had a TTC exceeding the per-event threshold, triggering Condition A. These represent activations where a target was detected but the time margin was too large to justify an emergency intervention.
2. **G3 is the largest group** (60 events, 48.8% of total) with 17 FP events (28.3% FP rate). In G3, no forward target was present; the 17 FP events were classified via Condition B (absent driver response). The 43 TP events had the driver already braking, preventing a Condition B trigger.

4.4.2 G1: High-TTC Events (Dominant FP Group)

The 38 G1 events formed the second-largest group and were all classified as FP (100%). Their smoothed TTC ranged from 1.46 s to 2.55 s (mean 1.88 s), exceeding the per-event threshold in every case, triggering Condition A. Mean forward target distance was 10.6 m, and mean speed was 25.1 km/h—similar to the fleet average—indicating that these activations occurred in normal urban driving rather than unusual speed scenarios.

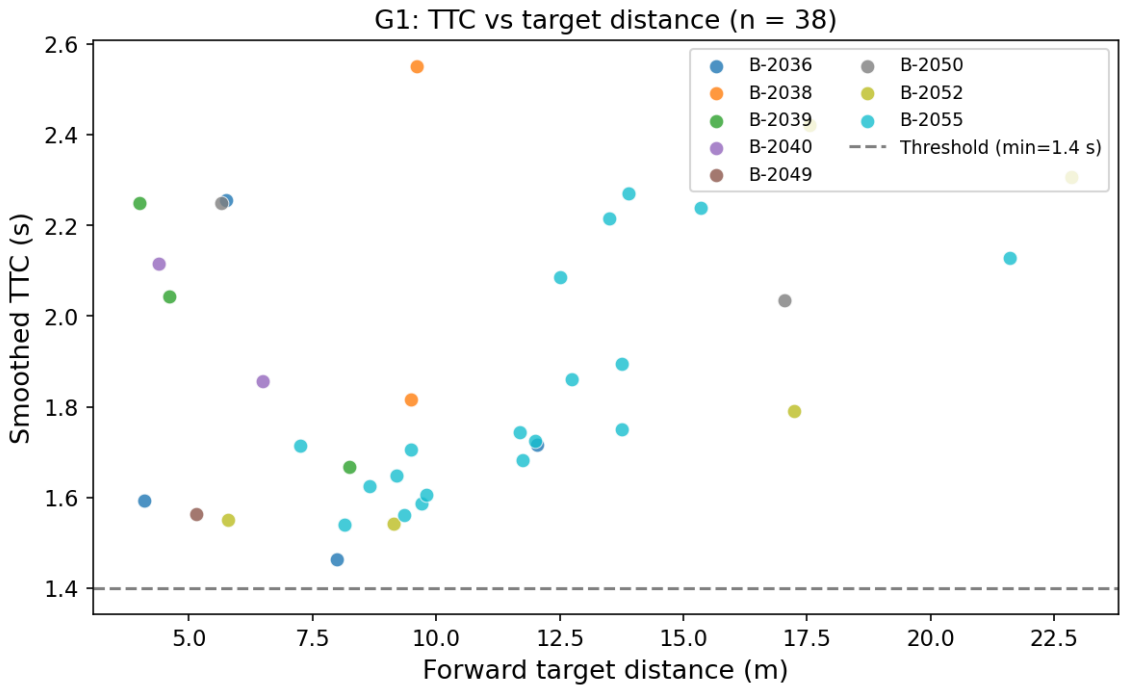

Figure 4.7: Scatter plot of smoothed TTC versus forward target distance for the 38 G1 events, coloured by vehicle. The dashed line marks the per-event threshold $\tau = \max(1.4,\ v/(2 \cdot 3.0))$.

4.4.3 G0: Normal Near-Target Activations

The 14 G0 events represent cases where a radar target was present and the TTC was at or below the per-event threshold, indicating a plausible collision risk. Eleven

events (78.6%) were classified as TP and three as FP.
The three FP events in G0 had TTC values of 1.48 s, 1.24 s, and 1.05 s—but were classified as FP via Condition B (absent driver response within 1.2 s).

4.4.4 G3: No-Target Events

The 60 G3 events had no identifiable forward radar target at the moment of AEB activation (longitudinal position reported as −179.25 m, a sentinel value indicating absence). No TTC could be computed for any of these events.
Despite the absence of a target, 43 events (71.7%) were classified as TP—the driver had applied the brake within 1.2 s (42 of 60 had `BrakeSwitch` = 1 at the anchor), and without a valid TTC, Condition A could not trigger. The 17 FP events in this group were cases where the driver did not brake within the 1.2 s window.
Steering activity was notably high in G3: 35 of 60 events (58.3%) showed steering wheel angle changes exceeding 3° within the 1.2 s post-anchor window, suggesting possible evasive manoeuvres or curve-related false detections.

4.4.5 G2: Target Present, TTC Undefined

Eleven events fell into G2, where a forward target was detected but no converging TTC could be computed (e.g. the target was moving away). Five were classified as FP (via Condition B) and six as TP. The small sample size limits further statistical analysis for this group.

4.4.6 Group Feature Comparison

Table 4.8 summarises the mean feature values across the four groups for the speed-gate scenario.

Table 4.8: Mean feature values per group (speed-gate scenario). '—' indicates insufficient data.

Feature	G0	G1	G2	G3
Speed (km/h)	25.8	25.1	18.4	22.6
Target distance (m)	8.4	10.6	9.5	—
TTC (s)	1.24	1.88	2.26	—
Long. accel. (m/s^2)	−0.72	−1.32	−1.65	−1.44
Braking duration (s)	0.43	0.77	0.56	0.61

G3 events showed slightly lower mean speed (22.6 km/h) and stronger deceleration ($-1.44\,m/s^2$) compared to G0, which may reflect more abrupt system interventions in the absence of a tracked target.

4.5 G0 Internal Analysis: FP versus TP

Group G0 contains the 14 speed-gate events where a forward radar target was present and the smoothed TTC did not exceed the per-event threshold — the most safety-relevant scenario. This section compares the three FP events against the 11 TP events to identify discriminating factors.

4.5.1 TTC and Target Geometry

Table 4.9 compares the key features between the FP and TP subsets within G0.

Table 4.9: G0 feature comparison: FP ($n = 3$) versus TP ($n = 11$), speed-gate scenario.

	FP ($n = 3$)		TP ($n = 11$)	
Feature	Mean	Range	Mean	Range
TTC (s)	1.26	[1.05, 1.48]	1.23	[0.92, 1.65]
Target distance (m)	9.0	[4.9, 13.8]	8.3	[3.4, 18.9]
Speed (km/h)	26.3	[18.0, 34.9]	25.7	[12.4, 37.9]
Long. accel. (m/s^2)	−0.59	[−1.31, 0.10]	−0.76	[−2.21, 0.12]
Braking duration (s)	0.27	[0.08, 0.66]	0.47	[0.07, 1.33]

The three FP events had TTC values (1.05–1.48 s), placing them in a region close to the threshold. All three were classified as FP solely via Condition B (absent driver response).

