Advances in Science and Technology Research Journal

Advances in Science and Technology Research Journal, 2027, 21(1), 503–521
https://doi.org/10.12913/22998624/236152
ISSN 2299-8624, License CC-BY 4.0

Received: 2026.08.10
Accepted: 2026.11.15
Published: 2026.12.01

Comparative analysis of static and dynamic calibration methods for an advanced driver assistance systems multifunction camera and front corner radar sensors under varying environmental conditions

Paulina Babuchowska[1], Ireneusz Pielecha[1*]

[1] Faculty of Civil and Transport Engineering, Poznan University of Technology, ul. Piotrowo 3, 61-138 Poznań, Poland
* Corresponding author's e-mail: ireneusz.pielecha@put.poznan.pl

ABSTRACT

Accurate calibration is essential for reliable perception in advanced driver assistance systems (ADAS), because sensor misalignment can degrade object detection, localization, and tracking. This study compared static and dynamic calibration of a forward-looking multifunction camera (MFK) and two front corner radar sensors, designated as nanoradars (NR) in the project documentation, installed on an MQB-based test vehicle. Static calibration was performed under controlled end-of-line conditions, whereas dynamic calibration was conducted during on-road operation. Camera yaw, pitch, and roll were evaluated in static, daytime, and nighttime scenarios; radar yaw was evaluated statically and during daytime dynamic trials at three locations. Because the road tests did not include an independent metrological ground truth, the analysis used deviation from the static reference, dispersion, repeatability, radar-pair descriptors, and utilization of project-specific acceptance limits, rather than absolute pose accuracy. All results met the applicable criteria. The largest camera offset and dispersion occurred in roll at nighttime Location C. The radar pair exhibited persistent opposite-sign yaw estimates with unequal magnitudes and sensor-specific dynamic responses, while reflector configuration affected the location-dependent results. Therefore, dynamic calibration provides an operationally representative but environment-dependent calibration state. Static calibration remains necessary as a controlled baseline, and dynamic calibration provides complementary in-use validation. Their combination supports more robust calibration strategies for automotive perception sensors. The principal scientific contribution is therefore a matched, repeated comparison of manufacturer-oriented static and dynamic calibration procedures, combined with an evaluation of how the calibration method and calibration conditions affect the resulting extrinsic parameters. The findings provide a technical basis for assessing whether an optimized dynamic procedure could serve as an alternative to static calibration in selected production contexts; however, this study did not establish superiority in terms of absolute accuracy or production cycle time.

Keywords: sensor calibration, advanced driver assistance systems, multifunction camera, front corner radar, static calibration, dynamic calibration, environmental conditions.

INTRODUCTION

ADAS have become a fundamental component of modern vehicle technology, contributing to road safety, driving comfort, and the gradual transition toward higher levels of driving automation. These systems rely on integrated perception architectures combining multiple sensor modalities, including cameras, radar, LiDAR, and ultrasonic sensors, to detect and interpret the surrounding environment [1, 2]. Their effectiveness depends strongly on the quality and reliability of sensor measurements, which in turn are closely linked to proper sensor calibration [3].

Technical perception systems provide the fundamental measurements required by numerous

driver-assistance functions, enabling the assessment of lanes, objects, free space, and relative motion. Cameras provide rich semantic and geometric information, but are sensitive to variations in illumination and contrast; radar provides direct range and radial-velocity measurements and is comparatively robust to visible-light degradation; and LiDAR provides accurate three-dimensional information, but remains sensitive to visibility conditions and surface properties. Multi-modal perception architectures exploit these complementary capabilities, with recent studies on camera-LiDAR and 4D radar-camera integration demonstrating improved object detection and vehicle-state estimation under challenging visibility conditions [4, 5]. Effective sensor fusion, however, depends on maintaining accurate and stable spatial relationships among the participating sensors. Sensor calibration ensures correct spatial alignment between sensing devices and the vehicle coordinate system. Even minor calibration errors may degrade perception performance, affecting object detection, localization, and tracking [6]. As reported in recent studies, improper calibration may also reduce the effectiveness of critical safety functions, such as lane-keeping assistance and automatic emergency braking, thereby compromising overall system reliability [7].

Two principal calibration approaches are commonly used in ADAS applications: static and dynamic calibration. Static calibration is typically performed in controlled environments using predefined targets and precise measurement setups. Although this approach provides high repeatability and well-defined boundary conditions, it does not fully reflect the variability of real-world driving. Dynamic calibration, by contrast, is performed during vehicle motion, allowing sensor alignment to be estimated under operational conditions, often based on environmental features such as lane markings or surrounding objects [8]. Owing to its ability to capture real driving conditions while reducing the need for specialized calibration infrastructure, this approach has attracted increasing research interest.

Calibration effectiveness is strongly influenced by environmental conditions. Illumination, precipitation, and road-surface characteristics may significantly degrade sensor performance, particularly in optical systems such as cameras and LiDAR [9, 10]. Rain, snow, and fog can reduce detection accuracy and increase measurement uncertainty, thereby lowering the reliability of perception systems [9]. Radar sensors are generally more robust to such disturbances, although their performance may still be affected by signal noise, interference, and reflector geometry [10].

Beyond weather-related influences, ADAS effectiveness also depends on road infrastructure and driving context. Studies have shown that system performance may decline in environments with insufficient road markings or inadequate infrastructure, which directly affects perception reliability [11]. In addition, the interaction among the driver, the vehicle, and the surrounding environment plays an important role in shaping the overall effectiveness of ADAS functions [12].

The broader safety context also includes the interaction between technical vision and human perception. Evidence from camera-monitor systems indicates that the presentation and placement of camera-based visual information can influence driver gaze behavior during maneuvering tasks [13]. Although such systems are not directly related to ADAS sensor calibration, these findings illustrate the broader importance of reliable and geometrically consistent visual information within camera-based vehicle systems.

BACKGROUND

Recent research has highlighted the need to evaluate ADAS under realistic operating conditions, as laboratory-based testing alone cannot capture the complexity of real-world environments [14]. Experimental studies further indicate that sensor performance and system reliability vary substantially with environmental context, road conditions, and traffic scenarios [15]. At the same time, the continuing development of ADAS technologies and sensor-fusion architectures introduces new challenges related to system validation, calibration robustness, and operational reliability [16].

Calibration approaches can be classified according to the reference information used and the timing of parameter estimation. Target-based static calibration relies on known target geometry and a controlled vehicle–target relationship [17]. Dynamic or targetless methods estimate extrinsic parameters from natural scene features, tracked objects, ego-motion, or cross-modal correspondences, whereas online approaches update or monitor sensor alignment during vehicle operation [18, 19]. For radar–camera and LiDAR–camera systems, published methods include trajectory

alignment, track association, shared-feature optimization, end-to-end automatic calibration, and external multi-modal calibration frameworks [20, 21]. These approaches primarily address cross-sensor extrinsic estimation or algorithmic automation, rather than repeated experimental comparison of established static and dynamic calibration procedures applied to the same vehicle.

Calibration and sensor fusion are related but distinct tasks. Fusion algorithms combine complementary measurements to improve object detection, scene interpretation, or vehicle-state estimation, whereas calibration establishes or verifies the transformations required to express those measurements in a common geometric frame. Radar–camera fusion has demonstrated improved robustness under adverse visibility and challenging operating conditions [5, 22]; however, fusion-level performance alone does not distinguish whether an observed deviation originates from sensor alignment, the calibration procedure, or the observation environment.

Environmental observability is modality-specific. Camera-based calibration depends on stable lane, edge, horizon, and perspective cues and may deteriorate when illumination, contrast, contamination, or adverse weather reduces their availability [23]. Radar is independent of visible-light conditions but remains sensitive to reflector geometry, clutter, multipath propagation, interference, and the spatial distribution of stationary objects [24, 25]. Consequently, the same route cannot be assumed to provide equivalent calibration information for both modalities, and formal acceptance does not necessarily imply comparable repeatability, robustness, or remaining tolerance margin. Despite this progress, comprehensive studies directly comparing static and dynamic calibration under varying environmental conditions remain limited, particularly for multi-sensor systems. The lack of standardized methodologies for evaluating calibration performance also remains a significant challenge in ADAS development. Much of the existing research focuses on individual sensor types or specific experimental setups, without offering a unified framework for assessing calibration effectiveness across different operating conditions [8, 26].

The specific knowledge gap bridged here is the limited experimental evidence showing how manufacturer-oriented static and dynamic calibration procedures, applied to the same installed sensor set, differ in their reported calibration state, within-scenario repeatability, and sensitivity to the observation environment. The central premise of the study was therefore to compare the two methods directly, evaluate how the calibration method and calibration conditions affect the resulting functionally relevant extrinsic parameters, and determine whether the findings justify further optimization of dynamic calibration as a potential alternative for selected production applications. The scientific contribution lies in the matched experimental design, randomized repeated trials, evaluation of dispersion and statistical uncertainty, comparison with a documented static procedural reference, and parallel but separate assessment of camera and radar modalities under controlled and operational conditions.