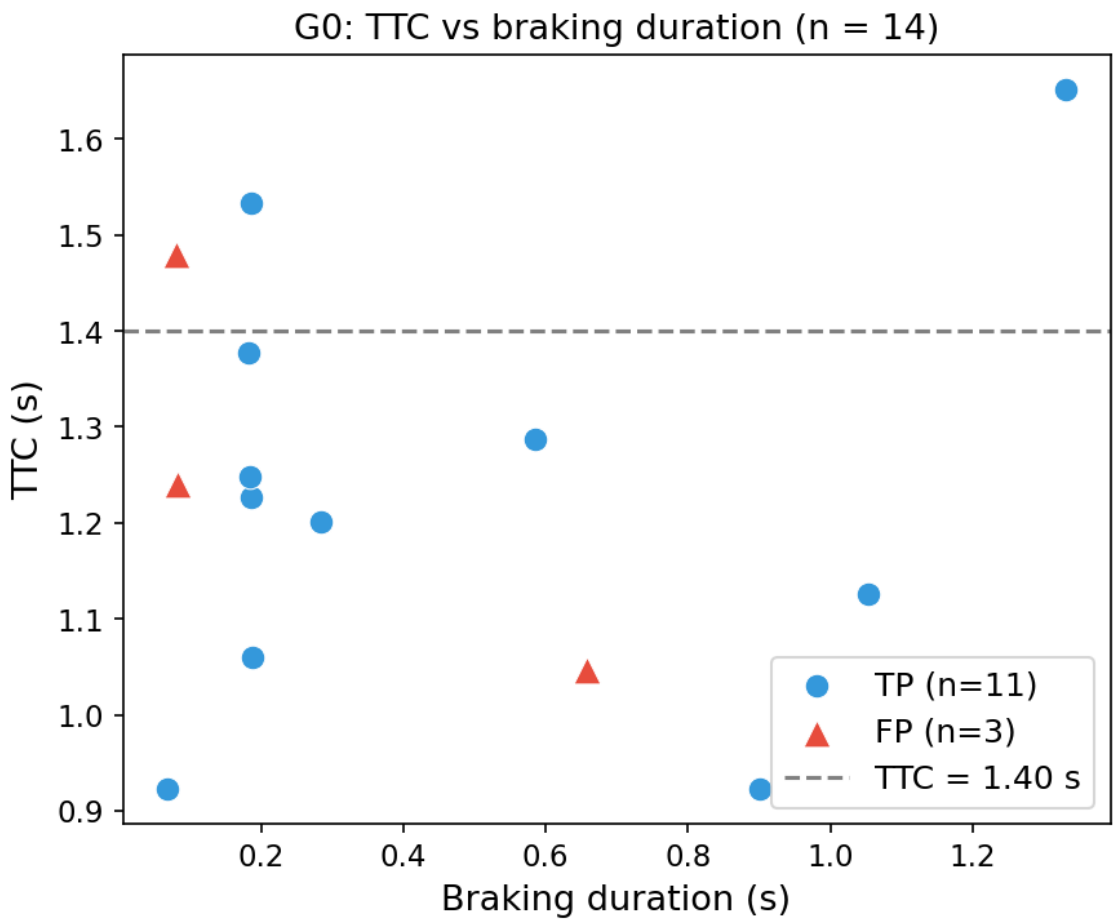

Figure 4.8: G0 events: TTC versus braking duration. Red triangles denote FP events ($n = 3$), blue circles denote TP events ($n = 11$). The horizontal dashed line marks TTC = 1.40 s.

4.5.2 Braking Behaviour

The most notable difference lies in braking behaviour. The three FP events had a mean braking duration of 0.27 s, whereas the TP events had a mean braking duration of 0.47 s. Among the 11 TP events, the majority had the brake switch already active at the AEB anchor, indicating the driver had recognised the threat independently. This pattern is consistent with the Condition B definition: in the FP events, the driver showed no braking intent within 1.2 s, while in the TP events the driver either was already braking or responded quickly.

4.5.3 Interpretation

The G0 FP events represent borderline cases: all three had a genuine forward target and a TTC close to the threshold, but the driver's lack of braking response triggered the FP classification. A possible explanation is that the driver perceived the situation as non-critical and chose not to brake. The small sample size ($n = 3$) limits the strength of these conclusions, but the pattern—FP associated with absent driver braking and shorter system braking duration—is consistent with the regulatory definition of an FP intervention.

4.6 Driver Behaviour and Attribution

This section examines what the driver was doing at the moment of AEB activation and in the 1.2 s window that followed. The attribution analysis does not change the FP/TP classification but provides context for understanding *why* events were classified as they were.

4.6.1 Brake Pedal State at AEB Activation

At the AEB anchor point, 74 of 123 events (60.2%) had the brake switch active, indicating the driver was already braking. Table 4.10 shows a strong association between brake state and classification.

Table 4.10: Brake switch state at AEB anchor versus FP/TP classification (speed-gate scenario).

	BrakeSwitch = 0	BrakeSwitch = 1	Total
FP	45	18	63
TP	4	56	60
Total	49	74	123

Among TP events, 93.3% (56/60) had the driver already braking, compared to only 28.6% (18/63) of FP events.

4.6.2 Accelerator Pedal Activity

Twenty-four events (19.5%) showed accelerator pedal activity exceeding 5% within the 1.2 s post-anchor window.

Table 4.11: Accelerator pedal activity within 1.2 s of AEB anchor (speed-gate scenario).

	Accel active	Accel inactive	Total
FP	23	40	63
TP	1	59	60
Total	24	99	123

The prevalence of accelerator activity in FP events (36.5% vs 1.7% for TP) strongly suggests that in FP cases the driver did not perceive an imminent collision threat and maintained or increased throttle input.

4.6.3 Steering Activity

Steering wheel angle changes exceeding 3° within the 1.2 s window were observed in 65 of 123 events (52.8%). Steering activity was present in both FP and TP events at broadly similar rates and therefore does not serve as a strong discriminator.

4.6.4 Driver Action Combinations

Table 4.12 presents the full combination of driver actions and their association with FP/TP classification.

Table 4.12: Driver action combinations within the 1.2 s window after AEB activation (speed-gate scenario). Brake = BrakeSwitch active at anchor; Accel = accelerator > 5%; Steer = $|\Delta\theta| > 3°$.

Brake	Accel	Steer	Events	FP	TP
No	No	No	17	16	1
No	No	Yes	8	6	2
No	Yes	No	13	13	0
No	Yes	Yes	11	10	1
Yes	No	No	28	7	21
Yes	No	Yes	46	11	35
Total			123	63	60

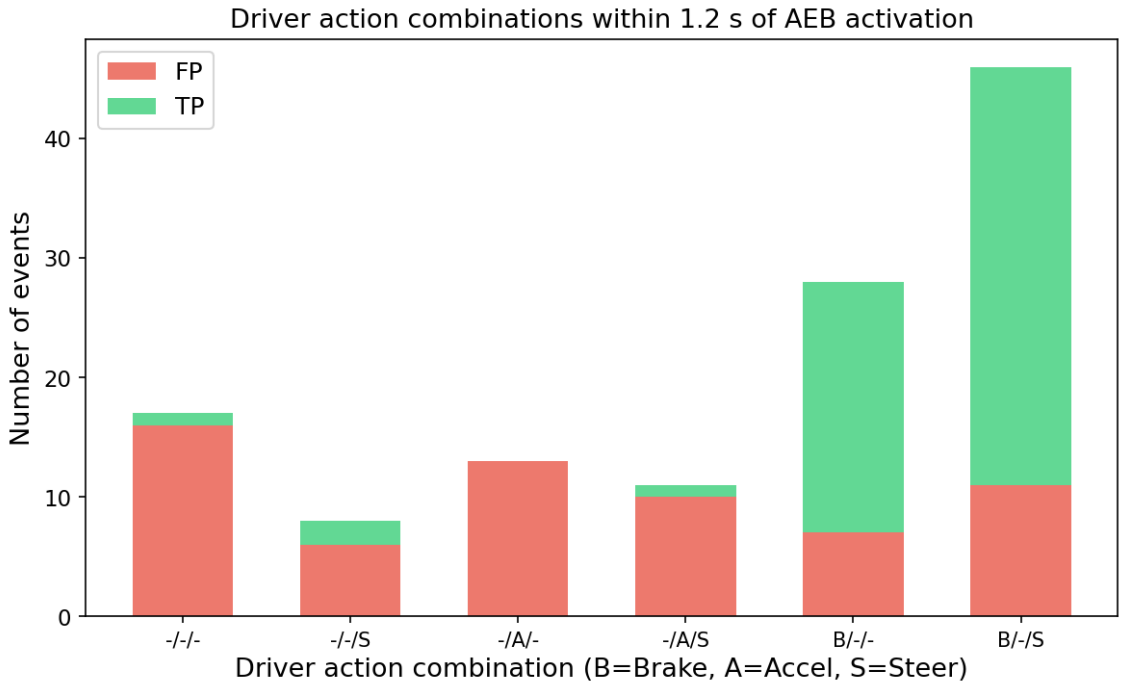

Figure 4.9: Driver action combinations within 1.2 s of AEB activation (speed-gate scenario). Each cluster corresponds to a Brake/Accel/Steer combination; red bars indicate FP events, green bars indicate TP events.

Two patterns stand out:

- **No brake + no accelerator + no steer** (17 events): FP rate of 94.1% (16/17). The driver showed no active control inputs, indicating no perceived threat.
- **Brake active + no accelerator** (74 events): predominantly TP (56/74 = 75.7%), with the driver already applying the brake pedal at the moment of AEB activation.

4.6.5 System Availability Context

All 123 events occurred with the AEB system operating at full capability:

- `CMReducedAbilityStatus` = 0
- `EmergencyBrakingInhibition` = 0
- `VDCFullyOperational` = 1

and across the board. No events were recorded under degraded system conditions, eliminating sensor degradation as a confounding factor in this dataset.

5
Discussion

This chapter interprets the results presented in Chapter 4 in the context of existing literature, discusses implications for AEB system development and fleet management, and identifies open questions that the current dataset cannot fully resolve.