RESEARCH AIM AND OBJECTIVES

The aim of the present study was to compare static and dynamic calibration procedures for a forward-looking multifunction camera and a pair of front corner radar sensors installed on the same vehicle. Particular emphasis was placed on determining how the calibration method and operating conditions, including location, scene geometry, reflector distribution, and day/night conditions, affect the resulting calibration state and repeatability of the estimated extrinsic parameters.

The scientific contribution lies in the matched experimental comparison of both established calibration procedures and in identifying calibration parameters and environmental conditions that are particularly relevant to calibration robustness. Since these extrinsic parameters determine the geometrical consistency of sensor information used by downstream ADAS functions, their sensitivity is also relevant to subsequent sensor operation. The findings provide a basis for further optimization of dynamic calibration and for evaluating its potential use as an alternative to selected static production procedures. Direct function-level ADAS performance was not evaluated in the present study. Accordingly, two research questions were formulated:

RQ1. How do static and dynamic calibration procedures and the conditions under which they are performed affect the calibration state and repeatability of the investigated camera and radar sensors?

RQ2. Which calibration parameters and environmental factors show the greatest sensitivity, and what is their potential relevance to the

geometrical consistency of sensor information used by downstream ADAS functions?

Because the study is exploratory, uses one vehicle, and lacks independent pose ground truth, these questions were evaluated descriptively, rather than as confirmatory hypotheses of absolute calibration accuracy or production-process superiority.

METHODOLOGY

Test vehicle and sensor systems

The investigation was performed using an MQB-based test vehicle and covered two perception subsystems: a forward-looking multifunction camera (MFK) and a pair of front corner radar sensors designated NR_1 and NR_2. The vision analysis was restricted to MFK, while the radar analysis included only the two front corner sensors. Rear and side radar sensors were present in the wider vehicle E/E architecture described in the source documentation, but they were outside the experimental scope of this study. The camera and radar subsystems were calibrated and evaluated independently. No direct radar-camera extrinsic calibration, joint optimization, or fusion-level performance assessment was performed.

The experimental configuration followed the project-specific geometry of the analyzed MQB variant. The front radar sensors represented the low-installation variant, with a nominal azimuth of 65 degrees, and project-defined reference coordinates in the vehicle coordinate system. NR_2 was treated as the mirrored counterpart of NR_1 about the vehicle centerline. Opposite signs were therefore interpreted with the sensor coordinate convention; residual pair asymmetry was identified from unequal magnitudes or a non-zero pair sum, not from sign alone.

Experimental setup

Both sensor subsystems were evaluated using dedicated production- and service-oriented calibration routines supported by diagnostic communication, precisely referenced targets or simulators, and stand-based vehicle-positioning equipment. Static MFK calibration was performed on an ALS/FAS calibration station equipped with a laser measurement system and the VAS6430/4 optical target (Figure 1).

Static calibration of the front corner radar sensors used Doppler generators as synthetic moving targets in combination with precise vehicle positioning and geometric referencing in the station coordinate system. The generators reproduced the Doppler response required by the calibration routine without a physically moving target. Figure 2 presents the sensor arrangement and the associated calibration infrastructure.

Instrumentation and data acquisition

Calibration outputs were recorded through diagnostic variables. For the camera, the relevant parameters were yaw, pitch, roll, and camera height. For the front radar pair, the principal output was the estimated yaw angle of each sensor, together with calibration status and progress information. These variables formed the basis for the comparative analyses presented in this study.

Camera height was monitored as a diagnostic control variable, but was not included in the comparative statistical analysis because the vehicle loading condition and camera installation



Figure 1. Schematic representation of the static MFK calibration setup, including the ALS/FAS station, optical target, and vehicle-positioning system [27]



Figure 2. Static front corner radar calibration station, including the alignment structure, vehicle-positioning system, and Doppler generators [27]

geometry were unchanged, and no relevant height variation was observed across the test scenarios. Camera height was therefore treated as a controlled invariant, rather than an experimental factor, and could not account for the between-scenario differences reported here. This assumption does not imply that camera height is physically irrelevant: a change in sensor height would alter the camera-to-road projection geometry and would require separate evaluation.

Experimental scope

A structured test matrix was defined to evaluate the influence of calibration procedure, test location, and illumination on the reported calibration state. For the MFK, the matrix comprised one static reference scenario and dynamic scenarios conducted during the day and at night at Locations A, B, and C. For the front radar pair, the matrix comprised one static reference and daytime dynamic scenarios at the same three locations. Table 1 summarizes the complete experimental design.

The six-trial dynamic radar series was used exclusively for the direct static-dynamic comparison presented in Figure 7 and was not stratified by Locations A-C. The location-specific radar series comprised ten independent trials per location.

In methodological terms, the static scenarios served as the controlled baseline, whereas the dynamic scenarios represented calibration under realistic environmental conditions with varying geometry, reflectivity, and illumination. This design enabled comparison not only between static and dynamic procedures, but also across different locations and, for the camera subsystem, between day and night operation.

To preserve project confidentiality, the dynamic test sites are identified as Locations A-C. Location A represented an industrial environment characterized by metallic infrastructure, machinery, parked vehicles, and building edges; Location B represented a comparatively open road environment with a less constrained reflector distribution; and Location C represented a structured proving-ground environment. These classifications describe the predominant visual and radar-observation characteristics of the respective routes.

All repeated calibrations were performed as independent trials, and the order of the trials was randomized to reduce sequence effects, temporal drift, and systematic bias associated with a fixed test order. Each calibration was initiated as a separate measurement run under the boundary conditions applicable to the corresponding scenario.

Daytime and nighttime tests were defined as categorical operating conditions based on the predominant source of scene illumination. Daytime tests were conducted under natural daylight, whereas nighttime tests were performed when the road scene was primarily illuminated by the vehicle headlamps and surrounding artificial light sources. Illuminance was not measured; therefore, the day/night distinction is not interpreted as a quantitative illumination scale. Other environmental variables, including precipitation, air temperature, road-surface properties, complete vehicle-speed histories, and the spatial distribution of radar reflectors, were likewise not recorded

Table 1. Test scenarios considered in the study

Sensor subsystem	Calibration mode	Time condition	Location	Scenario role	n
MFK front camera	Static	Controlled conditions	ALS/FAS station	Reference condition	10
MFK front camera	Dynamic	Day	A	Environmental comparison	10
MFK front camera	Dynamic	Day	B	Environmental comparison	10
MFK front camera	Dynamic	Day	C	Environmental comparison	10
MFK front camera	Dynamic	Night	A	Operating-condition comparison	10
MFK front camera	Dynamic	Night	B	Operating-condition comparison	10
MFK front camera	Dynamic	Night	C	Operating-condition comparison	10
Front corner radar pair (NR_1, NR_2)	Dynamic	Day	Not location-stratified	Direct static-dynamic comparison	6
Front corner radar pair (NR_1, NR_2)	Static	Controlled conditions	ALS station	Reference condition	6
Front corner radar pair (NR_1, NR_2)	Dynamic	Day	A	Environmental comparison	10
Front corner radar pair (NR_1, NR_2)	Dynamic	Day	B	Environmental comparison	10
Front corner radar pair (NR_1, NR_2)	Dynamic	Day	C	Environmental comparison	10

quantitatively. Route geometry and reflector distribution are therefore described qualitatively, while vehicle operation followed the procedural thresholds specified below.

Calibration procedures

Static calibration was conducted under controlled EOL conditions. For the MFK, the procedure corresponded to initial static calibration (iSCAL) on the ALS/FAS station. Its purpose was to identify and compensate installation tolerances after sensor assembly by estimating the camera extrinsic parameters relative to the vehicle coordinate system. The sequence comprised positioning the complete vehicle on the station, aligning the target, transferring the required adaptation values, initiating the diagnostic calibration routine, and verifying successful completion. The VAS6430/4 optical target was positioned at a nominal longitudinal distance of 1.5 m from the front-wheel center and adjusted to the project-specific target height.

Before initiation of the calibration routine, the vehicle had to satisfy the prescribed boundary conditions: complete assembly status, correct software parameterization, specified tire pressure, no additional payload, stationary vehicle state, an unobstructed optical path, and illumination of the target within the approved range. These prerequisites were intended to limit suspension-induced, geometric, and photometric disturbances during estimation of the extrinsic parameters.

For the radar subsystem, static calibration was executed only for NR_1 and NR_2. The vehicle was centered on the alignment station and its position was measured relative to the station reference frame. The Doppler generators assigned to the front sensors were then moved to their nominal coordinates, the measured coordinates were transferred to the diagnostic environment, the calibration routine was activated, and the reported result was verified. In practical terms, this EOL process determined compensation values for the yaw misalignment of each radar relative to its project-specific reference orientation.

Dynamic calibration was conducted while the vehicle was in motion and used features naturally present in the driving environment rather than a dedicated station target. For the MFK, the on-road calibration routine was performed on open roads and required a vehicle speed of at least 20 km/h. Successful execution also required activation of the appropriate calibration mode, a valid system state, sufficient feature quality, and a road segment that was adequately straight and stable for parameter convergence.