5.1 FP Rate in Context

The observed FP rate of 51.2% (63 of 123 events) provides one of the first published quantitative baselines for criteria-based AEB FP classification in an urban bus fleet. Direct comparison with existing studies is difficult because most published work reports crash-avoidance effectiveness rather than FP rates. Teoh [9] found that front crash prevention systems reduced police-reportable crashes by 12% in large trucks, while Woodrooffe et al. [5] estimated societal benefits of AEB in terms of crash and fatality reduction—but neither study reports activation-level FP rates.
The closest comparable results come from naturalistic studies of heavy vehicle crash avoidance systems. Svenson et al. [16] documented that stationary-object false alerts (overpasses, signs) constituted a substantial fraction of FCW activations in trucks, and Goodall and Ohlms [17] reported that transit-bus operators frequently cited false alarms as the primary complaint with collision avoidance warning systems. Our 51.2% rate, while not directly comparable due to differing definitions and thresholds, confirms that FP activations are a material concern in real-world urban bus operation.
Condition A triggered in 66.7% of FP events, while Condition B triggered at the same rate—with 21 events satisfying both conditions simultaneously. This balanced contribution aligns with the observation by Moran et al. [3] that AEB systems on city buses tend to activate at relatively long TTC values compared to passenger cars, partly due to the longer braking distances and higher sensor mounting positions of buses.

5.2 Role of the eTTC Model

Comparing the eTTC model (which incorporates relative acceleration) against a pure constant-velocity TTC reveals a nuanced picture. Of the 63 target-present events, the eTTC yielded a valid value for 52; the remaining 11 (Group G2) had no finite eTTC because the target's deceleration eliminated the collision trajectory. For Condition A evaluation, 7 events that would have triggered the threshold exceedance

under constant-velocity TTC did *not* trigger under eTTC—the acceleration term reduced the predicted collision time below the threshold. Conversely, 12 events triggered Condition A under eTTC but not under TTC_{CV}, because the relative deceleration was insufficient to offset the larger gap.
The net effect on the overall FP rate is modest (51.2% with eTTC versus 50.4% with TTC_{CV}), but the two models redistribute events across groups differently. This finding echoes the broader literature: Minderhoud and Bovy [21] argued that accounting for acceleration is essential for realistic safety assessment, and Flannagan et al. [10] noted that the relationship between measured TTC and system activation depends on the internal prediction model. The practical consequence for AEB calibration is that threshold-based FP definitions must specify *which* TTC model is evaluated; a constant-velocity model alone cannot capture acceleration-dependent encounter dynamics that are common in urban bus traffic.

5.3 Driver Behaviour as a Discriminator

The strong association between driver braking state and FP classification (93.3% of TP events had the brake switch active versus 28.6% of FP events) has practical implications. It suggests that a future refinement of the FP definition could incorporate driver state as a weighting factor rather than a binary condition, similar to the adaptive approach proposed by Shao et al. [15] for hazmat truck FCW systems, where driver behaviour and risk levels jointly determine warning thresholds.
Schindler et al. [12] found that truck drivers were already braking before full mitigation in the vast majority of AEB system events, suggesting that anticipatory braking reflects driver awareness of the forward threat. Our results corroborate this: when the driver was already braking at the AEB anchor without accelerator activity, the activation was classified as TP in 75.3% of cases (55/73 events in the brake-active, no-accelerator group). Conversely, the "no brake + accelerator active" combination yielded a 95.8% FP rate (23/24 events).

5.4 Passenger Safety Implications

Beyond the engineering classification of FP versus TP, each AEB activation—whether justified or not—subjects standing passengers to sudden deceleration. Albertsson et al. [13] documented that bus passengers sustain injuries primarily from falls and impacts during braking events, even at low speeds. Our dataset shows that 64 of 123 events occurred at speeds between 20–35 km/h, a range where standing passengers are routinely present.
This reinforces the argument by Moran et al. [3] that AEB requirements for city buses must balance collision mitigation against the risk of braking-induced passenger injuries. The 51.2% FP rate implies that approximately one in every two AEB activations in this fleet may have been unnecessary from a collision avoidance standpoint, potentially causing passenger discomfort or injury without a corresponding safety benefit.

5.5 G1 Events and System Calibration

Group G1 accounts for 60.3% of FP events (38 of 63) and represents the primary opportunity for AEB calibration improvement. These are events where a radar target was present but the smoothed eTTC exceeded the per-event threshold—indicating the system activated at a point where the collision risk, by the specification definition, was not yet imminent. Notably, every G1 event was classified as FP (100% FP rate), making this group the single most reliable indicator of unnecessary activation. Woodrooffe et al. [11] evaluated a commercial forward collision avoidance system on trucks and identified that the warning/braking threshold calibration is a critical trade-off parameter: too aggressive a threshold increases FP activations; too conservative a threshold reduces crash mitigation benefit. Our G1 results quantify one side of this trade-off: the 100% FP rate within G1 demonstrates that all activations with TTC above the threshold were unnecessary interventions under the adopted specification.
Jermakian [4] estimated that large-truck technologies could prevent 29% of rear-end crashes if calibrated appropriately. The implicit assumption is that systems are calibrated to avoid excessive false activations that would lead to driver override or system disablement. Our G1 analysis provides empirical data to inform such calibration decisions for urban bus AEB systems.

5.6 Generalisability and Fleet Context

The dataset covers 20 vehicles of a single platform operating in one city. Woodrooffe [14] highlights that large-vehicle safety challenges are highly dependent on operational context (route type, traffic density, vulnerable road user exposure), and Sweatman et al. [6] emphasise that voluntary AEB adoption on commercial vehicles requires confidence in low FP rates across diverse operating conditions.
The concentration of FP events in vehicle B-2055 (73.0% FP rate, 27 of 37 events) raises the question of whether vehicle-specific factors (sensor alignment, software version, route assignment) contribute to FP susceptibility. This vehicle-level variability cannot be fully explained by the available CAN data and warrants targeted investigation, possibly including sensor health diagnostics and route-level analysis.

5.7 Limitations of the Criteria-Based Approach

The Condition A / Condition B framework provides a traceable and reproducible classification, but it encodes specific assumptions and is subject to several analytical constraints.

- Condition A assumes that TTC above the threshold implies "no imminent collision risk". In reality, a high TTC does not preclude a collision—for instance, if the target subsequently brakes hard. The FP label is therefore a retrospective judgment based on the state at the activation instant.
- Condition B assumes that driver non-response within 1.2 s indicates the driver did not perceive a threat. As Schindler et al. [12] discuss that drivers may

intentionally rely on the AEB system (cooperative driving), which would produce a Condition B trigger despite the activation being operationally appropriate.

These ambiguities are inherent to any criteria-based offline analysis performed without video confirmation. The framework should be viewed as a *screening tool* that identifies candidate FP events for subsequent expert review, rather than a definitive ground-truth label.

Beyond the criterion assumptions above, three additional analytical constraints affect confidence in the reported results.

First, the pipeline's eTTC estimate was compared against an internal AEBS ground-truth TTC for a single track-test event, showing close agreement. However, systematic validation across all 123 fleet events has not been performed. The internal AEBS TTC incorporates proprietary sensor-to-function delays and safety margins that are not fully accessible; therefore, the degree to which Condition A labels might shift under an alternative TTC reference remains unquantified. The conservative threshold design ($\max(1.4\,\mathrm{s},\ v_{\mathrm{ego}}/2a_{\mathrm{LPoB}})$) provides a partial mitigation, but formal confidence bounds on the FP boundary require multi-event validation.

Second, with 123 events distributed across four groups (G0–G3), per-group cell sizes range from 11 (G2) to 38 (G1), limiting the reliability of within-group distributional comparisons. Per-vehicle FP rates are based on as few as 1–4 events for several vehicles, making vehicle-level conclusions indicative rather than statistically robust. Continued data collection through the pipeline's incremental processing mode is expected to improve statistical power without requiring methodological changes.

Third, only one video-embedded MF4 file was available; the corresponding event was classified as TP and visually confirmed. No FP event could be independently verified through video. The CAN-based labels should therefore be interpreted as candidate classifications subject to expert review, not as confirmed ground truth.

6
Conclusion

This thesis developed and demonstrated an end-to-end offline analysis pipeline for identifying candidate FP AEB braking events from urban bus fleet recordings. The pipeline ingests raw MF4 data from a 20-vehicle fleet operating on regular urban routes, applies automated event detection, TTC computation using the eTTC model, and a traceable two-condition FP classification. A total of 3 160 CAN recordings were processed with a 99.9% success rate, yielding 123 qualified AEB braking events. The main findings are summarised below, organised by the research questions.

6.1 Summary of Findings

RQ1 — FP rate. Under the adopted Condition A / Condition B definition with the eTTC model as the primary metric, the overall FP rate across 123 speed-gate-qualified AEB events is **51.2%** (63 of 123). Activations are concentrated in the 20–35 km/h range (mean 22.5 km/h), consistent with low-speed urban bus operations.

RQ2 — Distribution across rule-based groups. Group G1 (target present, TTC above threshold) dominates the FP landscape, containing 60.3% of all FP events (38 of 63) with a within-group FP rate of 100%. Condition A (excessive TTC) triggered in 66.7% of FP cases, and Condition B (absent driver response) triggered at the same rate. The role of the eTTC model in redistributing events across groups relative to a constant-velocity model is discussed in Section 5.2.

RQ3 — Driver behaviour signals and attribution. Driver behaviour signals discriminate strongly between FP and TP events: brake-switch activity at the AEB anchor is present in 93.3% of TP events versus 28.6% of FP events, and the "no brake + accelerator active" action combination yields a 95.8% FP rate compared to 24.7% for "brake active + no accelerator." The implications of these patterns for adaptive FP definitions are discussed in Section 5.3.

6.2 Contributions

The main academic and methodological contributions of this work are:

1. A **reproducible, scalable offline pipeline** that ingests raw MF4 fleet recordings from Amazon S3 and produces analysis-ready event catalogues with trace-

able FP/TP labels. The pipeline processed 3 160 files from 20 vehicles with 99.9% decode success rate, supports incremental resume for interrupted runs, and produces per-file intermediate outputs for audit.

2. A **structured event characterisation framework** (scenario buckets) that associates each qualified event with ego kinematics, target state, driver actions, perception quality indicators, and system availability context—enabling downstream grouping and root-cause analysis without re-processing raw data.
3. A **rule-based grouping methodology** (G0–G3) that segments events by target presence and TTC-to-threshold relationship, revealing that 60.3% of FP events concentrate in a single group (G1: high TTC with target present, 100% FP rate) and enabling targeted engineering investigation.
4. A **quantitative driver behaviour attribution** linking brake pedal state, accelerator activity, and steering input to FP/TP classification, demonstrating that driver action combinations are strong correlates of the FP outcome.

In addition to these academic and methodological contributions, the work delivers the following engineering value to the organisation, within the scope of the analysed fleet:

1. **Auditable event triage:** the pipeline converts the open-ended question "why did the bus brake?" into an auditable, specification-traceable FP/TP label for every qualified event, substantially reducing the manual effort previously required to triage individual recordings.
2. **Actionable calibration evidence:** the G1 finding for this fleet gives the AEB calibration team a quantified, prioritised target for investigating threshold behaviour, linking the offline analysis to concrete calibration questions.
3. **A reusable analysis framework, not transferable findings:** what transfers to other deployments is the methodology and software architecture (the reusable decode–detect–classify–group workflow), and not the numerical results or tuned parameters. The reported FP rates, group proportions, and the smoothing and brake-response calibrations are specific to the present Volvo bus platform, the London operating context, and the current firmware; applying the pipeline to a different sensor suite, vehicle model, or operating region would require recalibration of these parameters and independent revalidation against the new data before any result could be considered representative.

6.3 Limitations

The scope delimitations stated in Section 1.5 apply throughout. The main analytical limitations encountered during execution are discussed in Section 5.7 and summarised here:

- **Limited visual verification** — only one event had video available; FP labels are CAN-based candidate classifications pending expert review.
- **Condition B ambiguity** — driver non-response cannot disambiguate inattention from cooperative AEB reliance without additional context.
- **Preliminary TTC validation** — pipeline eTTC was verified against one track-test reference; systematic multi-event validation has not been performed.

- **Small sample size** — 123 events (20 vehicles) limit per-group and per-vehicle statistical power.

6.4 Future Work

- **Video annotation and FP confirmation:** As more video-capable MF4 recordings become available, use synchronised front-facing video to manually verify FP classifications. This will enable disambiguation of Condition B events into "driver inattention" (TP) versus "unnecessary activation" (confirmed FP), and provide ground truth for refining the Condition A threshold.
- **Ground-truth TTC validation:** Once GT TTC time-series data from the GTT team is available, validate the eTTC and TTC computations and establish confidence bounds on the FP classification boundary.
- **Root-cause clustering:** Extend the G0–G3 grouping with perception-level attributes (object class, track lifetime, existence confidence, dynamic classification) to enable automated FP root-cause attribution. Apply unsupervised clustering on the scenario bucket features to identify recurrent FP archetypes (stationary infrastructure, cut-in vehicles, crossing pedestrians at bus stops).
- **Confidence-scored FP labelling:** Replace the current binary FP/TP decision with a continuous FP confidence score that combines the margin by which the measured TTC exceeds (or falls short of) the per-event threshold, the radar existence confidence (`ExistConf`), the track-quality and slot-confidence indicators already produced by the conditioning and fusion stages, and the driver-response signals. Such a score would rank events by labelling certainty, allowing borderline cases (e.g. Group G0/G2 events near the threshold) to be flagged for prioritised expert review rather than forced into a hard class.
- **Probabilistic FP classification:** The current classification applies deterministic thresholds and therefore discards the uncertainty information already available in the pipeline—radar existence confidence, fusion slot confidence, the Mahalanobis gating statistics, and the Kalman/RTS state covariances. A natural extension is to propagate these quantities into the decision stage and estimate a posterior FP probability within a Bayesian or calibrated statistical framework (e.g. logistic regression with probability calibration). This would express each label as a graded risk rather than a hard class and, once video ground truth becomes available, the posterior could be probabilistically calibrated and evaluated with Receiver Operating Characteristic (ROC)/precision–recall analysis.
- **Threshold sensitivity analysis:** Systematically vary the Condition A TTC threshold and Condition B time window to produce ROC-style curves, quantifying the trade-off between FP deetection sensitivity and TP retention rate.
- **Longitudinal monitoring:** Deploy the pipeline in incremental mode on newly arriving recordings to track FP rate trends across AEB software updates, seasonal traffic pattern changes, and route modifications—supporting data-driven system improvement.

Bibliography

[1] European New Car Assessment Programme, "Autonomous emergency braking (AEB) systems," Available at: https://www.euroncap.com, 2023.

[2] National Highway Traffic Safety Administration, "Automatic emergency braking research report," U.S. Department of Transportation, 2022.

[3] D. Moran, R. Peck, T. Jiang, and B. Bain, "Assessing the case for requiring AEB on city buses and developing technical requirements and test procedures," in *Proceedings of the 26th International Technical Conference on the Enhanced Safety of Vehicles (ESV)*, no. 26ESV-000315. NHTSA, 2019. [Online]. Available: https://static.nhtsa.gov/esv/pdf/ESV/Proceedings/26/26ESV-000315.pdf

[4] J. S. Jermakian, "Crash avoidance potential of four large truck technologies," *Accident Analysis & Prevention*, vol. 49, pp. 338–346, 2012.

[5] J. Woodrooffe, D. Blower, S. Bao, and S. Bogard, "Societal benefit of automatic emergency braking and lane departure warning systems in large trucks," in *Proceedings of the 26th International Technical Conference on the Enhanced Safety of Vehicles (ESV)*, no. 26ESV-000167. NHTSA, 2019. [Online]. Available: https://static.nhtsa.gov/esv/pdf/ESV/Proceedings/26/26ESV-000167.pdf

[6] P. Sweatman, J. Sayer, D. Blower, and S. Bao, "Research and testing to accelerate voluntary adoption of automatic emergency braking (AEB) on commercial vehicles," University of Michigan Transportation Research Institute, Tech. Rep., 2020. [Online]. Available: https://rosap.ntl.bts.gov/view/dot/49335

[7] J. Shaikh and N. Lubbe, "Evaluating forward collision warning and autonomous emergency braking systems in India using dashboard cameras," SAE International, Tech. Rep. 2024-26-0031, 2024.

[8] D. Liu, W. Huang, R. Chu, Z. Li, X. Jin, H. Zhang, Y. Wang, and S. Ji, "Study on performance testing and evaluation of autonomous emergency braking system based on self-constructed comprehensive performance evaluation index model," *Sensors*, vol. 25, no. 7, p. 2171, 2025.

[9] E. R. Teoh, "Effectiveness of front crash prevention systems in reducing large truck real-world crash rates," *Traffic Injury Prevention*, vol. 22, no. 4, pp. 284–289, 2021, published online: 26 Mar 2021. [Online]. Available: https://doi.org/10.1080/15389588.2021.1893700

[10] C. A. Flannagan, D. J. LeBlanc, R. J. Kiefer, S. E. Bogard, A. Leslie, C. T. Zagorski, and C. S. Beck, "Field study of light-vehicle crash avoidance systems: Automatic emergency braking and dynamic brake support," National

Highway Traffic Safety Administration, Tech. Rep. DOT HS 812 615, 2018. [Online]. Available: https://www.nhtsa.gov/sites/nhtsa.gov/files/documents/812_615_fldstdy-lght-vhcle-crs-avoid-sys-aeb-dbs.pdf

[11] J. Woodrooffe, D. Blower, C. A. Flannagan, S. E. Bogard, and S. Bao, "Effectiveness of a current commercial vehicle forward collision avoidance and mitigation systems," in *SAE Technical Paper*, no. 2013-01-2394. SAE International, 2013.

[12] R. Schindler, G. Bianchi Piccinini, and L. Decoster, "An investigation of truck drivers' behaviour before and during real-world advanced emergency braking system interventions," *IATSS Research*, vol. 49, pp. 418–424, 2025.