For the front corner radar sensors, dynamic calibration was implemented as Service Dynamic Calibration (SDK). The procedure required an urban or comparable route of at least 1000 m containing stationary radar-reflective structures such as guardrails, fences, building edges, or parked vehicles. The recommended vehicle-speed range was 20 to 120 km/h, while calibration progress

was generated above 15 km/h. For stable estimation, the technical specification recommended approximately straight road sections, a vehicle yaw rate below 0.5 degrees/s, longitudinal acceleration below 12.5 m/s², and a lateral distance of approximately 5 to 8 m from the reflective roadside environment.

The radar SDK routine imposed explicit environmental boundary conditions. Radar covers had to be free from dirt, snow, mud, stickers, transport protection, and other attenuating or reflecting obstructions. Calibration was not recommended during heavy rain, snowfall, or hail, or on routes lacking a sufficiently distributed set of stationary reflectors. Interference from nearby active automotive radars was also to be minimized. The final evaluation used calibration progress, diagnostic status, timestamps, and estimated yaw. The applicable yaw acceptance limits were ±5 degrees for the low installation variant and ±3 degrees for the high variant; the sensors investigated here were the low variant.

The static and dynamic procedures differed in both boundary conditions and the physical information used by the estimator. Camera calibration relied on structured visual features and returned several rotational extrinsic parameters, whereas front-corner radar calibration primarily estimated yaw from Doppler-simulator responses or stationary environmental reflections. This modality-specific distinction is essential when comparing robustness, because reduced illumination primarily affects the availability of visual features, while radar calibration is more directly influenced by reflector distribution, clutter, interference, and multipath propagation.

Evaluation criteria

Calibration performance was evaluated using complementary criteria: the reported calibration state, deviation from the controlled static reference, within-scenario dispersion and repeatability, sensor-pair consistency, and compliance with project-specific acceptance limits. The post-calibration rotational parameters are illustrated in Figure 3. For MFK, the outputs were yaw, pitch, and roll; for NR_1 and NR_2, the principal output was yaw. The static result was used as a procedural reference rather than as an independent metrological ground truth. Consequently, the study quantifies the differences among calibration solutions, but does not claim absolute pose accuracy.

Calibration-state deviation was defined as the difference between an estimated parameter and the selected reference state. For MFK, this included changes in yaw, pitch, and roll between static and dynamic conditions, among locations, and between daytime and nighttime operation. For the front radar pair, deviation was represented by the difference between static and dynamic yaw estimates and by the location-dependent changes observed during on-road calibration. These quantities characterize the method- and environment-dependent offsets; they should not be interpreted as object-detection error or absolute calibration error in the absence of an external pose reference.

Stability and repeatability were assessed from repeated calibrations within each scenario. For every parameter and scenario, the arithmetic average, sample standard deviation, and coefficient of variation were calculated. The average represented the central calibration result, the sample standard

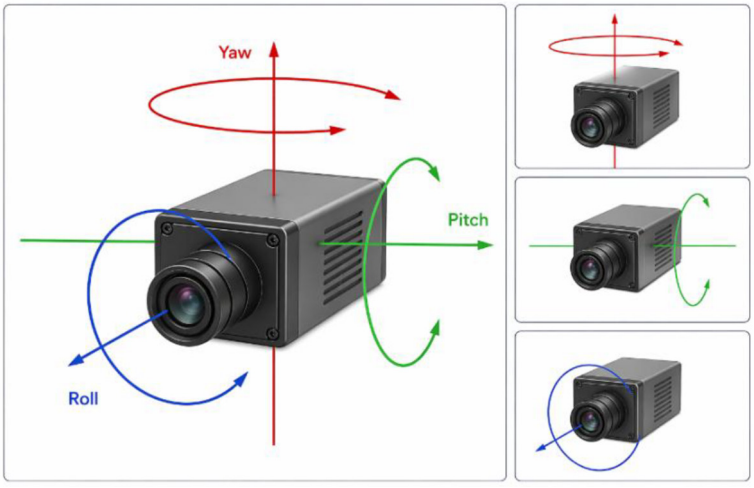


Figure 3. Rotational parameters yaw, pitch, and roll relative to the sensor coordinate system [27]

deviation quantified the absolute dispersion of repeated trials, and the coefficient of variation provided a normalized dispersion measure for the comparisons in which the average was sufficiently separated from zero. Lower dispersion was interpreted as better within-scenario repeatability, whereas higher dispersion indicated greater sensitivity to environmental or operational variation. Because angular averages can be close to zero, the coefficient of variation was interpreted only together with the corresponding average and standard deviation. The combined criteria separated the nominal calibration state, its deviation from the static procedural reference, and the robustness of repeated estimates.

Mathematical framework for calibration analysis

The mathematical framework was restricted to the quantities used directly to interpret Figures 4–8. For scenario s, parameter k, and repeated trial j, x denotes camera yaw, pitch, or roll, or the yaw estimate of one front corner radar sensor. The camera scenarios and the location-specific radar scenarios each comprised ten repetitions, whereas the direct static-dynamic radar comparison comprised six repetitions.

The arithmetic average was used to represent the central calibration result for each parameter and scenario:

$$\bar{x}_{k,s} = \frac{1}{n}\sum_{j=1}^{n} x_{k,s,j} \tag{1}$$

Repeatability was quantified using the sample standard deviation:

$$s_{k,s} = \sqrt{\frac{1}{n-1}\sum_{j=1}^{n}\left(x_{k,s,j} - \bar{x}_{k,s}\right)^2} \tag{2}$$

Dispersion relative to the magnitude of the average was expressed by the coefficient of variation. The absolute average in the denominator ensures a positive percentage; the result was interpreted together with the average because this measure becomes unstable when the average approaches zero [28, 29].

For calibration angles close to zero, the coefficient of variation was not treated as the primary repeatability metric, because division by a small absolute mean can produce disproportionately large or unstable percentage values. In such cases, the sample standard deviation and the corresponding confidence interval provide more interpretable measures of repeatability. The coefficient of variation was therefore retained only as a supplementary normalized descriptor and was not used independently to rank calibration scenarios.

$$CV_{k,s} = \frac{s_{k,s}}{\left|\bar{x}_{k,s}\right|} \times 100\% \tag{3}$$

The systematic change associated with a dynamic calibration scenario was calculated as the difference between the dynamic-scenario average and the corresponding static-reference average:

$$\Delta x_{k,s} = \bar{x}_{k,s}^{\mathrm{dyn}} - \bar{x}_{k}^{\mathrm{stat}} \tag{4}$$

Two additional descriptors were applied to the mirrored front radar pair. Pair separation quantified the distance between the average yaw estimates reported for NR_1 and NR_2:

$$D_s = \left|\bar{\psi}_{s,NR_1} - \bar{\psi}_{s,NR_2}\right| \tag{5}$$

The mirror-symmetry index quantified the residual departure from an ideal opposite-sign pair after accounting for the mirrored installation and sign convention. Values closer to zero indicate closer agreement in magnitude and should not be interpreted as proof that either individual sensor is correctly aligned:

$$S_s = \left|\bar{\psi}_{s,NR_1} + \bar{\psi}_{s,NR_2}\right| \tag{6}$$

Finally, tolerance utilization compared the absolute scenario average with the applicable directional magnitude of the project-specific acceptance limit. For an asymmetric interval, the denominator corresponds to the limit on the same side of zero as the scenario average:

$$\rho_{k,s} = \frac{\left|\bar{x}_{k,s}\right|}{L_k} \tag{7}$$

To quantify the statistical uncertainty of each scenario average, a two-sided 95% confidence interval was calculated using Student's t distribution:

$$\bar{x}_{s,k} \pm t_{0.975,n_s-1}\frac{s_{s,k}}{\sqrt{n_s}} \tag{8}$$

where: t is the critical value at the 97.5th percentile of Student's t distribution, with the degrees of freedom equal to the number of repetitions in the scenario minus one.

The calculation assumes independent and approximately normally distributed repeated estimates. The interval characterizes statistical uncertainty in the scenario average; it is not a substitute for a complete metrological uncertainty budget. A tolerance-utilization value below one confirms that the scenario average complies with the relevant acceptance limit, whereas a value approaching one indicates reduced margin to that limit. Tolerance utilization is not a substitute for evaluation of individual trials, uncertainty, or repeatability. Together, the eight statistical and comparative measures provide the quantitative basis for interpreting the reported calibration states [28, 29].

EXPERIMENTAL RESULTS

MFK calibration

MFK calibration was evaluated using yaw, pitch, and roll, which describe the rotational relationship between the camera coordinate system and the vehicle reference frame. Yaw is rotation about the vertical axis and governs horizontal pointing; pitch is rotation about the lateral axis and describes upward or downward inclination of the optical axis; roll is rotation about the longitudinal axis and represents lateral image-plane tilt relative to the vehicle horizon. These extrinsic parameters are safety-relevant because they influence lane-model projection, horizon estimation, object localization, and the geometric consistency of camera data used by downstream ADAS functions.