[13] P. Albertsson, T. Falkmer, and U. Björnstig, "Banging heads onboard buses: Assessment scheme to improve injury mitigation for bus passengers," *Traffic Injury Prevention*, vol. 20, no. 2, pp. 197–202, 2019.

[14] J. Woodrooffe, "Critical issues for large truck safety," in *International Encyclopedia of Transportation*. Elsevier, 2021, pp. 424–431.

[15] Y. Shao, X. Shi, Y. Zhang, Y. Zhang, Y. Xu, W. Chen, and Z. Ye, "Adaptive forward collision warning system for hazmat truck drivers: Considering differential driving behavior and risk levels," *Accident Analysis & Prevention*, vol. 191, p. 107221, 2023.

[16] A. L. Svenson, K. Grove, J. Atwood, and M. Blanco, "Using naturalistic data to evaluate heavy vehicle crash avoidance systems performance," in *Proceedings of the 25th International Technical Conference on the Enhanced Safety of Vehicles (ESV)*, 2017, paper No. 17-0231. [Online]. Available: https://trid.trb.org/View/1485114

[17] N. J. Goodall and P. B. Ohlms, "Evaluation of a transit bus collision avoidance warning system in Virginia," Virginia Transportation Research Council, Final Report FHWA/VTRC 22-R18, Feb. 2022, accessed: 2026-03-27. [Online]. Available: http://www.virginiadot.org/vtrc/main/online_reports/pdf/22-r18.pdf

[18] UDRIVE Consortium, "UDRIVE naturalistic driving study dataset," Available at: https://www.udrive.eu, 2018.

[19] Transportation Research Board, "SHRP 2 naturalistic driving study dataset," National Academy of Sciences, 2015.

[20] euroFOT Consortium, "Final results of the euroFOT project," European Commission, Tech. Rep., 2012.

[21] M. M. Minderhoud and P. H. Bovy, "Extended time-to-collision measures for road traffic safety assessment," *Accident Analysis & Prevention*, vol. 33, no. 1, pp. 89–97, 2001. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0001457500000191

[22] E. Tomasch and S. Smit, "Naturalistic driving study on the impact of an aftermarket blind spot monitoring system on the driver's behaviour of heavy goods vehicles and buses on reducing conflicts with pedestrians and cyclists," *Accident Analysis & Prevention*, vol. 192, p. 107242, 2023.

[23] S. Riedmaier, T. Ponn, D. Ludwig, B. Schick, and F. Diermeyer, "Survey on scenario-based safety assessment of automated vehicles," *IEEE Access*, vol. 8, pp. 87 456–87 477, 2020.

[24] J. Bärgman, C.-N. Boda, and M. Dozza, "Counterfactual simulations applied to SHRP2 crashes: The effect of driver behavior models on safety benefit estimations of intelligent safety systems," *Accident Analysis & Prevention*, vol. 102, pp. 165–180, 2017.

[25] N. Zhang, Y. Wu, A. Yang, and T. Liu, "Development of a bus real-time crash risk prediction framework by using a self-attention-based bidirectional long and short-term memory network with anomaly detection learning and mixed sequence features," *Accident Analysis & Prevention*, vol. 225, p. 108306, 2026.

[26] E. Agrell, "The structure of a technical report," Technical report, Tech. Rep., 2010, rev. 4.2, December.

[27] H. E. Rauch, F. Tung, and C. T. Striebel, "Maximum likelihood estimates of linear dynamic systems," *AIAA Journal*, vol. 3, no. 8, pp. 1445–1450, 1965. [Online]. Available: https://doi.org/10.2514/3.3166

[28] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 03 1960. [Online]. Available: https://doi.org/10.1115/1.3662552

[29] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.

[30] R. Rajamani, *Vehicle dynamics and control*. Springer, 2006.

[31] United Nations Economic Commission for Europe, "No. 131 (advanced emergency braking system)," 2022, UNECE document ECE/TRANS/WP.29/2022/76, submitted by the Working Party on Automated/Autonomous and Connected Vehicles, April 11, 2022. [Online]. Available: https://unece.org/sites/default/files/2022-08/ECE_TRANS_WP.29_2022_76E.pdf

[32] International Organization for Standardization, "Road vehicles – safety of the intended functionality," 2022, iSO 21448:2022. [Online]. Available: https://www.iso.org/standard/77490.html

A

CAN Signal Catalogue

This section provides the complete list of 54 active CAN signals used by the false-positive detection pipeline, grouped by functional category. For each signal the CAN channel, unit, and role in the pipeline are given. Signals marked with an asterisk (*) serve as fallback alternatives when the primary signal is unavailable in a particular MF4 recording.
Table A.1 summarises the six signal categories used by the detection pipeline. The complete per-signal listing follows in the subsections below.

Table A.1: Signal categories and their roles in the detection pipeline.

Category	Signal Group	Role
Core detection	`CM_Status`, ego acceleration, ego speed, `BrakeSwitch`	AEB activation detection and qualified brake event identification
Driver intent	Steering angle, yaw rate, accelerator pedal, turn signal	Driver avoidance/response assessment
Brake authenticity	`EBSBrakeSwitch`, `FoundationBrakeUse`	Verification of physical brake engagement
Radar perception	`FLRObj`*`XX`*`_LongPos`, `_LongVel`, `_LatPos`, `_ObjClass`, `_ExistConf`	TTC computation and target characterisation (FLR)
Camera perception	`FLCObject`*`XX`* position, velocity, classification	FLC object tracks; fused with FLR targets for cross-validation
System availability	`CMReducedAbilityStatus`, `EmergencyBraking-Inhibition`, `VDCFullyOperational`	System state and suppression context

A.1 Core Detection Signals

#	Signal Name	Ch.	Unit	Pipeline Role
1	CM_Status	3	enum	AEB activation flag; primary event segmentation trigger
2	LongitudinalAcceleration	2	m/s^2	Braking threshold verification ($a \leq -1.5\,m/s^2$)
3	EgoMotionXYplane_LongAcc*	4	m/s^2	Fused acceleration fallback
4	VehicleSpeed	2	km/h	Speed-range gating (speed gate > 10 km/h)
5	WheelBased-VehicleSpeed_BB1_X_V*	3	km/h	Redundant wheel-speed source
6	EgoMotion-XYplane_LongSpeed*	4	m/s	Fused speed fallback
7	BrakeSwitch	2	bin	Digital brake-pedal confirmation
8	BrakePedalPosition	3	%	Analog pedal travel for Condition B
9	XBR_A_B_External-AccelerationDemd	2	m/s^2	AEB brake demand reference
10	TTC	3	s	Pre-computed TTC (if available in log)

A.2 Driver Intent Signals

#	Signal Name	Ch.	Unit	Pipeline Role
11	AccelPedalPos1_BB2	3	%	Accelerator state at AEB anchor
12	AccelPedal-KickdownSwitch_BB2	3	bin	Kick-down detection (driver override)
13	SteeringWheelAngle1	3	rad	Swerve detection ($\lvert\Delta\theta\rvert > 3°$)
14	YawRate	2	rad/s	Vehicle path-change confirmation
15	LateralAcceleration	2	m/s^2	Lateral dynamics during AEB

A.3 Brake Authenticity Signals

#	Signal Name	Ch.	Unit	Pipeline Role
16	EBSBrakeSwitch	2	bin	EBS-level brake confirmation
17	FoundationBrakeUse	2	bin	Service brake engagement flag
18	BrkApplPressHiRaRear-Axle1R	2	bar	Brake pressure — rear axle 1 right

#	Signal Name	Ch.	Unit	Pipeline Role
19	BrkApplPressHiRaRear-Axle1L	2	bar	Brake pressure — rear axle 1 left
20	BrkApplPressHiRaRear-Axle2L	2	bar	Brake pressure — rear axle 2 left
21	BrkApplPressHiRaRear-Axle2R	2	bar	Brake pressure — rear axle 2 right
22	BrkApplPressHiRaRear-Axle3L	2	bar	Brake pressure — rear axle 3 left
23	BrkApplPressHiRaRear-Axle3R	2	bar	Brake pressure — rear axle 3 right

A.4 Radar Perception Signals (per object slot)

Radar signals follow the naming pattern `FLRObj`*`XX`*`_<field>` where *XX* ranges from `00` to `23` (24 simultaneous tracks). The table below lists the fields extracted per slot.