According to the project-specific calibration guidelines, the admissible ranges for both static calibration (iSCAL) and dynamic calibration (iOCAL) were –2.0° to 2.0° for roll, –3.5° to 2.5° for pitch, and –3.0° to 3.0° for yaw. All measured values obtained in the present study remained within these limits, indicating that every analyzed calibration procedure was formally successful. However, compliance with the acceptance criteria did not imply identical calibration behavior across scenarios.

The ALS-based static series established a compact and controlled reference distribution (Figure 4). Across ten repetitions, yaw was concentrated at approximately 0.5 to 0.8 degrees, pitch at 0.2 to 0.5 degrees, and roll at 0.4 to 0.6 degrees. The limited spread is consistent with a station-based procedure in which target geometry, vehicle position, target illumination, and diagnostic boundary conditions are prescribed. It should nevertheless be regarded as a procedural baseline rather than an independent measurement of the true camera pose.

The residual spread can plausibly be attributed to small differences in vehicle positioning, suspension settling, target illumination, or other station-related influences. The static estimates occupied only a narrow portion of the available acceptance intervals, showing that a binary pass/fail assessment would conceal useful information about the consistency of the calibration result. The subsequent numerical analysis also shows that one dynamic scenario, daytime Location B, achieved even lower within-scenario standard deviations; the controlled static condition should therefore not automatically be described as the most repeatable series.

A distinct change was observed during daytime dynamic calibration, where the estimated angles depended on test location (Figure 5). At Location A, yaw ranged from –0.33 to –0.18 degrees, pitch from –0.76 to –0.43 degrees, and roll from –0.39 to –0.15 degrees. Although all values satisfied the project limits, the distributions were systematically shifted from the ALS baseline. Therefore, the on-road routine converged to a condition-dependent calibration state derived from environmental features rather than

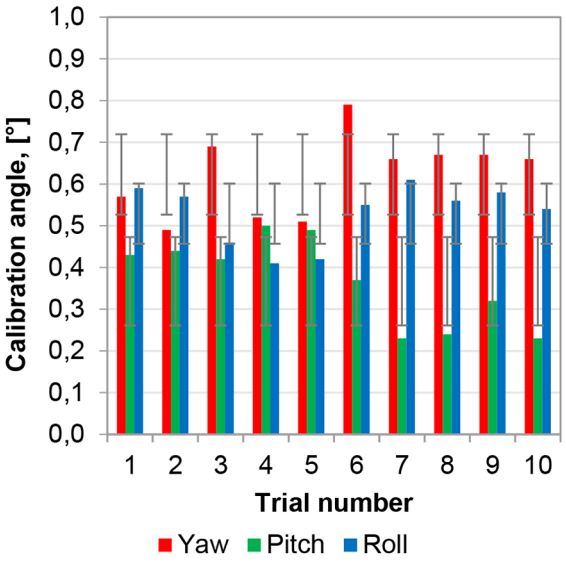


Figure 4. Yaw, pitch, and roll angles obtained in ten static MFK calibration trials under controlled EOL conditions. Error bars denote one standard deviation calculated across the ten repeated values of the corresponding parameter

reproducing the target-based station solution. The more negative pitch estimate may be associated with the local road profile, lane and edge visibility, perspective geometry, or the balance of vertical and longitudinal scene cues; without an external pose reference, the shift cannot be assigned uniquely to true installation error or algorithmic bias.

In Location B, the daytime values ranged from –0.46° to –0.36° for yaw, –0.72° to –0.62° for pitch, and –0.53° to –0.43° for roll. Compared with Location A, the values were more negative but substantially more tightly grouped. The calculated standard deviations confirm that Location B provided the highest repeatability among the daytime dynamic scenarios despite its systematic offset from the ALS reference. The deviation is therefore predominantly systematic rather than random, which suggests that this location provided consistent roadway edges, perspective cues, and lane structures.

The most demanding daytime condition was Location C, where yaw ranged from –0.30° to –0.15°, pitch from –0.72° to –0.54°, and roll from –0.97° to –0.40°. Although all values remained within project tolerance, the substantially larger roll spread indicates reduced repeatability and greater sensitivity of lateral tilt estimation to local scene conditions. Roll estimation depends strongly on lateral scene symmetry and stable horizon-related cues; nonuniform roadside structure, shadows, partial occlusions, irregular pavement geometry, or weak lateral references can therefore increase inter-trial variation. The relatively compact yaw distribution indicates that the scene still provided sufficient information for horizontal directional alignment.

The influence of illumination was more evident in the nighttime dynamic series (Figure 6). At nighttime, Location A, yaw ranged from –0.46 to –0.21 degrees, pitch from –0.53 to –0.40 degrees, and roll from –0.71 to –0.39 degrees. Relative to daytime operation, the average yaw became more negative, pitch shifted toward zero, and roll moved farther from the ALS reference. Therefore, reduced illumination changed both the central estimate and its dispersion rather than merely adding random noise. Continuous lane boundaries, road edges, distant geometry, and horizon cues were partly replaced by localized headlamp illumination, specular reflections, and point light sources.

A similar tendency was observed in nighttime Location B, where yaw ranged from –0.43° to –0.12°, pitch from –0.53° to –0.40°, and roll from –0.74° to –0.36°. All values remained within the admissible limits. Compared with daytime Location B, however, the larger yaw and roll standard deviations show that repeatability deteriorated under reduced illumination. The average pitch became less negative and the average roll more negative, indicating a systematic illumination-dependent shift in addition to the increased spread.

The most unfavorable behavior was observed in nighttime Location C, where yaw ranged from –0.37° to –0.21°, pitch from –0.58° to –0.45°, and roll from –1.30° to –0.76°. These were the most extreme camera results, especially for roll, although they remained within the admissible roll interval of –2.0° to 2.0°. The substantial negative shift and broad roll spread indicate that nighttime Location C provided insufficiently stable lateral and horizon-related visual references. The smaller changes in yaw and pitch support the conclusion

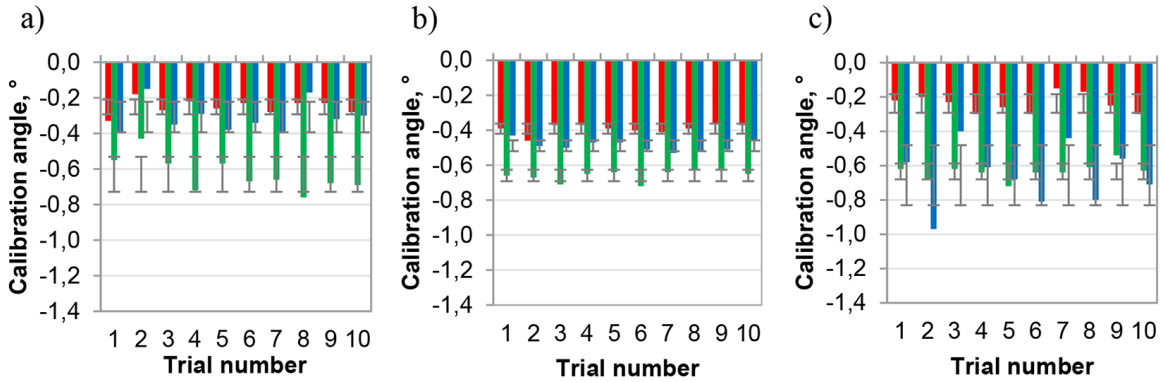


Figure 5. Yaw (blue), pitch (orange), and roll (green) angles obtained during daytime dynamic MFK calibration at: (a) location A, (b) location B, and (c) location C. Error bars denote one standard deviation calculated across the ten repeated values of the corresponding parameter at each location

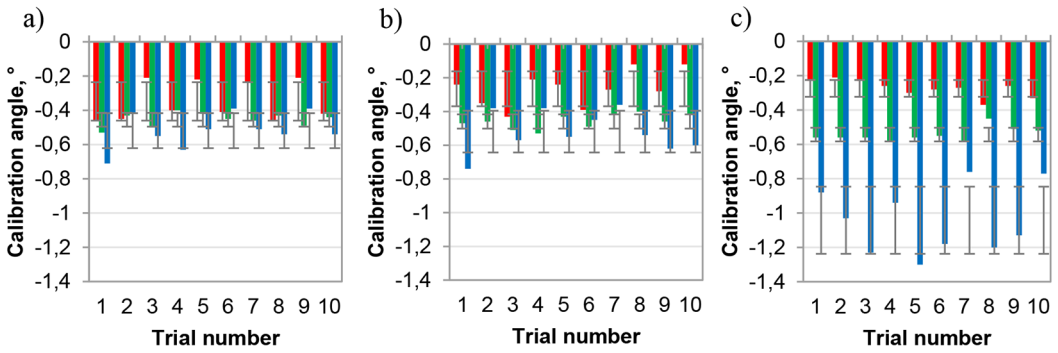


Figure 6. Yaw (blue), pitch (orange), and roll (green) angles obtained during nighttime dynamic MFK calibration at: (a) location A, (b) location B, and (c) location C. Error bars denote one standard deviation calculated across the ten repeated values of the corresponding parameter at each location

that lateral orientation estimation is the most vulnerable component of dynamic camera calibration under degraded visual conditions.

Overall, the measured yaw, pitch, and roll values remained within the project-defined tolerance ranges, and all calibrations were formally acceptable. The ALS procedure provided the controlled reference state, whereas the dynamic scenarios introduced location- and illumination-dependent shifts. Within the daytime tests, Location B had the lowest dispersion, and Location C had the largest roll dispersion. At night, roll moved farther from the static reference at all locations, with the strongest shift and dispersion in Location C. Roll was therefore the most environment-sensitive camera parameter.

Front corner radar calibration

The front corner radar pair was evaluated using estimated yaw, which represents the azimuthal compensation reported for each sensor relative to its project-specific reference orientation and is the principal output of both EOL static calibration and the SDK routine. The technical specification defines an acceptance interval of ±5 degrees for the low sensor variant and indicates a measurement uncertainty of approximately ±1.5 degrees. Every yaw observation in the dataset was within the formal project limits. As for the MFK, however, threshold compliance did not establish equivalence between calibration procedures or exclude systematic differences between the two sensors.

The static series already showed a persistent opposite-sign pattern (Figure 7). NR_1 ranged from approximately 1.79 to 2.75 degrees, whereas NR_2 ranged from approximately –2.04 to –1.66 degrees. The radar pair was therefore not centered near zero but formed an approximately mirrored set of compensation values. Because the sensors are installed as mirrored front corner units, the algebraic sign must be interpreted in the applicable coordinate convention; opposite signs alone do not demonstrate a physical installation defect. However, the unequal magnitudes and non-zero pair sum indicate residual pair asymmetry in the reported calibration state. This pattern was observed under controlled EOL conditions with station-referenced vehicle centering, predefined Doppler-generator coordinates, specified generator orientation, an unobstructed sensor-to-generator path, and restrictions on disturbing metallic structures. The persistence of the pattern is therefore more compatible with the complete installation and reference chain – sensor, bracket, bumper, guide profile, chassis, nominal azimuth, and coordinate convention – than with random environmental variability alone.

Static and dynamic calibration did not converge to identical yaw estimates, and the difference was sensor-dependent. For NR_1, dynamic yaw ranged from approximately 2.32 to 2.60 degrees and remained within the same positive region as the static observations, but with lower dispersion. Therefore, the SDK procedure retained the positive NR_1 tendency while producing a more compact estimate near 2.4 to 2.6 degrees. NR_2 showed a larger negative change: its static values ranged from approximately –2.04 to –1.66 degrees, whereas the dynamic values ranged from –3.60 to –3.00 degrees. The magnitude of the reported yaw consequently increased by approximately 1.1 to 1.6 degrees relative to the static baseline. Because the technical documentation provides an uncertainty value but the present

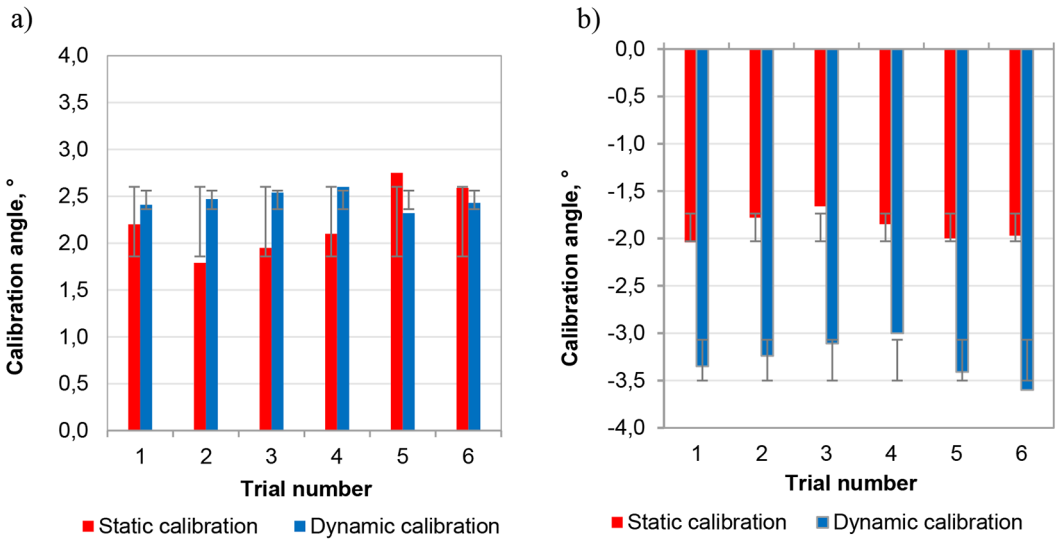


Figure 7. Comparison of yaw angles obtained in six static and dynamic calibration trials for: (a) NR_1 and (b) NR_2 Error bars denote one standard deviation calculated across the six repetitions for each calibration mode

study does not establish a complete uncertainty budget for the difference between two calibration procedures, strict metrological significance cannot be assigned to individual static-dynamic differences. The repeated and sensor-specific shift nevertheless indicates that the two procedures respond differently to their observation models. Static calibration references each radar to synthetic Doppler-generator geometry, whereas SDK estimates yaw from the real stationary-reflector field encountered during vehicle motion. A deviation can therefore arise from the installed azimuth, the nominal reference, vehicle dynamics, reflector asymmetry, or interactions among these factors even when both routines achieve formal acceptance.

The location-specific SDK results further demonstrate that estimated radar yaw depended on the environmental reflector configuration (Figure 8). At Location A, NR_1 ranged from 2.41 to 3.05 degrees and NR_2 from –4.03 to –1.63 degrees. The scenario averages were close to mirror-symmetric in magnitude, but the wide NR_2 distribution indicated poor repeatability for the second sensor. The industrial surroundings, including metallic infrastructure, machinery, parked vehicles, and building edges, plausibly produced a strong but spatially unbalanced reflector field that influenced NR_2 more than NR_1.

In Location B, NR_1 ranged from 1.83° to 3.19°, while NR_2 ranged from –4.06° to –2.48°. The more open environment provided enough stationary radar-reflective objects for successful calibration but did not eliminate the pair asymmetry. Both sensors showed greater dispersion than in the most repeatable location-specific series, which is consistent with a less constrained reflector geometry and a stronger environmental influence on the estimated yaw.

In Location C, NR_1 ranged from 2.83° to 4.01°, while NR_2 ranged from –3.12° to –2.61°. This location yielded the highest positive values for NR_1 and a compact negative interval for NR_2. The structured test-track reflector field therefore stabilized the second radar but did not center the sensor pair; instead, it reinforced the positive yaw tendency of NR_1. Repeatability and agreement with the static reference are consequently distinct evaluation criteria.

The differences among locations are physically consistent with the SDK boundary conditions defined in the technical specification. Calibration progress is generated only when vehicle speed exceeds 15 km/h, yaw rate remains below 0.5 degrees/s, acceleration remains below 12.5 m/s², radar interference is limited, and a sufficient set of stationary reflective targets is available. The documentation additionally recommends approximately straight routes, low-interference surroundings, clean sensor surfaces, and the absence of severe weather. The reported dynamic yaw is therefore a property of the complete sensor-vehicle-environment system during motion rather than of the sensor alone. Dense industrial clutter can overrepresent one side of the reflector field; an open environment can weaken geometric constraint and increase variability; and a

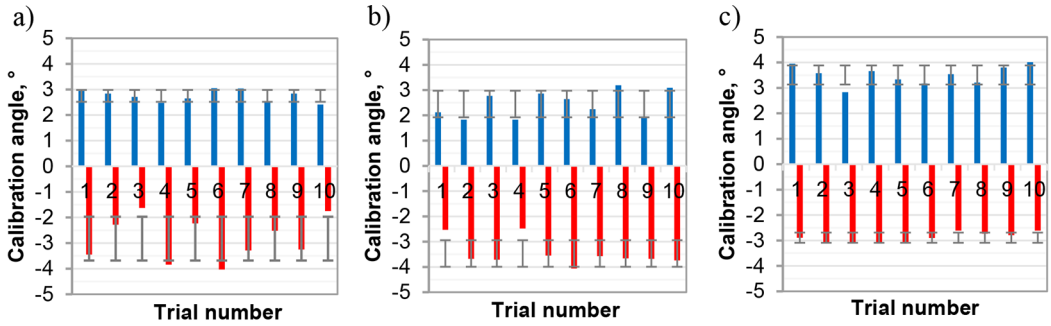


Figure 8. Yaw angles obtained during daytime dynamic calibration of the front-corner radar sensors NR_1 (blue) and NR_2 (red) at: (a) location A, (b) location B, and (c) location C. Error bars denote one standard deviation calculated across the ten repeated values for each sensor at each location

structured proving ground can support stable but shifted convergence. Locations A and B increased the negative NR_2 deviation and pair separation, while Location C reduced NR_2 dispersion but shifted NR_1 upward. These patterns show the interaction between the installed radar geometry, coordinate convention, and location-specific reflector distribution.