#	Field Suffix	Unit	Pipeline Role
24	LongPos	m	Longitudinal distance to ego (TTC numerator)
25	LongVel	m/s	Relative longitudinal velocity (TTC denominator)
26	LongAcc	m/s^2	Relative longitudinal acceleration (eTTC)
27	LatPos	m	Lateral offset for in-path check
28	ObjClass	enum	Object type (2=Car, 5=Pedestrian, 10=Bicycle, ...)
29	ObjDynClass	enum	Dynamic state (1=Stationary, 2=Movable, 3=Error)
30	TrackStatus	enum	Track validity (0=Tracked, 1=Predicted, 2=Invalid)
31	ExistConf	$[0, 1]$	Existence confidence score

A.5 System Availability Signals

#	Signal Name	Ch.	Unit	Pipeline Role
32	CMReducedAbilityStatus	3	enum	CM system degradation flag
33	EmergencyBrakingInhibition	3	bin	AEB suppression indicator

#	Signal Name	Ch.	Unit	Pipeline Role
34	VDCFullyOperational	3	bin	Vehicle Dynamics Control status

A.6 Auxiliary and Context Signals

#	Signal Name	Ch.	Unit	Pipeline Role
35	FrontCollisionDriverAlrt_BB1_X_D	3	bin	FCW alert phase marker
36	DriverVigilanceLevel_BB1_X_D	3	enum	Driver vigilance state
37	DrowsyDriverWarning_BB1_X_D	3	bin	Drowsiness warning active
38	TotalVehicleDistanceHighRes	2	m	Odometer for session identification
39	AmbientAirTemperature	2	°C	Environmental context
40	VehicleMode	3	enum	Vehicle operating mode
41	VIN_DACU_Byte00 – Byte16	3	byte	Vehicle identification (17 bytes)

B

Pipeline Configuration Details

This section documents implementation-level configuration details that supplement the methodology described in the main text. Parameters already presented in Chapter 2 (e.g. RTS smoother tuning, QBE thresholds, FP classification conditions, event grouping rules) are not repeated here.

B.1 CAN Database File Mapping

Table 2.2 in the main text lists abbreviated DBC filenames. The full filenames used by the pipeline are given below for reproducibility.

Table B.1: Full DBC filenames used for CAN signal decoding.

Channel	DBC File (full name)
2	`J1939_1 - B2_Q0_1_5 (1) 2.dbc`
3	`DBC_TEA2BP_Master_Backbone_2.dbc`
4	`VicinityNet2-T2_1.27.06.dbc`
5	`VicinityNet1-T2_1.27.06.dbc`

B.2 Event Segmentation Parameters

Table B.2: Parameters governing CM_Status event detection and merging.

Parameter	Value	Description
AEB active states	$\{2, 3\}$	CM_Status values indicating active braking
Merge gap Δt_{merge}	1.0 s	Maximum inter-event gap before merging
Analysis window (batch)	5.0 s	Pre/post context extracted around each event
Analysis window (interactive)	10.0 s	Extended context for manual inspection

B.3 Speed Unit Detection Heuristic

The pipeline automatically detects whether the primary speed signal is reported in m/s or km/h. The 95th percentile of the speed column is computed; if it is below 70, the values are assumed to be in m/s and are multiplied by 3.6. Otherwise they are taken as km/h directly.

B.4 Radar Quality Gating Criteria

Table B.3: Quality gates applied to each radar object slot before inclusion in the fusion step.

Gate	Criterion
Track validity	TrackStatus = 0 (Tracked)
Existence confidence	ExistConf > 0.5
Slot-switch filter	Reject frames where slot index changes mid-track
Coast-period filter	Suppress predicted-only frames between valid updates

B.5 Slot-Switch Detection and Fill Rules

At the per-slot conditioning stage (Section 3.2.1.2), a reassignment (slot switch) is declared when any of the criteria in Table B.4 is met between consecutive frames.

Table B.4: Slot-switch detection criteria applied between consecutive frames within a single radar object slot.

Criterion	Condition	Rationale
Position jump	$\lvert\Delta$`LongPos`$\rvert > 5\,\text{m}$ or $\lvert\Delta$`LatPos`$\rvert > 2\,\text{m}$	Kinematic discontinuity
Object ID change	`ObjID` changes (excl. idle value 63)	Sensor re-identification
Lifetime reset	`LifeTime` counter resets	New track initialised
Validity edge	`slot_valid` transitions $0 \rightarrow 1$	Slot re-activation

After switch detection, missing values are filled according to Table B.5.

Table B.5: Per-signal NaN fill strategy applied after hard/soft gating.

Signal category	Signals	Fill strategy
Continuous kinematic	`LongPos`, `LongVel`, `LongAcc`	Segment-bounded forward fill (interrupted at slot switches)
Lateral position	`LatPos`	Forward fill; remaining NaN → sentinel 5.0 m
Categorical	`ObjClass`, `ObjDynClass`	Set to −1 (out-of-range sentinel)

B.6 Signal Fallback and Channel-Retry Logic

Table B.6 lists the candidate signal names for each required category, ordered by selection priority.

Table B.6: Signal categories and CAN candidate names (in priority order).

Category	Candidates
AEB status	`CM_Status`, `CM_Status_BB1_X_I`
Longitudinal acceleration	`LongitudinalAcceleration`, `EgoMotionXYplane_LongAcc`
Vehicle speed	`VehicleSpeed`, `WheelBasedVehicleSpeed_BB1_X_V`, `EgoMotionXYplane_LongSpeed`, `GPSData_VehSpeed`
Brake switch	`BrakeSwitch`, `BrakePedalSwitch_BB1_X_V`, `EBSBrakeSwitch`

If decoding yields fewer than three signal categories on the default channel mapping, the pipeline retries with all CAN channel numbers shifted by −1, accommodating MF4 files with offset bus numbering.

C

Track Test Supplementary Details

The track test scenario overview is given in Table 3.9 in the main text. This section provides additional details on test targets and the MF4 file naming convention.

C.1 Target Specifications

Table C.1: VRU and vehicle targets used in track tests.

Target Type	Standard	Description
Adult pedestrian	Euro NCAP VRU target	Articulated adult pedestrian dummy
Child pedestrian	Euro NCAP VRU target	Articulated child pedestrian dummy
Bicycle + rider	Euro NCAP VRU target	Bicycle with articulated rider dummy
Stationary car	Euro NCAP GVT	Global Vehicle Target (soft car)

C.2 Naming Convention for MF4 Files

Each MF4 recording follows the naming pattern:

`TT_GOT_VBC_FD_V2_MS_1_SREC_1_shot`*`NNNN`*`_`*`YYYYMMDD`*`_`*`HHMMSS`*`[_CAN].mf4`

where:

- `TT_GOT` — Track Test, Gothenburg
- `VBC_FD_V2` — Volvo Bus City, Front Detection, Version 2
- `shot`*`NNNN`* — Run sequence number
- `_CAN` suffix — CAN-only file (no synchronised video); 50–150 MB. Without the suffix the file includes synchronised video and is substantially larger.

C.3 Brake Calibration Filtering and Detailed Fit Comparison

Calibration set selection. Three exclusion criteria were applied before fitting the cascade correction model (Eq. 3.30 in the main text). First, scenarios classified as FCW-only (BBLA25), bus-stop manoeuvres, and false-positive aborted crossings

were removed because they do not produce a full AEB braking intervention. Second, events where the real deceleration crossed the $-1.5\,\mathrm{m/s^2}$ threshold before the simulated signal (onset delay ≤ 0) were excluded, as this indicates driver-initiated braking preceding the AEB command. Third, events in which the driver's brake pedal signal was active were excluded to avoid confounding driver and AEB deceleration. After filtering, seven events remained in the calibration set.

Onset delay. The onset delay $\Delta t_{\mathrm{onset}}$ is defined as the time difference between the instants at which the simulated and real deceleration signals first cross $-1.5\,\mathrm{m/s^2}$. Across the seven calibration events, the median onset delay is 310 ms (range 240–377 ms, interquartile range (IQR) [276, 327] ms), indicating that the physical brake system consistently lags the commanded deceleration by approximately 0.3 s.

Peak gain ratio. The peak gain ratio $G = |a_{\mathrm{real,peak}}|/|a_{\mathrm{sim,peak}}|$ quantifies the amplitude discrepancy. The median ratio is 0.78 (range 0.59–1.29), reflecting that the real brake system generally undershoots the commanded deceleration at low-to-moderate speeds but can exceed it at higher speeds where the vehicle's kinetic energy engages the hydraulic system more aggressively.

Optimisation procedure. The four parameters were jointly optimised by minimising RMSE in the *onset window*—defined from the first sample where $a_{\mathrm{real}} \leq -0.1\,\mathrm{m/s^2}$ to 150% of the span until $a_{\mathrm{real}} \leq -1.5\,\mathrm{m/s^2}$. A multi-start L-BFGS-B solver with six diverse initial guesses was used to avoid local minima.

D

Brake Simulation Calibration Figures

This section presents the complete set of brake simulation calibration figures referenced in Section 3.6.2.4.