Three principal observations arise from the radar results. First, static calibration produced a persistent, approximately mirrored sign pattern, while unequal magnitudes indicated residual pair asymmetry associated with the installation and reference chain rather than random on-road variation. Second, the direct dynamic comparison remained close to the static result for NR_1 but increased the negative magnitude reported for NR_2. Third, the environment changed the balance of the pair: Location A provided the closest average mirror symmetry but poor NR_2 repeatability, Location B produced the largest residual average asymmetry, and Location C produced the greatest average separation while yielding the most compact NR_2 distribution. Although every angle satisfied the ±5-degree limit, these differences demonstrate why radar calibration quality cannot be reduced to a binary threshold decision.

Mathematical interpretation of calibration results

The calculations used the repeated values underlying Figures 4–8. The interpretation below reports the resulting averages, sample standard deviations, coefficients of variation, shifts from the static reference, radar-pair descriptors, and tolerance utilization. Equations are not repeated here because their definitions are given in the preceding mathematical framework.

For the static camera calibration, the average yaw, pitch, and roll angles were 0.623°, 0.367°, and 0.529°, respectively. Their standard deviations were 0.096°, 0.106°, and 0.072°, corresponding to coefficients of variation of 15.5%, 28.8%, and 13.6%. These results define the controlled reference state, but they also show why repeatability should be assessed numerically rather than inferred solely from the test environment.

During daytime dynamic calibration, Location B produced the lowest within-scenario dispersion: the standard deviations were 0.029° for yaw, 0.033° for pitch, and 0.031° for roll, with coefficients of variation between 5.0% and 7.4%. Its average angles were nevertheless shifted by approximately –1.014°, –1.026°, and –1.018° from the static yaw, pitch, and roll averages. Location A showed moderate dispersion, whereas Location C was clearly less repeatable in roll: its roll standard deviation reached 0.175°, and its average roll was displaced by –1.185° from the static reference. Therefore, the calculations distinguish repeatability from agreement with the static baseline: Location B was highly repeatable, but its calibration state remained systematically different from the static result.

At night, pitch became less negative than in the corresponding daytime scenarios, while roll became more negative at all three locations. The largest change occurred at Location C, where average roll moved from –0.656 degrees by day to –1.042 degrees at night. The nighttime roll standard deviation was 0.195 degrees, the shift from the static roll reference was –1.571 degrees, and the average used 52.1% of the magnitude of the negative roll limit. Nighttime yaw had coefficients of variation of 32.2% at Location A and 39.1% at Location B; because the associated yaw

averages were relatively small, these normalized values must be interpreted together with the absolute standard deviations. All camera averages remained acceptable, but the combined shift, dispersion, and tolerance utilization identify nighttime roll at Location C as the least robust MFK result.

In the six-trial static radar series, the average yaw angles were 2.230° for NR_1 and –1.883° for NR_2. The corresponding standard deviations were 0.371° and 0.147°. In the direct dynamic comparison, the averages changed to 2.462° and –3.285°. Thus, NR_1 moved only 0.232° from its static average and became more repeatable, with a standard deviation of 0.099°, whereas NR_2 shifted by –1.402° and had a standard deviation of 0.216°. The average separation of the pair increased from 4.113° to 5.747°, and the mirror-symmetry index increased from 0.347° to 0.823°. Dynamic calibration therefore preserved the positive orientation of NR_1 but amplified the negative offset of NR_2.

The location-specific calculations confirm that reflector geometry affected the two radars differently. Location A had average values of 2.749° and –2.827° and the best average mirror symmetry, with an index of 0.078°, but NR_2 was poorly repeatable there: its standard deviation was 0.856°, and its coefficient of variation was 30.3%. Location B produced the least symmetric average pair, with an index of 1.019°, and NR_2 used 69.3% of its tolerance. Location C produced the largest pair separation, 6.399°, because NR_1 increased to an average of 3.508° while NR_2 remained at –2.891°. At this location, NR_1 used 70.2% of its tolerance, the highest utilization in the radar dataset, whereas NR_2 was comparatively repeatable with a standard deviation of 0.201°.

All radar averages remained below the ±5° project limit. However, the calculations show three distinct properties that a pass/fail criterion cannot distinguish: the persistent opposite-sign mounting tendency, the unequal static-dynamic response of the two sensors, and the location-dependent trade-off between repeatability, mirror symmetry, and distance from the tolerance boundary.

Under the assumptions of independent and approximately normally distributed repeated estimates specified for Equation (8), the 95% Student's t confidence intervals further distinguish random dispersion from systematic calibration displacement. For the ten-trial static MFK series, the intervals were 0.554° to 0.692° for yaw, 0.291° to 0.443° for pitch, and 0.478° to 0.581° for roll. The calculated daytime Location B averages were –0.391° for yaw, –0.659° for pitch, and –0.489° for roll, with corresponding intervals of –0.412° to –0.370°, –0.683° to –0.635°, and –0.511° to –0.467°. The absence of overlap is consistent with scenario-dependent systematic displacement; however, non-overlap of marginal confidence intervals is not a direct hypothesis test of the difference between scenarios and is therefore interpreted as supporting descriptive evidence rather than as a standalone significance test. Expressing tolerance utilization as a remaining angular margin provides an additional engineering interpretation: NR_1 at Location C retained 1.492° to the +5° radar limit, NR_2 at Location B retained approximately 1.535° to the –5° limit, and nighttime MFK roll at Location C retained approximately 0.958° to the –2° limit. Consequently, threshold compliance alone does not imply equivalent calibration robustness; statistical precision, displacement from the controlled reference, and residual tolerance margin must be considered jointly.

The numerical analysis confirms that the differences visible in the figures are systematic and practically interpretable. For the MFK, dynamic calibration produced consistent shifts from the static reference, and reduced illumination primarily affected roll, with the largest degradation at nighttime Location C. For the front radar pair, the opposite-sign pattern was already present under static conditions, while dynamic calibration changed the magnitudes in a sensor- and location-dependent manner. Calibration quality should therefore be assessed based on the central result, absolute dispersion, change from the static baseline, pair consistency, and tolerance utilization, rather than from threshold compliance alone.

DISCUSSION

The results indicate that dynamic calibration is operationally more representative of the conditions in which the sensors are used, but its output is also more strongly conditioned by the observation environment [17, 18]. Static EOL calibration estimates sensor pose relative to a controlled station geometry, whereas dynamic calibration derives orientation from visual or reflective features observed during vehicle motion [19, 22].

The calculated differences between dynamic-scenario averages and the static reference show that the procedures need not converge to the same reported calibration state [27, 30]. Because no independent ground truth was available, this finding establishes procedural and environmental dependence, rather than superior absolute accuracy of either method.

The persistent positive NR_1 and negative NR_2 estimates in both static and dynamic conditions should be interpreted in the context of the mirrored sensor installation and coordinate convention. Opposite signs are compatible with mirror mounting; the relevant evidence of residual pair asymmetry is the unequal magnitude, non-zero mirror-symmetry index, and unequal response to SDK [27, 30]. Persistent inter-sensor asymmetry is also relevant to broader sensor-failure monitoring and diagnostic consistency assessment [31]. The project analysis associates these effects with the complete tolerance and reference chain, including the radar module, mounting bracket, bumper, guide profile, body structure, nominal azimuth, and achieved installation geometry [32]. Therefore, the calculated offsets, separation, and symmetry index show that the two front-corner sensors cannot be treated as perfectly identical mirrored channels under operational conditions.

Comparison of the two procedures further suggests that dynamic calibration is appropriate when the objective is to characterize the calibration state produced during actual vehicle operation, particularly for sensors the estimation process of which relies on natural scene structure [20, 21]. This interpretation is consistent with current research on targetless, automated, and online calibration, which uses naturally occurring features and motion to reduce dependence on dedicated targets [19, 20]. However, operational representativeness should not be equated with ground-truth accuracy: an environment-dependent solution can be highly repeatable while remaining shifted from a controlled baseline, as demonstrated by daytime Location B.

At the same time, the results show that dynamic calibration is not universally better in the sense of producing values closer to zero or to the static baseline; it is more operationally informative, but also more environment-sensitive [22, 30]. The differences among Locations A, B, and C demonstrate that reflector and scene structure affect both the average result and repeatability. Industrial or semi-open surroundings can introduce broad, unbalanced responses, whereas a structured test-track environment can produce a compact distribution for one sensor while maintaining a systematic pair bias. The calculated averages, standard deviations, and coefficients of variation make these distinctions explicit [29].