D.1 Simulated Brake Request vs Real Deceleration

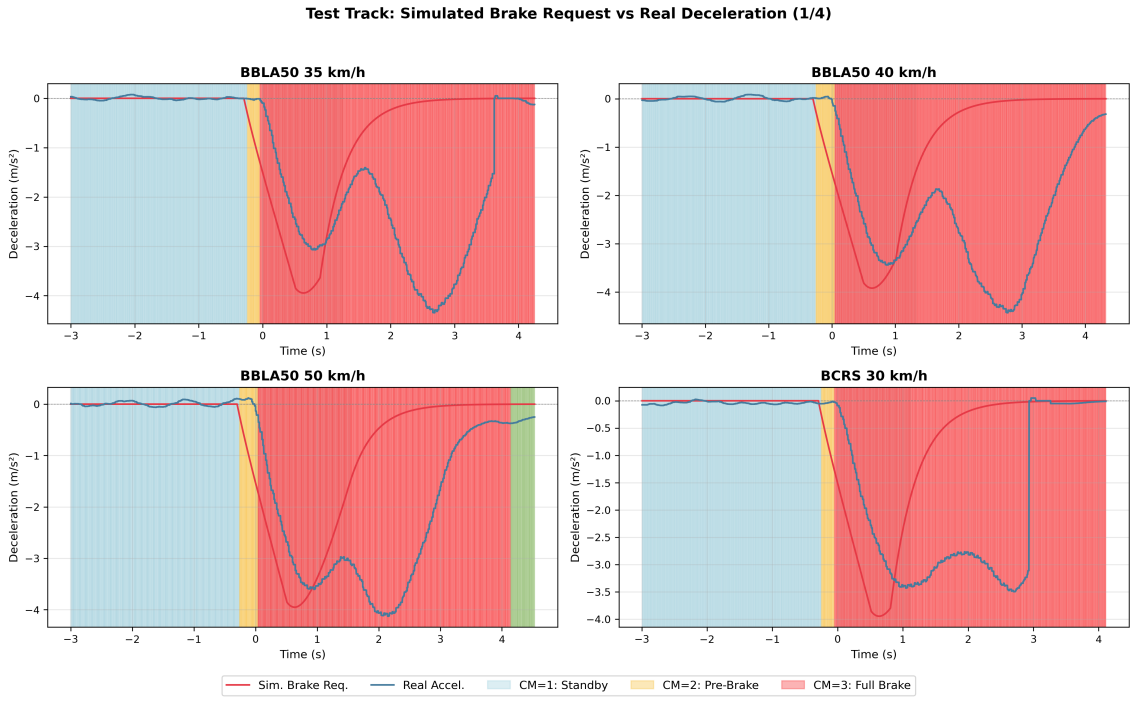

Figure D.1: Simulated brake request vs real deceleration for track test events (page 1 of 4).

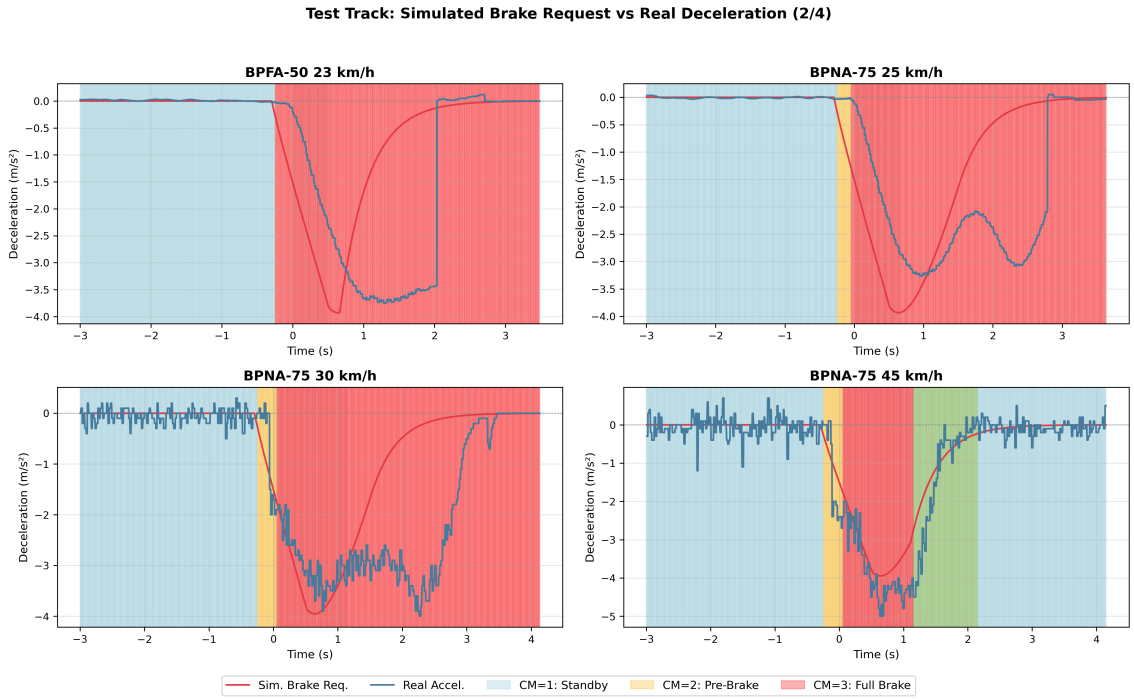

Figure D.2: Simulated brake request vs real deceleration for track test events (page 2 of 4).

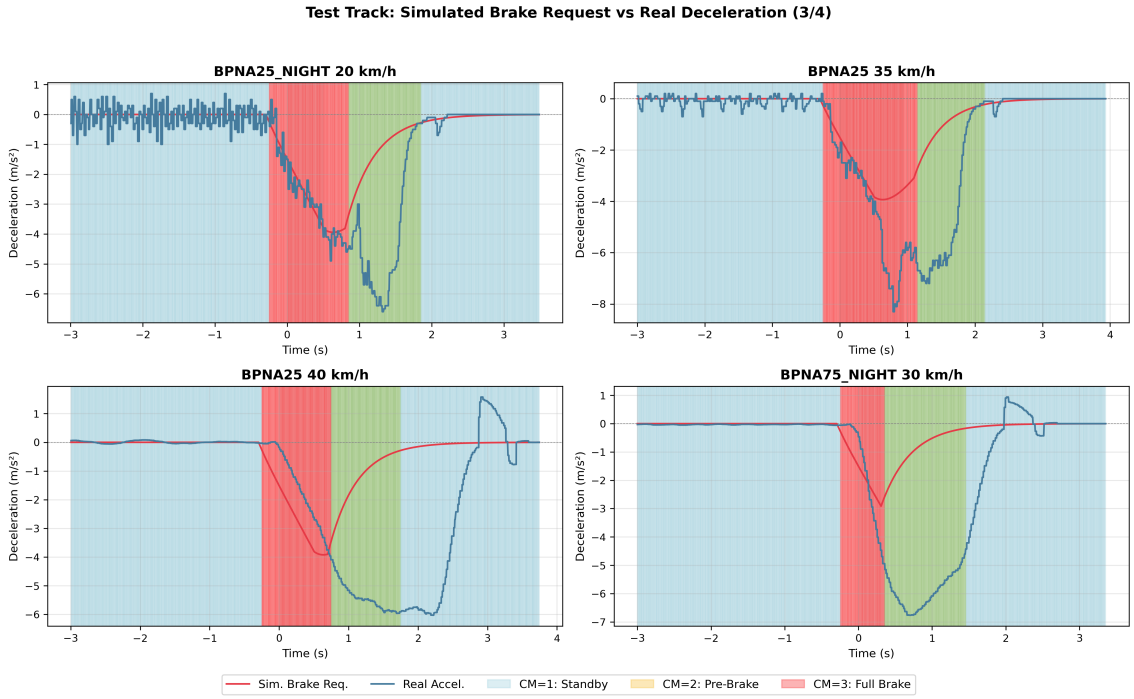

Figure D.3: Simulated brake request vs real deceleration for track test events (page 3 of 4).

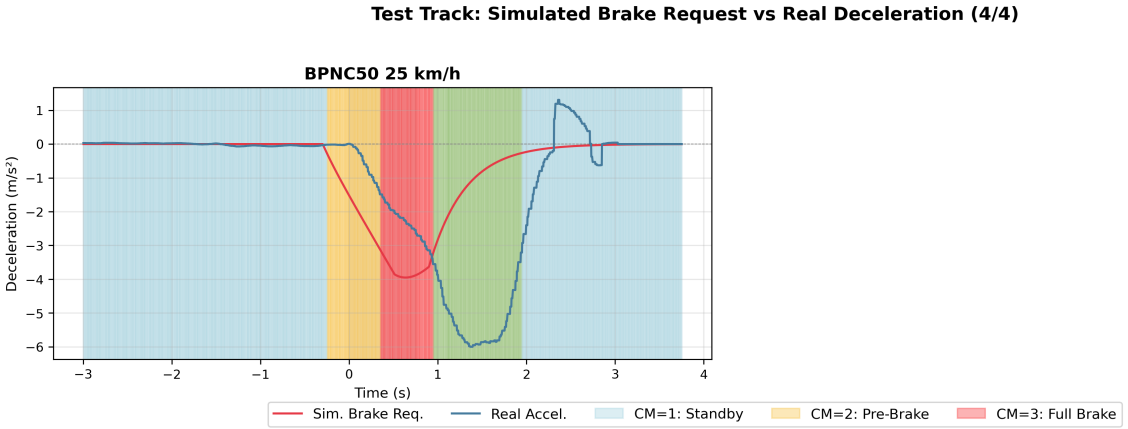

Figure D.4: Simulated brake request vs real deceleration for track test events (page 4 of 4).