The MFK results require a parallel interpretation. The ALS series provided the controlled target-based reference, while daytime Location B produced the lowest within-scenario dispersion. However, across the complete dynamic dataset the reported calibration state changed with location and illumination, most clearly in roll. This behavior is consistent with modality-specific evidence: automotive radar is insensitive to visible light but remains affected by reflector geometry, clutter, multipath propagation, and interference [23], whereas low illumination, adverse weather, optical contamination, and contrast loss can degrade camera-based perception and ADAS performance [9, 24, 33]. The largest combined shift and dispersion occurred for the MFK roll at nighttime Location C.

The environmental dependence observed in this study is therefore consistent with the sensing physics. At night, MFK results showed larger deviations and weaker repeatability, especially for roll, which is plausibly associated with reduced availability of stable lane, horizon, and perspective cues [24]. For radar, the dominant factor was not illumination but the geometry and balance of stationary reflective structures. This interpretation is consistent with the technical SDK conditions, radar-technology literature, and vehicle-level sensor-degradation studies [23, 25]. The mathematical descriptors show that environmental degradation can appear as a systematic change in the average, an increase in dispersion, or both; it should not be interpreted as random variation alone [29].

From a validation perspective, compliance with a calibration threshold is necessary but insufficient for a complete assessment [34, 35]. Calibration quality has several dimensions: acceptance-limit compliance, repeatability, agreement with a defined procedural reference, sensor-pair consistency, and relevance to the intended operating environment [20, 21, 28, 29]. All camera and radar averages were formally acceptable, yet the tolerance-utilization values and dispersions differed substantially. Two passing results can therefore have markedly different robustness and

remaining margin, even when neither can be classified as an absolute pose error without external ground truth.

Taken together, the results support treating static and dynamic calibration as complementary rather than mutually exclusive procedures [22, 30]. Static calibration remains indispensable for establishing controlled geometry, loading known adaptation values, and detecting gross installation or assembly errors [27]. Dynamic calibration characterizes how the estimator behaves in realistic use, including its interaction with vehicle motion, visual observability, and the environmental reflector field [21]. An industrial validation strategy should therefore extend beyond an EOL pass/fail decision and include repeated dynamic checks in representative environments, with direct comparison to a documented static baseline where practicable. This is especially important for mixed-modality perception systems, because camera and radar respond differently to illumination, scene structure, clutter, and weather, while sensor-fusion performance depends on consistent cross-sensor geometry and modality-aware validation [23, 36, 37].

Practical implications for production, validation, and service

The primary premise of this work was the direct comparison of static and dynamic calibration methods. Static calibration is widely applied at end-of-line stations and in service workshops, where it requires dedicated targets, positioning equipment, and controlled geometry. The observed calibration behavior of the dynamic procedure, including high repeatability in selected scenarios and estimation under the actual sensor-vehicle-environment interaction, supports its further investigation as a potential alternative for selected factory operations. This represents a production-feasibility proposition rather than a recommendation to replace the currently approved procedure.

An optimized dynamic workflow could reduce the dependence on stationary calibration infrastructure and potentially shorten the production process. However, task-resolved cycle time, throughput, and resource utilization were not measured in the present study; therefore, a claim that dynamic calibration is faster is not yet quantitatively supported. Likewise, the absence of an independent metrological ground truth precludes a conclusion that it is more accurate. The present results show that dynamic calibration can be operationally representative and, under selected conditions, highly repeatable, while its absolute accuracy and production-time benefits remain to be demonstrated.

The next stage should therefore comprise a controlled factory pilot using an independent pose reference and direct measurements of cycle time, throughput, first-pass yield, retry and failure rates, robustness, and resource demand across multiple vehicles, installation tolerances, software variants, operators, and environmental conditions. Replacement of static calibration should be considered only if the optimized dynamic procedure demonstrates predefined non-inferiority or superiority criteria for calibration quality, safety, robustness, and production efficiency. Such optimization and validation constitute a natural application-oriented continuation of the present research.

Until such evidence is available, the results do not justify an immediate formal modification of current maintenance or service procedures. They do, however, support supplementary validation based on the recorded calibration state, environmental conditions, proximity to acceptance limits, and repeatability where required.

A further implication concerns calibration development. Research is increasingly directed toward targetless, automated, and online approaches, including AI-assisted optimization and scene-based estimation using natural environmental features [38]. Future work should investigate adaptive calibration frameworks that combine camera, radar, and inertial information, quantify the uncertainty of estimated extrinsic parameters, and identify conditions under which visual features or radar reflector geometry are insufficient for reliable dynamic recalibration [39].

Study limitations

The findings should be interpreted considering several limitations. The experimental campaign was conducted on a single vehicle with one project-specific sensor installation, and therefore did not capture between-vehicle variation, production tolerances, software-version effects, or operator-related variability. In addition, no independent metrological ground truth was available, so the reported differences describe calibration

states relative to the controlled static procedural reference rather than the absolute accuracy of either procedure. Environmental conditions were characterized qualitatively, rather than through continuous measurements of illuminance, meteorological parameters, road-surface properties, complete vehicle-speed histories, or radar-reflector distribution. Finally, production-related metrics such as task-resolved calibration time, throughput, infrastructure utilization, and implementation cost were outside the scope of the present study. Taken together, these constraints limit the generalizability of the results and prevent a clear causal separation of route geometry, environmental observability, calibration-algorithm response, true changes in installed sensor pose, and production-process efficiency.

CONCLUSIONS

This study showed that dynamic calibration characterizes the sensor calibration state under actual vehicle operating conditions, whereas static calibration characterizes the state relative to a controlled station geometry. The two procedures produced systematically different estimates, and the results confirmed that environmental observability is an intrinsic input to the dynamic estimator. In the absence of an independent pose reference, dynamic calibration should be described as operationally representative rather than intrinsically more accurate.

For the MFK subsystem, calibration behavior was evaluated using yaw, pitch, and roll. All static and dynamic observations remained inside the project-specific acceptance intervals, but the calculations revealed clear differences in deviation from the static baseline and in within-scenario repeatability. Daytime Location B had the lowest dispersion, while roll became increasingly sensitive under degraded visual conditions and reached its largest shift and standard deviation at nighttime Location C. Therefore, a binary acceptance result cannot distinguish a highly repeatable calibration state from one that is strongly dependent on scene structure or illumination.

For the front corner radar pair, all static and dynamic yaw observations satisfied the ±5 degree project limit. The results nevertheless showed an approximately mirrored sign pattern with residual unequal magnitude: NR_1 remained positive, NR_2 remained negative, and the direct dynamic comparison increased the negative magnitude reported for NR_2. Location-specific calculations further demonstrated that repeatability, mirror symmetry, pair separation, reference agreement, and tolerance margin did not improve simultaneously. Radar calibration quality therefore depends on the installed and referenced sensor geometry together with the spatial distribution of reflectors in the calibration environment.

Three principal conclusions follow. First, dynamic calibration is valuable when the objective is to characterize the calibration solution obtained in use, because it incorporates the actual sensor-vehicle-environment interaction. Second, environmental context must be treated as a calibration variable rather than only as an external disturbance, because it affects both the average result and the repeatability of repeated trials. Third, validation should not rely solely on threshold compliance; it should also quantify static-dynamic deviation, absolute dispersion, repeatability, sensor-pair consistency, and utilization of the relevant acceptance limit.

For industrial applications, the findings support a combined calibration strategy. Controlled static calibration remains necessary as an EOL baseline, for loading compensation values, and for detecting gross installation deviations, while dynamic validation is required to assess the calibration state generated during representative vehicle operation. Future research should incorporate independent metrological ground truth, a complete uncertainty budget, repeated routes across a wider range of weather and road conditions, and cross-sensor analysis. Such a design would allow procedural bias to be separated from route-specific observability and from true changes in the installed sensor pose. AI-assisted estimation, online recalibration, and sensor-fusion-based consistency monitoring are promising directions, provided that the system can recognize when environmental conditions are unsuitable for a reliable calibration update. These findings reflect the currently validated calibration process, rather than a permanent technological requirement. The comparative results provide a rationale for a dedicated production study of dynamic calibration as a potential alternative to static end-of-line calibration in selected applications. Such a study should incorporate independent pose ground truth and direct measurements of cycle time, throughput, failure rates, and resource demand. Therefore, the present dataset supports further feasibility

assessment and optimization, but does not establish that dynamic calibration is already faster or more accurate than the static procedure.

Acknowledgements

This work was supported by Volkswagen Poznan sp. z o.o. The investigations were carried out within the Implementation Doctorate Program of the Ministry of Education and Science, and were supported by the Research Subsidy: 0415/SBAD/0374.

REFERENCES

1. Neumann T. Analysis of advanced driver-assistance systems for safe and comfortable driving of motor vehicles. Sensors (Basel). 2024;24(19):6223. https://doi.org/10.3390/s24196223
2. Prochowski L, Szwajkowski P, Ziubiński M. Research scenarios of autonomous vehicles, the sensors and measurement systems used in experiments. Sensors (Basel). 2022;22(17):6586. https://doi.org/10.3390/s22176586
3. Qian H, Wang M, Zhu M, Wang H. A review of multi-sensor fusion in autonomous driving. Sensors (Basel). 2025;25(19):6033. https://doi.org/10.3390/s25196033
4. Nowakowski M, Kurylo J, Dang PH. Camera based AI models used with LiDAR data for improvement of detected object parameters. In: Modelling and Simulation for Autonomous Systems. Cham: Springer; 2024. p. 287–301. https://doi.org/10.1007/978-3-031-71397-2_18
5. Yang S, Choi M, Han S, Choi KH, Kim KS. 4D radar-camera sensor fusion for robust vehicle pose estimation in foggy environments. IEEE Access. 2024;12:16178–88. https://doi.org/10.1109/ACCESS.2023.3345375
6. Markiewicz P, Długosz M, Skruch P. Review of tracking and object detection systems for advanced driver assistance and autonomous driving applications with focus on vulnerable road users sensing. In: Mitkowski W, Kacprzyk J, Oprzędkiewicz K, Skruch P, editors. Trends in Advanced Intelligent Control, Optimization and Automation. Cham: Springer; 2017. p. 224–37. https://doi.org/10.1007/978-3-319-60699-6_22
7. Ascential Technologies. The importance of ADAS calibration for collision prevention and road safety [white paper]. Grand Rapids (MI): Ascential Technologies; 2025.
8. Durmaz O, Cevikalp H. Targetless radar-camera calibration via trajectory alignment. Sensors (Basel). 2025;25(24):7574. doi: 10.3390/s25247574.
9. Feng S, Cai X, Li L, Wang W, Ying S. A review of research on vehicle detection in adverse weather environments. J Traffic Transp Eng Engl Ed. 2025;12(5):1452–83. https://doi.org/10.1016/j.jtte.2025.06.001.
10. Rahman T, Liu A, Cheema D, Chirila V, Charlebois D. ADAS reliability against weather conditions: quantification of performance robustness. In: Proceedings of the 27th Enhanced Safety of Vehicles Conference; 2023 Apr 3–6; Yokohama, Japan. Washington (DC): NHTSA; 2023. Report 23-0306.
11. Peiris S, Newstead S, Berecki-Gisolf J, Chen B, Fildes B. Quantifying the lost safety benefits of ADAS technologies due to inadequate supporting road infrastructure. Sustainability. 2022;14(4):2234. https://doi.org/10.3390/su14042234
12. Bosurgi G, Pellegrino O, Ruggeri A, Sollazzo G. The effects of ADAS on driving behavior: a case study. Sensors (Basel). 2023;23(4):1758. https://doi.org/10.3390/s23041758
13. Skórska H. A controlled pilot study in a practice yard to assess driver perception with camera-monitor systems vs. mirrors: eye tracking insights from a pilot study, part one. Transport Problems. 2025;20(1):111–24. https://doi.org/10.20858/tp.2025.20.1.10
14. Skokan A. Field testing of ADAS technologies in naturalistic driving conditions. Vehicles. 2025;7(4):135. https://doi.org/10.3390/vehicles7040135
15. Masello L, Castignani G, Sheehan B, Murphy F, McDonnell K. On the road safety benefits of advanced driver assistance systems in different driving contexts. Transp Res Interdiscip Perspect. 2022;15:100670. https://doi.org/10.1016/j.trip.2022.100670.
16. Cheni VK. The evolution and impact of advanced driver-assistance systems. Int J Sci Adv Technol. 2025;16(1):1–13. https://doi.org/10.71097/IJSAT.v16.i1.2915
17. Luu VT, Cai YL, Tran VH, Chiu WC, Chen YT, Huang CC. RC-AutoCalib: an end-to-end radar-camera automatic calibration network. In: Proc IEEE/CVF Conf Comput Vis Pattern Recognit (CVPR); 2025. p. 6700-9. https://doi.org/10.1109/CVPR52734.2025.00628
18. Liu X, Deng Z, Zhang G. Targetless radar-camera extrinsic parameter calibration based on track association. Sensors (Basel). 2025;25(3):949. https://doi.org/10.3390/s25030949
19. Cheng L, Cao S. Online targetless radar-camera extrinsic calibration based on the common features of radar and camera. In: Proc IEEE Natl Aerosp Electron Conf (NAECON); 2023. p. 294-9. https://doi.org/10.1109/NAECON58068.2023.10366051

20. An P, Ding J, Quan S, Yang J, Yang Y, Liu Q, et al. Survey of extrinsic calibration on LiDAR-camera system for intelligent vehicle: challenges, approaches, and trends. IEEE Trans Intell Transp Syst. 2024;25(11):15342–66. https://doi.org/10.1109/TITS.2024.3419758
21. Qiu Z, Martínez-Sánchez J, Arias-Sánchez P, Rashdi R. External multi-modal imaging sensor calibration for sensor fusion: a review. Inf Fusion. 2023;97:101806. https://doi.org/10.1016/j.inffus.2023.101806
22. Yao S, Guan R, Huang X, Li Z, Sha X, Yue Y, et al. Radar-camera fusion for object detection and semantic segmentation in autonomous driving: a comprehensive review. IEEE Trans Intell Veh. 2024;9(1):2094–2128. https://doi.org/10.1109/TIV.2023.3307157.
23. Ersü C, Petlenkov E, Janson K. A systematic review of cutting-edge radar technologies: applications for unmanned ground vehicles. Sensors (Basel). 2024;24(23):7807. https://doi.org/10.3390/s24237807
24. Chen Z, Zhang Z, Su Q, Yang K, Wu Y, He L, et al. Object detection for autonomous vehicles under adverse weather conditions. Expert Syst Appl. 2026;296:128994. https://doi.org/10.1016/j.eswa.2025.128994
25. Stowe L, Alcazar J, Huggins S, Beale GT, Palmer M, Doerzaph ZR. Safety implications of potential advanced driver assistance systems sensor degradation. Washington (DC): NHTSA; 2025. Report DOT HS 813 740.
26. Gulino MS, Fiorentino A, Vangi D. Prospective and retrospective performance assessment of advanced driver assistance systems in imminent collision scenarios: the CMI-Vr approach. Eur Transp Res Rev. 2022;14:3. https://doi.org/10.1186/s12544-022-00527-4
27. IAV GmbH. Anleitung – Statische Kalibrierung: Nanoradar E3 1.2. Version 1.0.6 [internal document]. 2023.
28. Joint Committee for Guides in Metrology. Guide to the expression of uncertainty in measurement – Part 6: developing and using measurement models. JCGM GUM-6:2020. Sèvres: BIPM; 2020.
29. Taylor BN, Kuyatt CE. Guidelines for evaluating and expressing uncertainty in NIST measurement results. Gaithersburg (MD): NIST; 1994. Technical Note 1297.
30. IAV GmbH. Anleitung – Service Dynamic Kalibrierung SDK: Nanoradar E3 1.2. Version 1.0.8 [internal document]. 2023.
31. Matos F, Bernardino J, Durães J, Cunha J. A survey on sensor failures in autonomous vehicles: challenges and solutions. Sensors (Basel). 2024;24(16):5108. https://doi.org/10.3390/s24165108
32. Volkswagen AG. Technical guidelines for front nanoradar EOL calibration issues in project VW_02/03 [internal document]. 2025 Jun 27.
33. MITRE Corporation. A study on real-world effectiveness of model year 2015–2023 advanced driver assistance systems. McLean (VA): MITRE; 2025. Report PR-25-0114.
34. Ullah Z, da Silva JAC, Nunes RR, Reis A, Filipe V, Barroso J, et al. Performance of advanced rider assistance systems in varying weather conditions: a systematic review. Vehicles. 2025;7(4):105. https://doi.org/10.3390/vehicles7040105
35. Ding S, Abdel-Aty M, Yang S, Wang C, Horrey WJ. Assessing the safety effectiveness of advanced driver assistance systems. J Safety Res. 2025;95:345–61. https://doi.org/10.1016/j.jsr.2025.10.014
36. Xu C, Sankar R. A comprehensive review of autonomous driving algorithms: tackling adverse weather conditions, unpredictable traffic violations, blind spot monitoring, and emergency maneuvers. Algorithms. 2024;17(11):526. https://doi.org/10.3390/a17110526
37. Feng D, Haase-Schütz C, Rosenbaum L, Hertlein H, Glaeser C, Timm F, et al. Deep multi-modal object detection and semantic segmentation for autonomous driving: datasets, methods, and challenges. IEEE Trans Intell Transp Syst. 2021;22(3):1341–60. https://doi.org/10.1109/TITS.2020.2972974
38. Zhao C, Shi Y, Du Y, Jiang S, Ji Y, Zhao X. A rapid and convenient spatiotemporal calibration method of roadside sensors using floating connected and automated vehicle data. IEEE Trans Intell Transp Syst. 2024;25(9):10953–66. https://doi.org/10.1109/TITS.2024.3366758
39. Ibrahim A, Pointner W, Raab A, Kyamakya K. Exploring radar-camera extrinsic calibration: a deep learning perspective and challenges. WSEAS Trans Comput. 2024;23:289–299. https://doi.org/10.37394/23205.2024.23.29