D.2 Original vs Corrected Simulation — Remaining Calibration Events

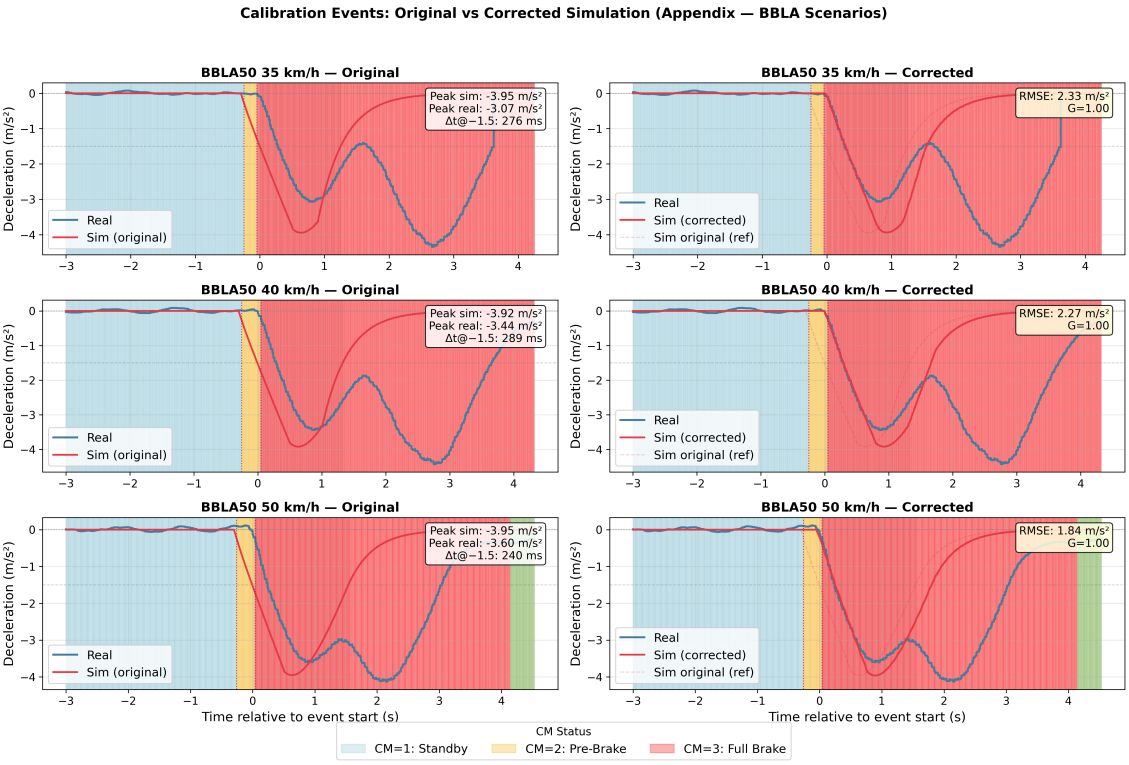

Figure D.5: Original vs corrected simulation for BBLA50 scenarios (bicycle crossing, 50% overlap) at 35, 40, and 50 km/h.

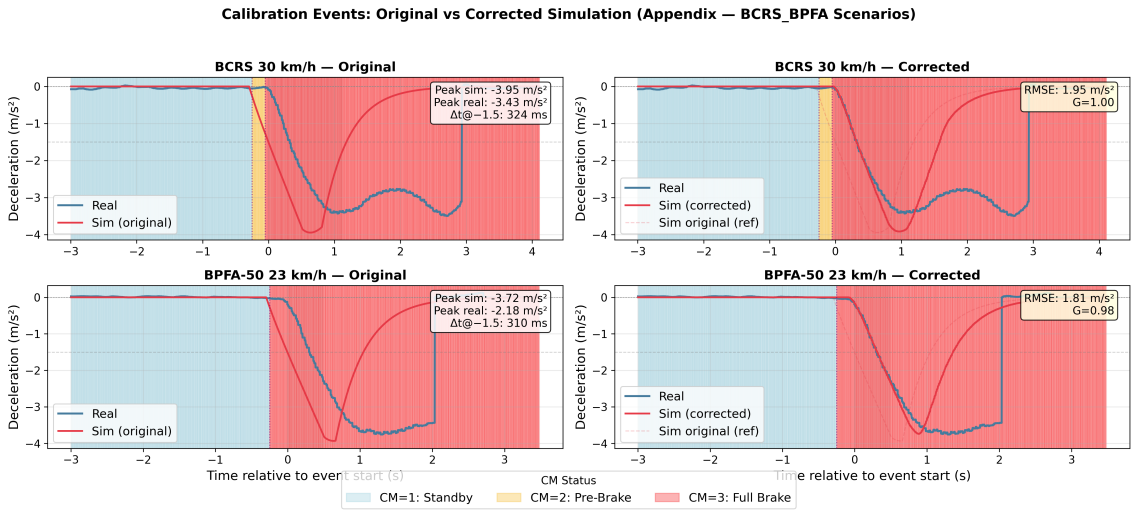

Figure D.6: Original vs corrected simulation for BCRS (car rear stationary, 30 km/h) and BPFA-50 (pedestrian far-side, 23 km/h).

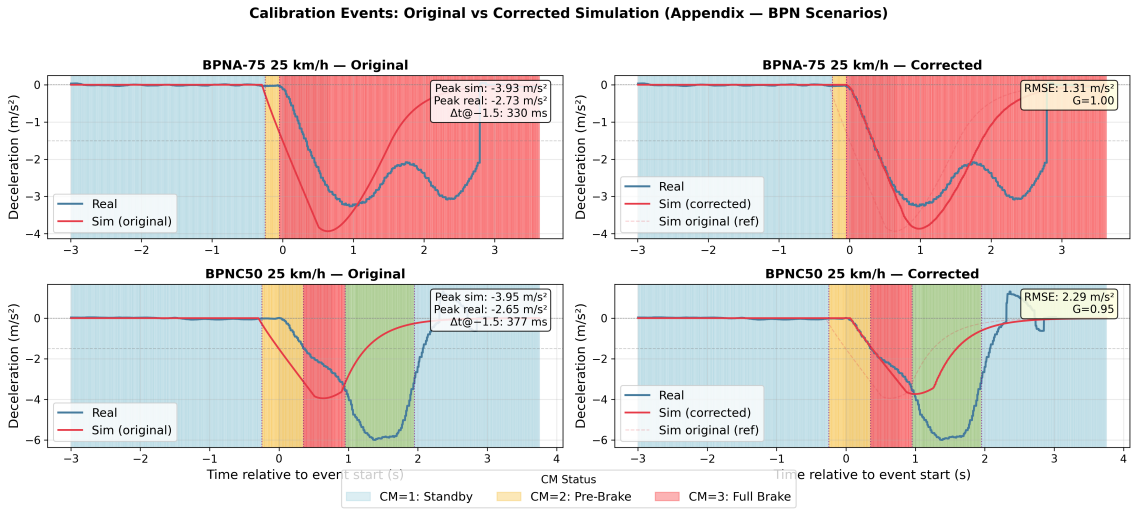

Figure D.7: Original vs corrected simulation for pedestrian near-side scenarios: BPNA-75 (75% overlap, 25 km/h) and BPNC50 (child, 25 km/h).

D.3 Regulatory-to-Pipeline Traceability

Table D.1 maps each regulatory requirement cited in the main text to the pipeline stage that implements it.

Table D.1: Traceability matrix between regulatory requirements and pipeline implementation.

Requirement	**Source**	**Pipeline Stage**
$a \leq -1.5\,\mathrm{m/s^2}$ gate	TfL §4.3.2.3	D1 (Qualified Brake Event)
Speed gate ($> 10\,\mathrm{km/h}$)	R-131-02 §5.2.1.3	D1
Speed range (20–60 km/h)	R-131-02 §5.2.2.3	D1
Condition A (TTC excess)	TfL §4.3.2.3	D2 (FP Judgment)
Condition B (no driver brake)	TfL §4.3.2.3	D2
Attribution indicators 1–4	TfL §4.3.2.4	D3 (Legitimacy Verification)
Driver override (kick-down, swerve)	R-131-02 §5.3.2	D3

DEPARTMENT OF SOME SUBJECT OR TECHNOLOGY
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY