UWB-Based Localization of Smartphones inside a Vehicle to Prevent Distracted Driving

Kailai Cui
University of Michigan Ann Arbor
Ann Arbor, USA
kailaic@umich.edu

Ke Sun
University of Michigan Ann Arbor
Ann Arbor, USA
kesuniot@umich.edu

Kang Geun Shin
University of Michigan Ann Arbor
Ann Arbor, USA
kgshin@umich.edu

Abstract

Smartphone-induced driver distraction is a major contributor to road accidents. A promising mitigation approach is to localize the smartphone within the vehicle and automatically disable distracting functions when the phone is detected in the driver's seat while the car is in motion. However, existing solutions either require additional hardware deployment or lack the accuracy needed for seat-level, multi-phone localization. We propose `UWB-PTrac` (UWB-based Phone Tracking), a real-time seat-level smartphone localization system that operates using only the existing single Ultra-Wideband (UWB) anchor infrastructure commonly found in modern vehicles for keyless entry. `UWB-PTrac` analyzes multipath profiles extracted from the Channel Impulse Response (CIR) to distinguish phone locations between the driver and passenger seats. It runs continuously in the background, using low-power IMUs to detect cross-seat movement and trigger UWB ranging, then classifies the phone's location using a machine learning model. To resolve spatial ambiguity in CIR features, we propose two strategies to enhance signal propagation asymmetry: (1) optimizing UWB anchor placement during vehicle design, and (2) applying a directional RF shield in retrofit scenarios. We further improve robustness by leveraging temporal correlations across consecutive CIR frames to filter noise and stabilize peak features. Implemented as an Android app, our extensive evaluations show that `UWB-PTrac` achieves over 96% and 93% seat-level localization accuracy using an anchor placed above the steering column and a shielded anchor, respectively. End-to-end driving tests confirm consistent real-time, multi-phone localization performance in dynamic conditions.

CCS Concepts

• **Human-centered computing → Ubiquitous and mobile computing**.

Keywords

UWB, smartphone localization, in-cabin sensing, driving safety

ACM Reference Format:
Kailai Cui, Ke Sun, and Kang Geun Shin. 2026. UWB-Based Localization of Smartphones inside a Vehicle to Prevent Distracted Driving. In *ACM/IEEE International Conference on Embedded Artificial Intelligence and Sensing Systems (SenSys '26), May 11–14, 2026, Saint Malo, France.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3774906.3800472

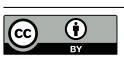

SenSys '26, Saint Malo, France

ACM ISBN 979-8-4007-2309-4/26/05
https://doi.org/10.1145/3774906.3800472

1 Introduction

Smartphone-induced driver distraction is a major contributor to road accidents, raising serious safety concerns [2]. Studies indicate that drivers engaged with smartphones are 3.6 times more likely to be involved in crashes due to delayed reaction times and diverted attention [15]. In 2022 alone, 3,308 fatalities and an estimated 289,310 injuries were attributed to crashes involving phone-distracted drivers [2]. To mitigate these risks, 29 U.S. states have enacted laws banning handheld cellphone use while driving, and 49 states prohibit text messaging behind the wheel [4].

In addition to legislation, smartphone manufacturers have introduced driving mode features designed to reduce distraction [3, 20], and insurance telematics systems offer incentives for safe driving based on phone usage patterns [34, 42]. Current commercial solutions, such as iOS Driving Focus and Android Auto, primarily rely on GPS or Bluetooth to detect in-vehicle use. However, they can not determine whether the phone belongs to the driver and often require manual activation, making them prone to user oversight or intentional circumvention.

A more effective strategy is to *automatically* disable certain phone features, such as texting, *only when* the device is in the driver's seat. Since smartphones can be moved around during a trip (e.g., handed between occupants or retrieved from a different seat), *the system must continuously and accurately track each phone's location at seat-level granularity.* This ensures that restrictions are applied precisely, preventing passengers from being limited and blocking drivers from bypassing restrictions by using another phone.

Building such a robust, real-time in-cabin smartphone localization system remains challenging. Prior efforts have explored various approaches, each with key limitations. For instance, [46] uses acoustic ranging via car speakers, achieving 95% accuracy. However, it requires Bluetooth connection with the car speakers, thus supporting the localization of only one smartphone. [12] leverages Bluetooth Received Signal Strength Indicator (RSSI) and IMU data to classify phone movements, but it depends on a known initial position and fails under complex movements like shaking and moving back-and-forth.

Ultra-Wideband (UWB) technology, increasingly available in modern vehicles, offers promising capabilities for in-cabin smartphone localization. Automotive OEMs such as BMW, Tesla, and Hyundai [8, 23, 40] have adopted UWB-based keyless entry systems to defend against relay attacks [16, 17] and to enable convenience features like proximity unlocking [6, 19]. We aim to repurpose this existing infrastructure to support accurate, real-time localization of multiple smartphones inside the vehicle.

The key challenge is that existing UWB keyless entry systems typically include a single in-cabin UWB anchor with one antenna

[31, 38], which is insufficient for traditional localization techniques. State-of-the-art (SOTA) UWB localization solutions requires either multiple anchors using Time-Difference-of-Arrival (TDoA) [11, 28, 41] or a single anchor equipped with multiple antennas leveraging Angle-of-Arrival (AoA) [7, 10, 18, 48]. These hardware assumptions do not hold in current vehicle deployments. This mismatch motivates a more constrained yet practical question: *how can we localize in-cabin smartphones using only a single UWB anchor?*

We propose UWB-PTrac, a system that leverages UWB *Channel Impulse Response* (CIR) measurements captured by the smartphone's UWB module for seat-level self localization. As signals propagate from the vehicle's interior UWB anchor to a smartphone, they reflect off cabin surfaces, such as the dashboard, seats, and windows, creating rich multipath profiles. The key insight is that the pattern of these reflections varies with the phone's location, and these spatial differences are captured in the CIR. By analyzing these signatures, UWB-PTrac can distinguish the phone's position at seat-level granularity. To support multiple devices, a round-robin protocol manages anchor access, enabling real-time localization across smartphones.

To build a robust and practical system without requiring any hardware modifications to the vehicle infrastructure, we must address the following challenges. First, UWB-PTrac must operate using the existing in-cabin UWB anchor, which offers limited sensing resolution and varies in placement across different vehicle models. These anchors are often installed along the vehicle's central axis to ensure signal coverage throughout the cabin and its surroundings [13, 14]. However, the CIR data, which captures the multipath profile of the environment, has a temporal resolution of 1ns [1, 43] (30cm in signal path length at the speed of light). This resolution is insufficient to capture fine-grained asymmetries, such as the presence of the steering wheel, that distinguish the driver's seat from the front passenger's seat. Thus, centrally placed anchors tend to generate similar CIR signatures for both seats due to spatial symmetry. To address this, we consider two complementary approaches. In design-time deployment scenarios, where UWB-PTrac can guide anchor placement prior to vehicle manufacturing, we identify optimized anchor locations that amplify asymmetries in CIR signatures across seats. For retrofit deployments in existing vehicles with fixed anchor placement, we introduce a lightweight, 3D-printed RF shield that integrates seamlessly with the anchor. This shield introduces artificial asymmetry in multipath propagation, improving CIR discriminability without altering the vehicle's electronic hardware.

Second, the vehicle cabin is a highly dynamic environment during driving. Drivers move their arms while steering, and both occupants may interact with their smartphones or reposition them, introducing significant motion artifacts. These variations can introduce dynamic multipath and distort the CIR and degrade localization accuracy. To ensure robustness under motion, we develop a motion-resilient feature extraction method that leverages the temporal correlation between consecutive CIR frames. This method tracks the inter-frame evolution of CIR peaks to construct a stable profile (e.g., by retaining temporarily vanished peaks) and performs localization only when the phone remains steady.

Third, UWB-PTrac is implemented as a smartphone app to perform self-localization in real time. This design must be energy-efficient and scalable to multiple devices. However, continuous UWB ranging is power-intensive, with packet reception consuming around 160 mW [5, 35]. To minimize energy overhead, UWB-PTrac employs an event-driven ranging strategy that triggers UWB measurements only when necessary, specifically, during significant phone movements such as cross-seat transitions or when the device is picked up. We utilize the phone's low-power IMU to monitor motion and develop a lightweight location-change detection algorithm. This approach significantly reduces energy consumption and scales efficiently to multi-device scenarios, as each phone initiates UWB ranging independently only upon substantial movement.

We implemented UWB-PTrac as an Android app that performs processing of CIR data and executes the self-localization algorithm directly on the smartphone in real time. We use two UWB modules (DWM3001CDK [1]): one emulates the smartphone's UWB transceiver, while the other serves as the in-cabin anchor. We evaluate two complementary anchor deployment strategies designed to enhance signal propagation asymmetry: (1) the steering-top anchor, placed above the steering wheel column, and (2) the shielded anchor, positioned near the rearview mirror with a half-cylindrical RF shield that introduces directional propagation. We evaluated UWB-PTracś classification accuracy across different phone placements on the driver and passenger seats, varying orientation and position to reflect natural usage. UWB-PTrac achieves > 95% accuracy using the steering-top anchor and > 93% using the shielded anchor. In an end-to-end (e2e) tests during city driving, UWB-PTrac correctly triggers UWB localization in over 94% of cross-seat movements and achieves 96% and 91% real-time classification accuracy using the steering-top and shielded anchor, respectively, while localizing two smartphones simultaneously.

In summary, we make the following contributions:

- Leveraging temporal continuity in CIR measurements to construct stable multipath profiles and suppress transient disturbances caused by environmental changes or smartphone movements;
- Proposing two complementary methods to resolve CIR ambiguity: (1) identifying optimized anchor placements in design-time deployments, and (2) introducing a lightweight directional RF shield in retrofit scenarios, enabling accurate seat-level classification using a single UWB anchor with one antenna;
- Designing a lightweight phone motion detection method based on low-power IMU sensing to trigger UWB localization only when significant movement (e.g., cross-seat transfer) is detected;
- Implementing and evaluating UWB-PTrac comprehensively, demonstrating > 95% classification accuracy using the steering-top anchor, > 93% using the shielded anchor, and consistent performance in dynamic, multi-device scenarios.

2 Background and Challenges

This section introduces the background of in-cabin smartphone localization using UWB CIR and presents preliminary insights that reveal the key challenges in designing UWB-PTrac.

2.1 In-Cabin UWB CIR for Phone Localization

Recent advancements in vehicular sensing and digital key systems have led to the integration of UWB transceivers into modern vehicles [8, 23, 40]. Their primary function is to enable secure, proximity-based access control by exchanging UWB signals with smartphones

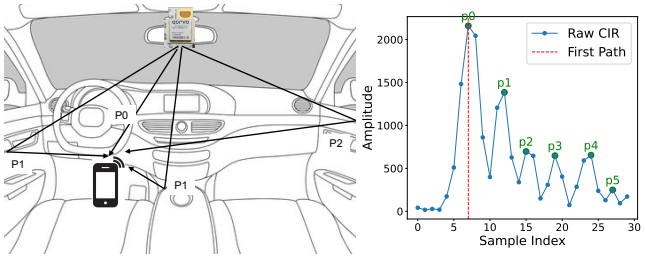

Figure 1: CIR peaks represent signal reflections from vehicle interior.

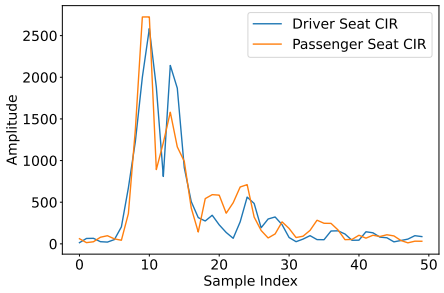

Figure 2: Two CIRs from driver and passenger seat.

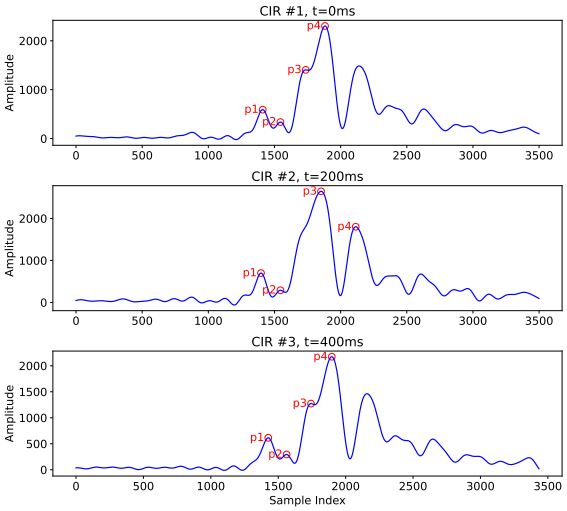

Figure 3: Three consecutive CIRs (200 ms apart) showing the temporary vanishing and reappearance of the third peak.

or smart keys. They generally operate by determining whether a smartphone or key is outside, inside, or near the vehicle, supporting secure access during unlocking and ignition, using basic distance measurements derived from UWB signals [8, 38, 40].

In contrast, UWB-PTrac repurposes this existing UWB infrastructure to achieve seat-level smartphone localization, enabling detection of driver smartphone usage during driving. To support it, UWB-PTrac exploits the phone-captured raw CIR signals as the primary feature for spatial localization within the vehicle cabin.

Fig. 1 illustrates our preliminary setup using a commodity UWB anchor development kit (DWM3001CDK) [1], with the anchor positioned near the rearview mirror and the smartphone serving as the receiving device. The module supports two-way ranging and exposes the CIR via USB interface for post-processing. CIR captures the time-domain response of the transmitted signal, representing its multipath propagation through the vehicle cabin. Each CIR frame contains 1,016 complex samples, and Fig. 1 shows a 50-sample segment around the first path index. The x-axis in Fig. 1 denotes time index in nanoseconds, and the y-axis represents the magnitude of the CIR sample, reflecting the relative signal strength of each path.

Fig. 1 shows a typical CIR when the smartphone is located on the driver's seat. The first peak (p0) corresponds to the direct line-of-sight (LOS) path. Subsequent peaks represent reflected paths: p1 arises from early reflections off nearby surfaces such as the center console and front doors, while p2, p3, and p4 reflect off more distant regions such as the driver's side door, vehicle floor, rear seating area, or trunk. These CIR traces encode rich, location-dependent multipath characteristics. UWB-PTrac leverages these features to train learning-based models capable of distinguishing seat-level positions with high accuracy.

2.2 Challenges of In-Cabin Phone Localization

UWB-PTrac aims to reuse the existing in-vehicle UWB infrastructure without requiring additional hardware, but this design goal introduces several technical challenges as follows.

Challenge (1): Spatial ambiguity due to number and placement of the UWB anchor. Most existing vehicles are equipped with only a single in-cabin UWB anchor [38], which is sufficient for identifying whether the smartphone/key is in cabin, but limits spatial diversity for fine-grained CIR interpretation. When the anchor is centrally located (e.g., at the rearview mirror in Fig. 1), symmetry in the vehicle layout exacerbates ambiguity. Fig. 2 compares two CIRs measured from the driver and the front passenger seats. Despite vehicles' structural differences (e.g., presence of the steering wheel on the driver's side), the CIRs exhibit nearly identical peak features. This is due to the limited time resolution of the commercial UWB transceivers, which typically sample CIRs at 1 ns intervals, equivalent to a spatial resolution of approximately 30 cm [1, 43]. Within a 2–3 ns window (60-90 cm), reflections from multiple nearby surfaces (e.g., steering wheel, dashboard, and driver side door) merge into a single peak in CIR. Consequently, left and right seating regions may produce indistinguishable CIR profiles.

UWB-PTrac *needs to resolve such ambiguity using only a single UWB anchor without deploying additional in-cabin UWB anchors.*

Challenge (2): Interference caused by human body movement. In real-world driving scenarios, both drivers and passengers introduce motion artifacts that significantly impact CIR stability. Drivers frequently move their arms while steering, while passengers often interact with their smartphones. In addition, both occupants exhibit large body movements during vehicle maneuvers such as turns or sudden stops. Furthermore, occupants commonly adjust their seat positions for comfort, resulting in changes in in-cabin structure, human posture and relative device position.

These actions cause noticeable perturbations in CIR measurements, especially in weaker multipath components. Fig. 3 illustrates an example of this phenomenon: the third peak in the first CIR disappears in the second frame and then reappears in the third frame. Such transient changes in peak presence highlight the need for robust tracking and smoothing strategies to handle frame-to-frame signal variability caused by human motion.

UWB-PTrac *needs to generalize across a wide range of occupant motions and in-cabin layout variations.*

Challenge (3): Smartphone battery life. In UWB-PTrac, smartphone functions as both the UWB signal receiver and the processing unit for localization inference. However, continuous UWB reception and UWB-PTrac processing can lead to significant battery drain, making the system impractical. The UWB module we use consumes about 160 mW during packet reception [5, 35], while Bluetooth consumes 50–70 mW during active transmission [24]. In practice, smartphones minimize UWB active time. For example, the UWB-based car Digital Key system use Bluetooth LE for presence detection and wake the UWB radio only for a quick distance verification [6].

UWB-PTrac *must, therefore, incorporate an intelligent triggering mechanism for phone localization only when necessary.*

3 System Design

This section describes the system design of UWB-PTrac. Fig. 4 shows the e2e workflow of UWB-PTrac, which is implemented as a smartphone app. First, the smartphone uses its onboard IMU sensors to intelligently decide when to initiate self-localization (Challenge (3), Section 3.1). Then, it processes the received UWB CIR to perform seat-level localization, ensuring robustness against motion artifacts (Challenge (2), Section 3.2). Finally, we propose a method to create signal asymmetry while minimizing modifications to the existing UWB anchor infrastructure, thereby resolving CIR ambiguity (Challenge (1), Section 3.3).

3.1 Detection of Seat-Level Phone Location Changes

UWB-PTrac aims to maintain accurate seat-level localization of the smartphone while minimizing energy consumption. Once the phone's initial location is established, continuous localization is unnecessary. UWB-PTrac leverages low-power IMU sensors to opportunistically trigger UWB-based localization only when it detects the phone has moved across seats. Unlike prior work that detects general phone movement to trigger localization [29, 49], UWB-PTrac targets fine-grained detection of cross-seat movements within a vehicle cabin. Our design goal is to avoid false negatives, which would result in undetected seat changes and ultimately localization and tracking failures, while minimizing false positives that incur unnecessary localizations.

To meet this requirement, UWB-PTrac employs a lightweight IMU-based detection model that differentiates cross-seat phone movements from common motion sources such as normal vehicular dynamics and in-hand phone manipulations. To guide the model design, we first conducted real-world experiments to characterize IMU signatures under three scenarios: 1) cross-seat phone movement; 2) phone manipulation while held in hand, and 3) in-hand phone movement during vehicle acceleration.

Each scenario was performed 40 times inside a moving vehicle along urban and suburban roads. For cross-seat movement, we tested transitions both during vehicle stops and while the car was in motion, including cases where the user reached across the cabin, passed the phone to a neighboring seat, or moved the phone to and from cup-holder area. For in-hand manipulation, we performed typical phone usage actions such as unlocking, scrolling, or adjusting the phone's orientation. For the vehicle-induced condition, users held the phone loosely in a natural posture during acceleration, braking, or turns without actively moving it.

Our design is based on three core observations.

(1) Cross-seat phone relocations typically occur over short durations. We empirically chose a window size of 1.2 seconds with a step size of 0.2 second for IMU processing, with data sampled at 50Hz. This is sufficient to capture the full duration of most cross-seat movements, while minimizing the inclusion of unrelated motion.

(2) Motion generated by the vehicle itself (e.g., acceleration or cornering) rarely produces strong gyroscope activity when the phone is stationary. In contrast, user-initiated movements (e.g., picking up or moving the phone) result in significant rotational changes. UWB-PTrac first filters out vehicle-only motion by applying a gyroscope magnitude threshold to detect phone pick-up events.

(3) Cross-seat movements tend to be brief, abrupt, and generate high accelerometer variability, distinguishing them from more subtle in-hand manipulations. UWB-PTrac uses the standard deviation of accelerometer magnitude within the detection window to further discriminate between in-seat handling and actual relocation. A conservative threshold is chosen based on the empirical experiments (see Fig. 5) to suppress false positives while ensuring reliable detection of seat-level transitions.

3.2 Motion-Resilient Phone Localization

Once triggered, UWB-PTrac processes the raw CIR data to extract multipath features for seat-level smartphone localization.

CIR Preprocessing: We begin by identifying the index of the first significant CIR peak (fp_{index}), typically corresponding to the LoS path from the UWB anchor to the smartphone. Given the constrained in-cabin environment, we narrow the analysis window to a segment around the LoS path, specifically $[\text{fp}_{\text{index}} - 10, \text{fp}_{\text{index}} + 40]$, which corresponds to propagation distances within approximately 12m. To improve spatial resolution, we upsample the CIR waveform by a factor of 64 using interpolation techniques, enabling finer peak discrimination [9, 27].

CIR Feature Extraction: As discussed in Section 2, CIR peaks characterize reflections from in-cabin structures. The relative positions (index differences) and amplitudes of these peaks encode the propagation delay and signal attenuation of multipath components.

We use a set of features proposed in [27]. UWB-PTrac extracts the top $n = 4$ CIR peaks based on amplitude and sorts them by index. For each non-primary peak, it computes the index offset and amplitude ratio relative to the first peak. These features reflect the spatial arrangement and reflectivity of nearby surfaces. Fig. 6 shows the features extracted from the top four peaks of a CIR.

Peaks beyond the fourth are typically the result of long-path reflections, often from the vehicle floor, trunk, or exterior structures. They are excluded as these late components exhibit lower amplitude and higher variability. In addition to CIR-based features, we also include the ranging distance between the anchor and smartphone, which provides a coarse location prior. The complete set of features used in classification is summarized in Table 1.

Robust Inter-frame CIR Peak Tracking: While peak indices are key features for location classification, they are susceptible to disruption from occupant motion or environmental changes.

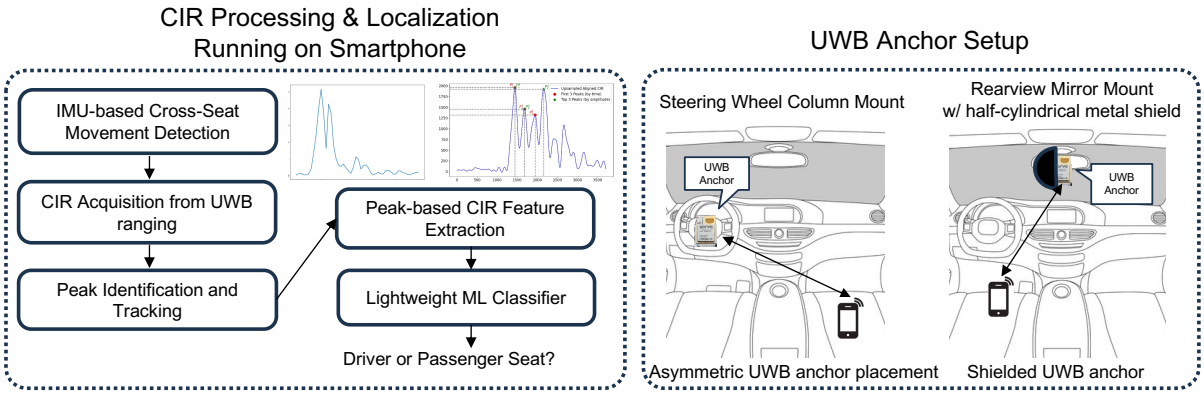

Figure 4: System overview.

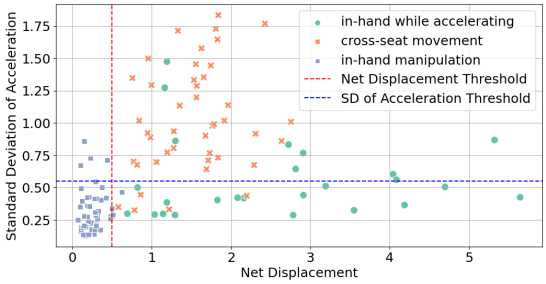

Figure 5: Net displacement vs. standard deviation of acceleration in 120 movement events.

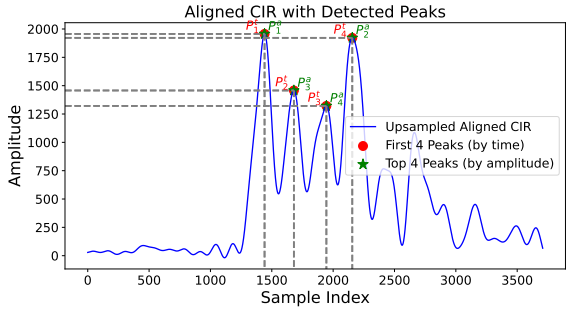

Figure 6: Features based on top 4 peaks.

Table 1: Features used for classification.

Features description
Ranging distance between anchor and smartphone
Index separation from first peak (first n peaks by time)
Amplitude ratio to first peak (first n peaks by time)
Index separation from highest amplitude peak (top n by amplitude)
Amplitude ratio to highest amplitude peak (top n by amplitude)
Number of detected peaks
Index of highest amplitude peak
Amplitude of highest peak

These disruptions may cause peaks to temporarily vanish, reappear, or shift between consecutive CIRs, even when the phone remains stationary. Fig. 3 shows an example of this phenomenon: three CIRs are captured 200 ms apart while the phone is held at the same location, but the third peak in the first CIR disappears in the second frame and reappears in the third frame. This variability arises from two common effects: (i) temporary occlusion of reflection paths by body parts (e.g., a driver's arm), and (ii) resolution limitations when multiple reflection paths arrive within a narrow time window (e.g., < 1ns), leading to peak merging or separation across frames [9]. Because `UWB-PTrac`'s classification features rely directly on peak stability, such fluctuations degrade the consistency of extracted feature vectors. To mitigate this, `UWB-PTrac` implements **inter-frame peak tracking**, which stabilizes peak detection by preserving temporarily missing peaks and smoothing out transient variations. This method leverages the strong resemblance in peak positions across consecutive CIRs when the phone's location and the in-cabin propagation environment are quasi-static.

For each new CIR measurement, local maxima are identified beyond the first peak index using an amplitude threshold. Each detected peak is then compared against a record of previously tracked peaks. To avoid incorrect matches with adjacent peaks, if a new peak lies within a tolerance window of an existing peak, it is treated as the same and its position is updated. Since two peaks are separated by at least 2 indices (128 after upsampling by a factor of 64), we empirically set tolerance window to 100 indices to accommodate minor peak jitter while ensuring that the peak is not matched with an adjacent peak. This window is wide enough to accommodate minor peak jitter from small motions, yet narrow enough to avoid incorrect matches with adjacent peaks. If no match is found, a new peak entry is added. Unmatched old peaks are retained against brief disappearance. Fig. 7 shows three consecutive CIR measurements, where the algorithm successfully preserves the third peak despite its temporary disappearance in CIR #2.

We analyze the stability of our extracted peak indices with and without inter-frame peak tracking by using consecutive five CIR measurements from our testing dataset (details in Section 4.1). Fig. 9 shows results of three inter-peak time difference features: Δt_1, Δt_2, and Δt_3, representing the index separations between the first peak and the second, third, and fourth peaks, respectively. The standard deviation of Δt_3 drops from 127.1 (≈0.6 m) to 82.7 (≈0.4m) indices after tracking is applied, with similar improvements observed for Δt_1 and Δt_2. This enhanced stability improves classification robustness under motion artifacts (see detailed experiments in Section 4.2.2), especially when a passenger is present.

However, when the smartphone changes position, CIR peak indices need to reflect this change. If old peaks are blindly preserved, they may distort the localization output. `UWB-PTrac` needs to distinguish between temporary peak disappearance and peak shifting due to change in the smartphone's location. Our empirical observations suggest that a genuine phone movement typically causes

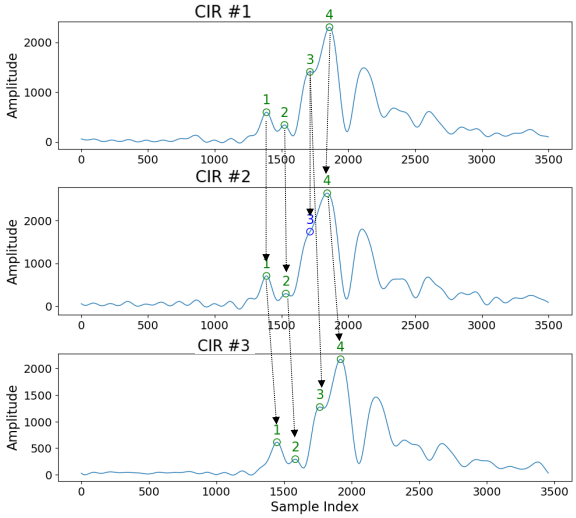

Figure 7: Three consecutive CIRs with peaks tracked.

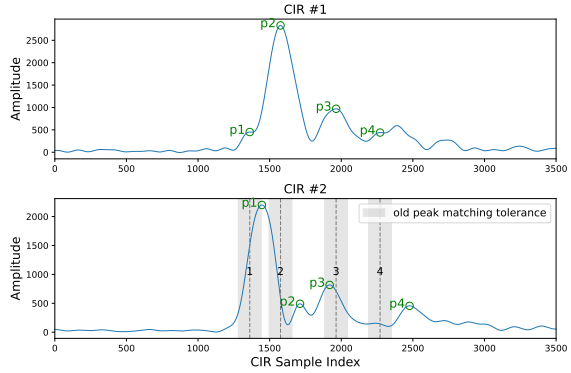

Figure 8: Two peaks become unstable in CIR 2.

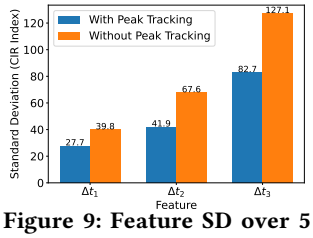

Figure 9: Feature SD over 5 consecutive CIRs.

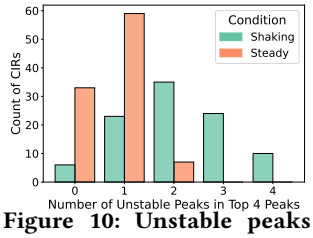

Figure 10: Unstable peaks when shaking vs. steady.

multiple peaks to shift beyond the matching tolerance, reflecting a true change in the spatial configuration of reflective paths, as shown in Fig. 8, peak 2 and peak 4 of the second CIR has shifted beyond the matching tolerance.

In contrast, during stationary conditions, most peaks remain stable across frames. We verify this by collecting 200 CIRs, with 100 the phone is stationary and another 100 during deliberate phone movement, respectively. Fig. 10 shows the distribution of unstable peaks in these two scenarios, where the unstable peaks mean that the CIR peak indices exceed the matching tolerance. We record the number of unstable peaks in each frame. This measurement helps `UWB-PTrac` distinguish true location changes from transient noise, ensuring both responsiveness and resilience in the presence of motion. When more than one of the first 4 peaks becomes unstable, the CIR is marked unreliable for localization and will be skipped.

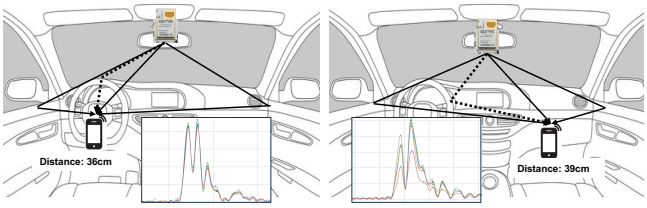

Figure 11: Two locations with overlapping CIR and distance.

We summarize the CIR peak tracking pipeline as follows.

(1) **Detect and Compare:** For each new CIR frame, detect all candidate peaks and compare them against a list of previously tracked peaks.
(2) **Assess Stability:** Count the number of tracked peaks that have shifted beyond a predefined matching tolerance window, indicating a significant change in the multipath profile.
(3) **Update or Skip:** If there are more than one shifted peaks in a frame, the frame is marked as **unreliable**, and the localization is skipped to avoid classifying during motion. Otherwise, the list of tracked peaks is updated with their new positions, and this stable list is used for feature extraction.

Overall, through inter-frame peak tracking, we maintain a steady multipath profile at one location under environment noises and detecting smartphone's actual location change, and only localize the phone when it remains steady, reducing misclassifications when smartphone is moving.

Multi-Phone Localization: The extracted features (Table 1) allow us to train a lightweight ML model that can be run in real time on the phone. Since our extracted CIR features are nonlinear, we choose the Random Forest model (100 estimators) as it handles nonlinear decision boundaries effectively through its ensemble of decision trees. The classifier outputs a binary label indicating whether the phone is on the driver or passenger side.

`UWB-PTrac` can also localize multiple smartphones in a round-robin manner. While the off-the-shelf anchor can only respond to one phone at a time, we implement a wait-and-reconnect protocol in the smartphone app. If the anchor is currently ranging with another phone, the waiting device retries connection after a short delay. In practice, since localization is triggered infrequently, only when the phone moves across seats, the round-robin scheduling does not introduce noticeable ranging delays (Section 4.3).

3.3 Resolving Ambiguity by Creating Asymmetry for In-cabin UWB Sensing

`UWB-PTrac` is designed to repurpose the existing single-anchor UWB infrastructure for in-cabin smartphone localization. A core challenge in this setting arises from the inherent spatial symmetry of the vehicle cabin, particularly between the driver and front passenger seats, which results in similar UWB propagation characteristics and introduces localization ambiguity (see Section 2.2). To overcome this, we explore two complementary approaches: (1) optimizing the anchor location at design time, and (2) applying a directional RF shield in retrofit scenarios. We focus on resolving ambiguity between driver and passenger seat on the front row, as back-row locations can be easily distinguished by their significantly longer UWB ranging distances. Together, these strategies enable

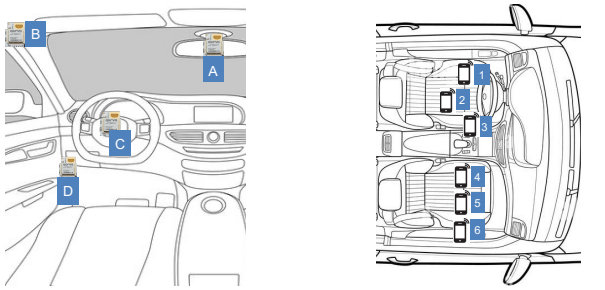

(a) Anchor placements. (b) Smartphone locations.

Figure 12: Anchor and smartphone Locations.

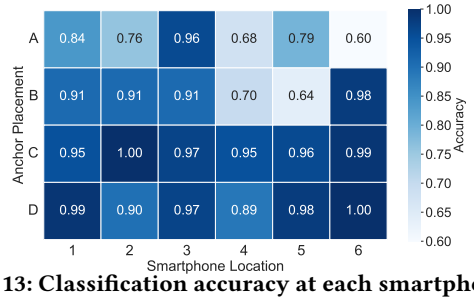

Figure 13: Classification accuracy at each smartphone location across different anchor placements.

robust and low-cost in-cabin smartphone localization without modifying the UWB hardware.

3.3.1 Optimizing UWB Anchor Placement. We first consider a design-time deployment scenario, where UWB-PTrac can influence the placement of the UWB anchor prior to vehicle manufacturing. This allows us to strategically position the anchor to break symmetry in multipath propagation and enhance classification separability.

Fig. 11 illustrates how CIR and ranging measurements from symmetric seat positions can exhibit substantial overlap when the anchor is centrally placed, such as at the rearview mirror. The reflection from steering wheel (dashed line) and the reflection from the driver side door arrives within 1ns, appearing as one peak (the second peak) on the CIR. This motivates repositioning the anchor to amplify geometric asymmetries in signal paths. Specifically, we examine three alternative placements that are practical for wiring and integration in modern vehicles (Fig. 12a):

(1) Rearview mirror (centered, baseline);
(2) Top of driver-side door (to maximize ranging asymmetry);
(3) Steering-wheel column;
(4) OBD-II port in the driver footwell.

For each configuration, we collect CIR and ranging measurements across six smartphone locations in the front row (Fig. 12b) with the phone held at heights ranging from seat cushion level (≈45cm above floor) to passenger chest level (≈90 cm above floor) to reflect common usage scenarios. Each location is sampled with 1, 200 measurements for model training and 600 for evaluation.

Fig. 13 shows the classification testing results with different anchor placement and smartphone locations. When the anchor is placed at the rearview mirror, several driver- and passenger-side regions produce overlapping CIR and distance measurements, degrading classification accuracy, especially in the passenger left and driver right areas. To break this symmetry, we test relocating

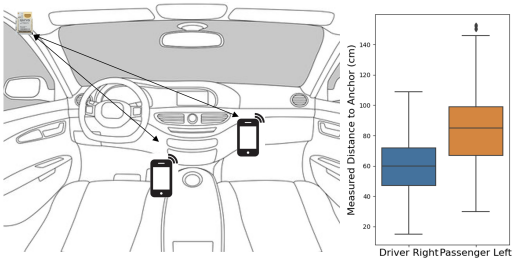

Figure 14: Ranging results at two smartphone locations.

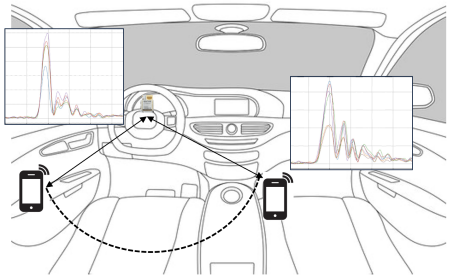

Figure 15: CIR distinguishes 2 overlapped-ranging locations.

the anchor to the top of the driver-side door. This placement introduces clear distance asymmetry between driver and passenger seats. However, UWB ranging alone is unreliable in this environment due to in-cabin multipath effects, especially when multiple reflective surfaces (e.g., seats, windows, dash) lie close to the signal path. As shown in Fig. 14, this placement suffers from overlapping distance distributions across seat regions. CIR features remain ambiguous in such configurations due to merged or indistinct peaks, especially when the reflective paths from symmetric seat regions remain geometrically similar (Fig. 11).

To mitigate this, we explore anchor placements that simultaneously amplify asymmetries in both ranging and CIR profiles. As shown in Fig. 13, positions, such as the steering-wheel column and OBD-II port, show better performance than other positions as they introduce significant geometric differentiation across all front-row positions. These placements shift the overlap region in ranging space to locations where CIR features remain distinguishable, enabling joint disambiguation. For example, with the anchor mounted on the steering-wheel column, driver-left and passenger-right positions may have similar path lengths, but their CIRs differ substantially due to proximity to reflective surfaces like the dashboard or vehicle pillars. As illustrated in Fig. 15, CIR peaks shift based on local reflectors: when the smartphone is close to surfaces such as the door, the non-primary peaks appear earlier due to short reflection paths; when positioned near the cabin center, those peaks are more delayed.

Our results show that anchor positions such as the steering wheel column and OBD-II port provide more distinguishable propagation characteristics due to structural asymmetry. These placements jointly improve the separability of both CIR and distance features and serve as effective alternatives when design-time control is available.

3.3.2 Shielding UWB Anchor. In scenarios where UWB-PTrac is retrofitted into existing vehicles with built-in UWB-based keyless entry systems, relocating the anchor is impractical. FCC filings show that some manufacturers install UWB transceivers along the

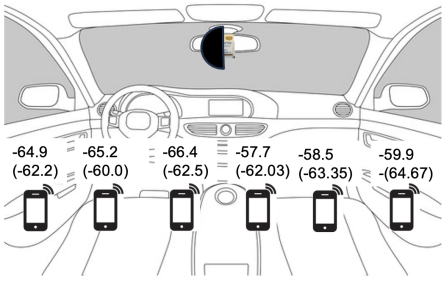

Figure 16: Average RSSI values w/ and w/o shield (dBm).

vehicle's centerline, typically near the rearview mirror or rear brake light, to maximize in-cabin and exterior coverage [13, 14]. However, this centralized placement creates geometric symmetry between driver and passenger seats, leading to overlapping CIR and ranging features that limit localization precision.

To address this challenge without modifying vehicle hardware, we propose a lightweight, non-invasive solution: attaching a small metal shield near the UWB anchor antenna to induce directional asymmetry in signal propagation. Inspired by RF metasurface principles, the shield is designed to attenuate signals toward one side of the cabin while enhancing reflections toward the other, improving separability in multipath signal characteristics.

Shield's Effect on Signal Propagation: Metals such as aluminum are effective at attenuating and reflecting RF signals due to their high conductivity [25, 39]. The curvature of the reflective surface also influences propagation directionality, a concept leveraged in prior beam-shaping systems [22, 45].

To evaluate this effect, we placed a half ping-pong ball (radius 20 mm) coated with aluminum foil adjacent to the driver-facing side of the UWB antenna. We compare signal features with and without the shield across six front-row locations (Fig. 12b) by collecting 1,000 real-world CIRs and RSSIs at each location. The RSSI is computed by the hardware as a function of the CIR power [1].

The results show clear directional asymmetry. As shown in Fig. 16, signals toward the blocked (driver) side are attenuated (RSSI ≈3 dBm lower on average), while reflections toward the open (passenger) side are enhanced (RSSI ≈4 dBm higher on average). CIR profiles also shift accordingly: on the blocked side, early peaks are attenuated as the direct path is obstructed by the shield; on the open side, late peaks are suppressed due to the lack of energy reflected back from the now-blocked region (Figs. 17a and 17b).

To quantify these effects, we extract two features: RSSI and late energy ratio, the proportion of CIR energy arriving more than 1,000 samples (4.7m) after the first peak. As shown in Fig. 18, both features exhibit strong side-dependent separation. However, intra-seat variation caused by phone orientation, anchor distance, and environmental reflections introduces overlap. Even within a single location, RSSI can fluctuate by ±3 dBm, and tail energy ratio can vary widely, potentially degrading classification performance.

These findings raise a key design question: *How can we design a metal shield that consistently produces sufficient signal feature separation across all front-row locations to enable robust driver-passenger smartphone localization?*

Shield Design: To answer this, the shield must satisfy two criteria: (i) induce sufficient signal feature separation across all seat positions, and (ii) minimize the transitional zone, or borderline

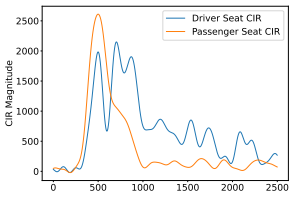

(a) CIRs from driver and passenger seat with directional shield.

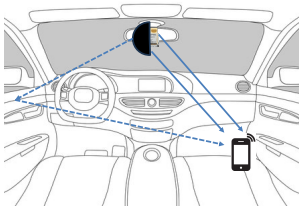

(b) Shield enhances signal to passenger while attenuating signal to driver side.

Figure 17: Impact of directional shielding on CIR and propagation path.

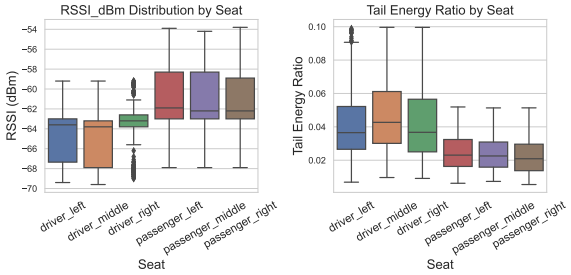

Figure 18: RSSI and Late Energy Ratio at 6 front row locations.

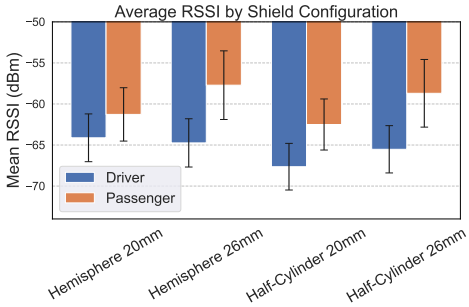

Figure 19: Driver and passenger RSSI values of different shield sizes.

region, where signals diffraction or partial reflection blurs classification boundaries. We explore variations in shield shape and size to achieve these objectives.

Because UWB antennas are omnidirectional, the shield must partially enclose the antenna to block propagation in a specific direction. To support driver-passenger classification, either side can be attenuated; in our implementation, we focus on shielding the driver side. We prototype multiple 3D-printed shield geometries, including hemispheres, half-cylinders, and half-open rectangles, each wrapped in aluminum foil, a lightweight and effective RF reflective material [22, 45]. These shields are mounted near the most typical UWB anchor location, i.e., at the rearview mirror. The design principles can be generalized to other anchor locations.

Shield Size: Larger shields reflect and block more radiated energy, producing stronger asymmetry. However, they must remain compact and visually unobtrusive, ideally fitting within the transceiver's form factor to blend into the vehicle interior. We tested hemisphere and half-cylinder shields with radii of 20 mm and 26 mm (52 mm height for half-cylinders), based on the 28 × 20 mm footprint of the UWB module. For each size, we measure RSSI at all driver- and passenger-side seat locations. As shown in Fig. 19, larger shields

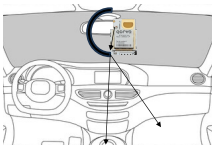

(a) Sector ambiguity due to diffraction.

Shield Shape	Borderline Angular Span
Hemisphere	20° ± 5°
Half-Open Rectangle	20° ± 5°
Half-Cylinder	15° ± 5°

(b) Angular spans under different shield shapes.

Figure 20: Measured diffraction effects caused by shielding.

Table 2: Additional features for shielded anchor setup.

Features description
Received signal strength indicator
Ratio of energy 1000 indices after first peak to total energy
Number of CIR peaks 1000 indices after first peak.
Position of the latest detected peak in the CIR.

yielded greater RSSI separation and more separation between two sides, resulting in higher classification accuracy (see Fig. 21).

Shield Shape: Since the UWB antenna is omnidirectional, partial shielding creates a borderline region, a transitional angular zone between the blocked and open sides. In this region, phones may receive diffracted signals or weakened reflections, resulting in ambiguous RSSI and CIR features for localization (Fig. 20). To quantify the extent of this region, we conducted controlled experiments with the phone placed along a circular arc (−25° to +25° in 5° increments) centered on the UWB antenna. At each angle, we measured signal features and evaluated classification accuracy using a pre-trained model. We then identified the angular range where accuracy dropped below 80%, marking it as the borderline region. We repeated this for three shield shapes: hemisphere, half-cylinder, and half-open rectangle. As shown in Fig. 20, the half-cylinder has the narrowest ambiguity zone, suggesting it best suppresses diffraction while maintaining signal reflection strength.

End-to-end Localization with Shield: To assess how each shield design supports robust end-to-end seat-level classification, we evaluate classification accuracy across six front-row smartphone locations (Fig. 12b). We select four representative designs based on shield size and shape insights: the hemisphere (r = 20mm) as the baseline, hemisphere (r = 26mm) to study effect of shield size, cylinder (r = 26mm) for optimal performance, and a half-open rectangle (26mm × 52mm). For each shield configuration, we collected approximately 1,200 samples per smartphone location, spanning 6 locations shown in (Fig. 12b). The data was randomly split into 80% for training and 20% for testing. Besides the features in Table 1, we include RSSI and features that reflect late peaks (Table 2).

We apply temporal smoothing by averaging the additional features (Table 2) in over a short window of 7 frames (1.4 seconds). Both RSSI and late CIR peaks exhibit high frame-to-frame variability (Fig. 18) as the late peaks are the result of distant, weaker multipath components that more susceptible to fluctuations caused by small movements of the phone. Smoothing over multiple frames helps mitigate such noise and produces more consistent inputs for the classifier. Unlike the steering-top anchor configuration, peak tracking is not applied in the shielded setup, as late-peak-based features do not benefit meaningfully from inter-frame peak consistency. We train a classifier (same model as in Section 3.2) to distinguish between driver-side and passenger-side locations.

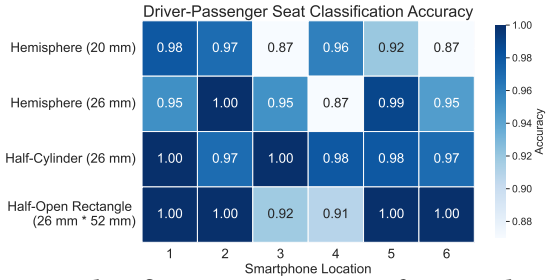

Figure 21: Classification accuracy at six front row locations for four shield designs.

(a) Anchor on top of steering column.

(b) Anchor next to rearview mirror with a shield attached.

Figure 22: Prototypes of two proposed anchor setups.

The hemisphere (r = 20mm) shows sub-90% accuracy at multiple positions, primarily due to the limited feature separation between the blocked and open sides. Both the hemisphere (r = 26mm) and the half-open rectangle perform well when the phone is far from the car's midline but degrade at edge positions near the center, where signal features overlap due to diffraction and partial occlusion. In contrast, the cylinder (r = 26mm) maintains > 97% accuracy at all 6 positions, confirming its small borderline ambiguity region and large enough feature separation. Thus, we adopt the 26 mm half-cylinder as our default shield design.

Guaranteeing the normal usage of UWB anchor: The internal UWB transceiver is primarily designed for vehicle access control applications such as keyless entry [13, 14, 38]. To ensure that our shielding design does not interfere with this functionality, we tested UWB signal coverage by moving a smartphone within a 1-meter radius around the vehicle. Even on the shielded side, the phone was able to maintain UWB connectivity and initiate ranging exchanges with the in-vehicle transceiver. These results confirm that the proposed shield does not compromise the original functionality of the keyless entry system.

4 Evaluation

4.1 Experimentation Methodology

Phones and Cars Used. While recent smartphones are equipped with UWB chips, their operating systems restrict access to low-level APIs necessary for capturing raw CIR data [47]. Additionally, the UWB module integrated into a vehicle's keyless entry system is controlled via the in-vehicle CAN bus, making it inaccessible for the continuous, large-scale data collection required by our experiments. Thus, we use two DWM3001CDK UWB development kits, one is fixed inside the vehicle as the anchor, and the other is connected to a Samsung Galaxy S21 Ultra, emulating the smartphone's UWB module. Our evaluation spans three different vehicles: a BMW 3

Table 3: Classification Accuracy Comparison.

Location	Steering-Top Anchor	Shielded Anchor	Yang et al. [46]
Driver - Left	0.95	0.98	1.00
Driver - Middle	1.00	1.00	1.00
Driver - Right	0.97	0.97	—
Passenger - Left	0.95	0.97	—
Passenger - Middle	0.96	0.93	0.93
Passenger - Right	0.99	0.95	1.00
Rear - Left	1.00	0.95	1.00
Rear - Right	1.00	0.99	1.00

Series (compact sedan), a Honda Accord (mid-size sedan), and a Honda CR-V (SUV).

Implementation. We develop an Android app for UWB-PTrac. It continuously monitors IMU data to detect cross-seat phone movements. Upon movement detection, it triggers UWB ranging and, captures and processes CIR data to infer the phone's seat-level location in real time. We evaluate two anchor deployment configurations separately: (1) steering-top anchor: the UWB anchor is placed at the identified optimal location, i.e., above the steering wheel (Fig. 22 (a)); (2) shielded anchor: the anchor is positioned next to the rearview mirror with a custom half-cylinder RF shield attached (Fig. 22 (b)). For both configurations, we collect approximately 12,000 CIR and ranging samples from the smartphone across 8 predefined regions, including those in Fig. 13 and two rear seat locations (left side and right side). Each sample is labeled as corresponding to either the driver or passenger seat, and a lightweight Random Forest classifier is trained to distinguish between the two.

Localization Accuracy. We first evaluate the localization accuracy of the trained model across different smartphone locations. At each of these 8 predefined regions, we collect 500 test samples with the phone held in two poses, i.e., lifted up and lying flat, and moved within a sphere of approximately $\approx$ 15 cm in diameter to capture realistic variations in position and orientation. The model outputs a binary label, "driver" or "passenger," for each sample. These samples are collected by running continuous UWB ranging. We evaluate performance under the following three experimental settings:

Single-Occupant Setting: One person is present in the vehicle and sits at the corresponding seat being tested.

Multi-Occupant Setting: Two people are present in the vehicle. The driver seat is always occupied, emulating real-world driving conditions. The second occupant sits at the target location (e.g., front passenger or second-row seat) during data collection.

Cross-Vehicle Deployment: To assess generalizability, we evaluate UWB-PTrac across different vehicles. For the asymmetric anchor placement (above the steering column), a model trained on one vehicle is directly applied to two other vehicles without retraining. For the central-shielded anchor configuration, where the anchor placement and nearby structures are vehicle-specific, we train a separate model for each vehicle using the same shield configuration.

For each setting, we compute the driver–passenger classification accuracy by recording the number of correctly classified samples at each test location.

E2E Performance. We evaluate UWB-PTrac's real-time e2e performance during real-world driving scenarios, focusing on two key metrics: (1) the ability to detect cross-seat movements and trigger UWB ranging; and (2) the accuracy of seat-level localization once the phone stabilizes. We conduct three driving sessions while using smartphones: Trip 1 (10 minutes): The phone is actively manipulated in driver or passenger's hand to simulate typical handling behaviors. Trip 2 (15 minutes): Two smartphones are simultaneously moved in and out of the driver's seat, front passenger seat, and cup holder during the drive to test multi-phone handling and movement dynamics. Trip 3 (15 minutes): The phone is repeatedly moved between the driver and passenger seats, focusing specifically on localization performance using the shielded anchor configuration.

Table 4: Classification accuracy in multi-occupant setting across two anchor configurations.

Location	Steering-Top w/ tracking	Steering-Top w/o tracking	Shielded no tracking
Driver - Left	0.97	0.97	0.98
Driver - Middle	0.99	0.99	0.98
Driver - Right	0.95	0.96	0.90
Passenger - Left	0.93	0.85	1.00
Passenger - Middle	0.97	0.94	0.98
Passenger - Right	0.99	0.99	0.87
Rear - Left	1.00	1.00	0.97
Rear - Right	1.00	1.00	0.91

4.2 Seat-level Phone Localization

4.2.1 Accuracy across Locations. As shown in Table 3, UWB-PTrac demonstrates robust performance across all tested positions, achieving an average classification accuracy of 0.978 with the steering-wheel-top anchor and 0.968 with the shielded rearview mirror configuration. Despite reduced spatial diversity, the shielded design maintains strong accuracy, achieving at least 0.93 in all positions, the lowest observed rate across all regions. Compared to the multi-speaker acoustic baseline by Yang et al. [46], UWB-PTrac achieves comparable localization accuracy while using only a single commercial UWB anchor, highlighting its practicality and efficiency. In contrast, Fig. 13 shows that centrally placed anchors without shielding suffer from spatial ambiguity, with classification accuracy dropping below 0.70 in some regions. This comparison underscores the necessity of creating directional asymmetry in signal propagation. UWB-PTrac *addresses this challenge using the steering-top anchor and the shielded anchor configurations, achieving consistently high seat-level classification accuracy.*

4.2.2 Impact of Vehicle Seat Occupants. Next, we evaluate localization robustness when additional vehicle occupants introduce interference. Table 4 shows results for both anchor configurations. For the steering-top anchor, we observe a drop in classification accuracy at the passenger seat's left and middle regions. This is primarily due to signal occlusion caused by the occupant's body, which blocks reflected paths and leads to frequent peak disappearance in the CIR. By using peak tracking algorithm that retains temporarily vanished peaks, UWB-PTrac achieves improved accuracy in these regions. Meanwhile, the driver and rear seat performance remains stable, showing resilience to interference. In contrast, the shielded anchor setup relies more on RSSI and CIR late peak features. These late peaks are typically weaker, long-path reflections that traverse the cabin and are more susceptible to absorption or distortion by the occupant's body. We average CIR features across multiple frames, and omit peak tracking, as it provides minimal benefit. While the shielded anchor configuration maintains high overall accuracy, its

Table 5: Classification Accuracy (Steering-Top Anchor).

Location	Accord (Unseen)	CR-V (Unseen)	BMW (Trained)
Driver - Left	0.98	0.92	0.95
Driver - Middle	1.00	0.99	1.00
Driver - Right	0.97	0.91	0.97
Passenger - Left	0.97	0.90	0.95
Passenger - Middle	1.00	0.91	0.96
Passenger - Right	1.00	1.00	0.99
Rear - Left	1.00	1.00	1.00
Rear - Right	1.00	1.00	1.00

Table 6: Classification Accuracy (Shielded Anchor).

Location	Accord	CR-V	BMW
Driver - Left	0.98	0.95	0.97
Driver - Middle	1.00	0.95	0.95
Driver - Right	0.97	0.91	0.89
Passenger - Left	0.97	0.96	0.90
Passenger - Middle	0.93	1.00	0.97
Passenger - Right	0.95	0.94	0.97
Rear - Left	0.96	0.99	1.00
Rear - Right	0.99	1.00	0.92

Table 7: Deployment Effort Trade-Offs for New Vehicles.

	Steering-Top	Shielded
Anchor Placement	Specific	Flexible
Model Retraining	Not Required	~40 min. data collection

performance slightly degrades at boundary regions such as the driver-right and passenger-right regions. In these areas, occupant-induced signal attenuation causes the distinguishing features to shift across classification decision boundaries.

4.2.3 Generalization to Other Vehicles. Table 5 provides the localization accuracy of our trained model using the steering-top anchor setup across different vehicles. The model was trained exclusively on data from the BMW 3 Series (sedan) and evaluated on two unseen vehicles: the Honda Accord (sedan) and the Honda CR-V (SUV). Interestingly, the Honda Accord achieves even higher accuracy than the BMW, likely due to its less compact interior. Greater spacing between key features, such as the steering wheel, doors, and center console, reduces overlapping multipath reflections, resulting in cleaner and more separable CIR features. In contrast, the Honda CR-V exhibits slightly lower accuracy (0.90-–0.92) at several front-row regions. As an SUV, its larger cabin dimensions introduce increased vertical distances between the smartphone and surrounding reflective surfaces. This alters the multipath profile and introduces unseen CIR patterns, especially in overlapping regions such as the left and right areas of both the driver's and passenger's seats, leading to degraded performance. Table 6 presents results using the shielded anchor setup. Unlike the steering-top model, each vehicle is evaluated using a model trained on its own data to account for variations in cabin layout and signal propagation. This is necessary because the decision boundaries based on RSSI and CIR features are sensitive to the exact placement and orientation of the shielded anchor near the rearview mirror.

The results confirm robust performance across all vehicles, with only minor accuracy drops in boundary regions such as the driver-right and passenger-right seats, where signal attenuation and feature shifts occur due to occupant or structural interference. The evaluation of our system across these three distinct vehicle types (compact sedan, mid-size sedan, and SUV) provides a representative baseline for generalizability to a wider range of common passenger vehicles.

Table 8: Confusion matrix: movement detection.

Actual / Predicted	Cross-Seat	In-Hand
Cross-Seat	0.94 (TPR)	0.06 (FNR)
In-Hand	0.11 (FPR)	0.89 (TNR)

In practice, UWB-based keyless entry systems differ across vehicles in terms of layout, mounting constraints, and anchor location. Therefore, our goal is not zero-shot generalization, but to demonstrate that adding a directional RF shield to a centrally mounted anchor enables consistently accurate seat-level localization across vehicles with different interior geometries. We summarize the trade-off between the two anchor configurations in Table 7.

4.3 End-to-end Evaluation

4.3.1 End-to-end Accuracy. We conducted three real-world driving sessions to evaluate the IMU-based movement detection. In each session, smartphone was moved between seat regions and manipulated in-hand (see details in Section 4.1). We performed 100 cross-seat movements, including 50 between the driver's seat and passenger seat, and 50 between the driver's seat and cup holder area, as well as 100 in-hand manipulations (e.g., lifting, rotating).

As shown in Table 8, `UWB-PTrac` correctly detected 94 out of 100 cross-seat movements. Of the six missed events, five occurred during movements between the driver's seat and the cup holder, and one during a driver-to-passenger transition. For in-hand manipulations, `UWB-PTrac` correctly ignored 89 out of 100 cases. The 11 false positives were all triggered during vehicle braking or turning, which caused transient IMU signals similar to cross-seat movements. *Overall, the high true positive rate ensures that most genuine movements are captured, minimizing missed localization events. At the same time, the low false positive rate reduces unnecessary localizations, improving system efficiency.*

Following a movement detection trigger, `UWB-PTrac` initiates UWB ranging and collects CIR data every 200ms. To improve robustness, it aggregates predictions over multiple CIR frames. With the steering-top anchor, `UWB-PTrac` make predictions using majority voting over 3 consecutive stable CIRs. With the shielded anchor, `UWB-PTrac` applies temporal smoothing by averaging the extracted features (Table 2) over 7 CIRs.

In the trip using steering-top anchor setup (on the BWM sedan), of 94 triggered localizations, 90 were correctly classified, yielding a classification accuracy of 95.7%. In a separate trip using the shielded anchor configuration (on the Honda sedan), 82 out of 90 triggered localizations were correct, resulting in an accuracy of 91%. This slight performance drop is attributed to the shielded anchor's reliance on late CIR peaks and RSSI, which are more susceptible to occlusion and body-induced attenuation in dynamic driving conditions.

4.3.2 Response Time. We evaluate `UWB-PTrac`'s response Time using data recorded during the driving sessions. With the shielded anchor, `UWB-PTrac` requires 7 frames (1.4 seconds) to make a classification decision. For the steering-top anchor, the variable waiting condition for 3 stable CIRs results in an average of 7.04 CIRs per decision, or approximately 1.41 seconds response time. When two phones are moved simultaneously, `UWB-PTrac` handles one localization at a time. In such cases, one phone must wait for the other to complete its classification. Across all concurrent activation events

Table 9: Comparison of Seat-Level Smartphone Localization Schemes. A = Driver–Passenger Classification Accuracy; B = Robust to Small Movements; C = Fast Response; D = Multi-Smartphone Support; E = Uses Existing Hardware; Y = Yes; N = No. For works with multiple configurations, the lowest reported classification accuracy is shown.

Related Work	A	B	C	D	E
[46] Acoustic ranging w/ speakers	95%	Y	Y	N	Y
[44] Centripetal acceleration w/reference device	>91%	Y	N	Y	N
[26] Pitch and roll dynamics	100%	N	Y	Y	Y
[12] Bluetooth RSSI and phone accelerometer	93%	N	Y	Y	Y
[32] Engine noise and vibration	>80%	Y	N	Y	Y
UWB-PTrac	>93%	Y	Y	Y	Y

(steering-top anchor), the average response time increases to *3.2 seconds*, which still satisfies real-time operational requirements.

5 Related Work

In-cabin Smartphone Localization: Table 9 summarizes the representative techniques used for in-cabin smartphone localization. [46] requires Bluetooth connection between the smartphone and the car's audio system and thus limits its use to the connected phone only. [44] depends on data collected during multiple turn to make accurate seat-level classification. [26] requires the smartphone to maintain its orientation during detection. [12] tracks phone location based on the prior knowledge of initial or prior locations and the movement trajectory. It struggles with complex movements like back-and-forth shifting or shaking. A more recent work [32] achieves an 80% accuracy in the car without using extra hardware, but only in long offline datasets.

Indoor Localization Using UWB CIR: [21] proposes SALMA, a UWB single anchor localization system. It localizes a UWB tag by comparing theoretical CIR profiles calculated from floor plan against the measured CIR. [21] identified ambiguity when using a single basic UWB transceiver as the anchor. They proposed to use a custom multi-antenna UWB transceiver as anchor, achieving a 90% probability of less than 20cm error. [36] proposed a representation of CIR used to evaluate effectiveness of anchor placements in reducing localization ambiguity. [37] proposed SALOS, a single anchor localization system based on CIR, which uses a basic UWB transceiver as anchor and tries to find an optimal anchor position given a floor plan. Unlike these systems, UWB-PTrac targets the vehicle interior, where compact geometry and overlapping multipath make model-based CIR matching impractical. Instead, we formulate in-cabin seat-level localization as a binary classification task.

UWB Sensing in Vehicles: UWB CIR has been used for various in-vehicle sensing tasks. [33] infer the location of a UWB-enabled key around a car using CIR from anchor-to-key ranging. [27] detects car states (e.g., door/window open or closed) using CIRs from 14 transceivers placed inside and outside the vehicle. Their peak-based CIR feature extraction inspired our system design. [30] senses seat occupancy using UWB CIRs. Their system collects CIRs from eight transceivers inside the car and feeds processed CIRs into a multi-input multi-output (MIMO) CNN, achieving 94.6% accuracy. [43] proposes UMusic, a vehicle per-seat occupancy detection system using a high-resolution power delay profile (PDP) extracted from CIR. It achieves accurate (99%) per-seat occupancy detection by setting UWB transceivers around the car such that passengers sit on signal propagation paths. PDP may help identify merged peaks, one cause for driver–passenger seat ambiguity discussed. However, our goal and transceiver setup differ from this. The asymmetric vehicle components (steering wheel, dashboard) are not in the path between anchor and smartphone, and one (smartphone) of the UWB transceivers can move.

6 Limitations

Failure Cases. Our end-to-end evaluation revealed two primary failure cases. First, classification errors are most common at borderline locations (e.g., a central cupholder) where the multipath signature is ambiguous. Real-world accuracy may be higher as users typically hold phones in less ambiguous viewing positions. Second, a user aware of the system could potentially bypass the IMU trigger by moving the phone with deliberate slowness. A practical deployment could address both cases by incorporating a periodic fail-safe check (e.g., every few minutes, regardless of IMU activity).

Impact of Environmental Factors. While our evaluation spanned three representative vehicle types, for a practical deployment, the system's performance would need to be validated on each target vehicle model with different seat location/inclination settings to account for its unique interior materials and levels of in-vehicle clutter. Furthermore, although our multi-occupant tests show robustness to a second passenger, a more systematic evaluation of the signal-shielding effects of different passenger body types or the impact of more occupants may be useful.

Phone API Access. UWB-PTrac relies on access to UWB CIR data. However, current commercial smartphone OSs restrict access to low-level UWB data including CIR [47]. Therefore, the practical deployment of UWB-PTrac is contingent on future platform support from mobile OS vendors to provide API access to this raw data.

User Study and Acceptance. We conducted e2e driving tests to validate UWB-PTrac's performance. For practical deployment, a wider user study is necessary to (1) assess real-world usability and uncover unforeseen edge cases, and (2) gauge user acceptance. Given that our goal is to disable phone functions, potentially against users' will, it is vital to understand the users' perceptions and privacy concerns, as well as the acceptability of UWB-PTrac to drivers.

7 Conclusion

We have presented UWB-PTrac, a real-time in-vehicle smartphone localization system that uses CIR multipath profiling captured by a single UWB anchor. To improve driver-passenger seat discriminability, we explore two complementary solutions to enhance CIR asymmetry: using optimized anchor placement and attaching a lightweight, 3D-printed RF shield. By incorporating temporal smoothing across CIR frames and an IMU-based movement detection mechanism, UWB-PTrac achieves robust seat-level localization under practical deployment constraints. Our evaluation shows UWB-PTrac achieves accurate seat-level classification in real-world conditions, supporting real-time operation with multiple smartphones.

Acknowledgments

The work reported in this paper was supported in part by the National Science Foundation under Grant No. CNS-2245223.

References

[1] 2019. DW3000 Family User Manual. https://www.qorvo.com/products/d/da008154

[2] 2024. Driver Electronic Device Use in 2022. https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/813531.pdf

[3] 2024. Use the Driving Focus on your iPhone to concentrate on the road. https://support.apple.com/en-us/108384

[4] 2025. Distracted driving. https://www.iihs.org/research-areas/distracted-driving

[5] 2025. DWM3001C Data Sheet. https://www.qorvo.com/products/d/da008154

[6] Apple Inc. 2021. Explore UWB-based car keys. https://developer.apple.com/videos/play/wwdc2021/10084. Accessed: 2025-04-16.

[7] Aditya Arun, Shunsuke Saruwatari, Sureel Shah, and Dinesh Bharadia. 2023. XRLoc: Accurate UWB localization to realize XR deployments. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems*. 459–473.

[8] Bmw. 2022. What's the deal with Ultra wideband technology and what will it do for your car? https://www.bmw.com/en/innovation/bmw-digital-key-plus-ultra-wideband.html

[9] Yifeng Cao, Ashutosh Dhekne, and Mostafa Ammar. 2024. UTrack3D: 3D Tracking Using Ultra-wideband (UWB) Radios. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services*. 345–358.

[10] Yanjun Cao, Chenhao Yang, Rui Li, Alois Knoll, and Giovanni Beltrame. 2020. Accurate position tracking with a single UWB anchor. In *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2344–2350.

[11] Witsarawat Chantaweesomboon, Charuwalee Suwatthikul, Supatra Manatrinon, Krit Athikulwongse, Kamol Kaemarungsi, Ratchasak Ranron, and Prapun Suksompong. 2016. On performance study of UWB real time locating system. In *2016 7th International Conference of Information and Communication Technology for Embedded Systems (IC-ICTES)*. IEEE, 19–24.

[12] Chun-Yu Chen and Kang G Shin. 2022. In-vehicle phone localization for prevention of distracted driving. *IEEE Transactions on Mobile Computing* 22, 6 (2022), 3365–3379.

[13] Continental Automotive GmbH. 2021. RF Test Report - Human Exposure: BMW FBD5. https://fcc.report/FCC-ID/KR5FBD5/5296648.pdf. FCC ID: KR5FBD5, Report No.: T46614-00-09FX.

[14] Continental Automotive GmbH. 2021. UWBtrx22 and UWBBLE22 User Manual. https://fcc.report/FCC-ID/KR5UWBTRX22/5954381.pdf. FCC ID: KR5UWBTRX22, Initial version dated 21.09.2021.

[15] Thomas A Dingus, Feng Guo, Suzie Lee, Jonathan F Antin, Miguel Perez, Mindy Buchanan-King, and Jonathan Hankey. 2016. Driver crash risk factors and prevalence evaluation using naturalistic driving data. *Proceedings of the National Academy of Sciences* 113, 10 (2016), 2636–2641.

[16] FiRa Consortium. 2024. Turning UWB-enabled Smartphones into Keys. https://www.firaconsortium.org/resource-hub/blog/turning-uwbenabled-smartphones-into-keys Accessed: 2025-04-16.

[17] Aurélien Francillon, Boris Danev, and Srdjan Capkun. 2011. Relay attacks on passive keyless entry and start systems in modern cars. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*. Eidgenössische Technische Hochschule Zürich, Department of Computer Science.

[18] Feng Ge and Yuan Shen. 2021. Single-anchor ultra-wideband localization system using wrapped PDoA. *IEEE Transactions on Mobile Computing* 21, 12 (2021), 4609–4623.

[19] Riad Ghabra and Keith A Christenson. 2020. Apparatus and method for detecting location of wireless device to prevent relay attack. US Patent 10,645,596.

[20] Google. 2025. Android Auto. https://www.android.com/auto/. Accessed: 2025-04-16.

[21] Bernhard Großwindhager, Michael Rath, Josef Kulmer, Mustafa S Bakr, Carlo Alberto Boano, Klaus Witrisal, and Kay Römer. 2018. SALMA: UWB-based single-anchor localization system using multipath assistance. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*. 132–144.

[22] Sihui Han and Kang G Shin. 2017. Enhancing wireless performance using reflectors. In *IEEE INFOCOM 2017-IEEE Conference on Computer Communications*. IEEE, 1–9.

[23] Hyundai Motor Europe. 2025. Hyundai Digital Key. https://www.hyundai.com/eu/driving-hyundai/owning-a-hyundai/bluelink-connectivity/digital-key.html Accessed: 2025-04-16.

[24] Texas Instruments. 2015. CC2564MODN Bluetooth® Host Controller Interface (HCI) Module Datasheet. https://www.mouser.com/datasheet/2/405/cc2564-269802.pdf. Rev. E, November 2015.

[25] Nuruddeen Iya, Ali Muqaibel, Umar Johar, and Mohamed Adnan Landolsi. 2011. Ultra-wideband characterization of obstructed propagation. In *2011 7th International Wireless Communications and Mobile Computing Conference*. IEEE, 624–629.

[26] Gregory Johnson and Rajesh Rajamani. 2020. Smartphone localization inside a moving car for prevention of distracted driving. *Vehicle system dynamics* (2020).

[27] Avinash Kalyanaraman, Yunze Zeng, Sushanta Rakshit, and Vivek Jain. 2020. CaraoKey: Car states sensing via the ultra-wideband keyless infrastructure. In *2020 17th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*. IEEE, 1–9.

[28] Benjamin Kempke, Pat Pannuto, Bradford Campbell, and Prabal Dutta. 2016. Surepoint: Exploiting ultra wideband flooding and diversity to provide robust, scalable, high-fidelity indoor localization. In *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM*. 137–149.

[29] Donnie H Kim, Younghun Kim, Deborah Estrin, and Mani B Srivastava. 2010. Sensloc: sensing everyday places and paths using less energy. In *Proceedings of the 8th acm conference on embedded networked sensor systems*. 43–56.

[30] Yongsen Ma, Yunze Zeng, and Vivek Jain. 2020. CarOSense: Car occupancy sensing with the ultra-wideband keyless infrastructure. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (2020), 1–28.

[31] Christian McCaig and Simon Ma. 2021. *FCC PART 15, SUBPART F and ISEDC RSS-220 Compliance Test Report*. Test Report R2009181-247 Rev A. Bay Area Compliance Laboratories Corp., 1274 Anvilwood Avenue, Sunnyvale, CA 94089, USA. https://fcc.report/FCC-ID/2AEIM-1614291/5078003.pdf FCC ID: 2AEIM-1614291, IC: 20098-1614291.

[32] Sugandh Pargal, Neha Dalmia, Harsh Borse, Bivas Mitra, and Sandip Chakraborty. 2024. Zero-Configuration Alarms: Towards Reducing Distracting Smartphone Interactions while Driving. *ACM Journal on Computing and Sustainable Societies* 2, 3 (2024), 1–30.

[33] JiWoong Park, Hong-Beom Choi, Young-Bae Ko, and Keun-Woo Lim. 2021. Locate UWB smart keys: smart and faster. In *Proceedings of the 22nd International Workshop on Mobile Computing Systems and Applications*. 169–171.

[34] Jun-geun Park. 2018. Using telematics data to identify a type of a trip. US Patent 9,900,747.

[35] Tommaso Polonelli, Simon Schläpfer, and Michele Magno. 2022. Performance Comparison between Decawave DW1000 and DW3000 in low-power double side ranging applications. In *2022 IEEE Sensors Applications Symposium (SAS)*. IEEE, 1–6.

[36] Sven Ole Schmidt, Marco Cimdins, and Horst Hellbrück. 2019. On the effective length of channel impulse responses in UWB single anchor localization. In *2019 International Conference on Localization and GNSS (ICL-GNSS)*. IEEE, 1–6.

[37] Sven Ole Schmidt, Marco Cimdins, Fabian John, and Horst Hellbrück. 2024. Salos—a uwb single-anchor indoor localization system based on a statistical multipath propagation model. *Sensors* 24, 8 (2024), 2428.

[38] David Schnaufer. 2023. Ultra-Wideband & CCC® DigitalKey. https://www.qorvo.com/design-hub/blog/ultra-wideband-and-ccc-digital-key-enabling-a-better-driving-experience

[39] Dipankar Shakya, Mingjun Ying, Theodore S Rappaport, Hitesh Poddar, Peijie Ma, Yanbo Wang, and Idris Al-Wazani. 2024. Wideband penetration loss through building materials and partitions at 6.75 GHz in FR1 (C) and 16.95 GHz in the FR3 upper mid-band spectrum. *arXiv preprint arXiv:2405.01362* (2024).

[40] Tesla, Inc. 2025. Model 3 Owner's Manual. https://www.tesla.com/ownersmanual/model3/en_us/ Software version 2025.8.

[41] Janis Tiemann, Fabian Eckermann, and Christian Wietfeld. 2016. Atlas-an open-source tdoa-based ultra-wideband localization system. In *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 1–6.

[42] Johan Wahlström, Isaac Skog, and Peter Händel. 2017. Smartphone-based vehicle telematics: A ten-year anniversary. *IEEE Transactions on Intelligent Transportation Systems* 18, 10 (2017), 2802–2825.

[43] Shuai Wang, Yunze Zeng, Vivek Jain, and Parth Pathak. 2025. *UMusic: In-car Occupancy Sensing via High-resolution UWB Power Delay Profile*. Association for Computing Machinery, New York, NY, USA, 116–129. https://doi.org/10.1145/3715014.3722049

[44] Yan Wang, Jie Yang, Hongbo Liu, Yingying Chen, Marco Gruteser, and Richard P Martin. 2013. Sensing vehicle dynamics for determining driver phone use. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*. 41–54.

[45] Xi Xiong, Justin Chan, Ethan Yu, Nisha Kumari, Ardalan Amiri Sani, Changxi Zheng, and Xia Zhou. 2017. Customizing indoor wireless coverage via 3D-fabricated reflectors. In *Proceedings of the 4th ACM International Conference on Systems for Energy-Efficient Built Environments*. 1–10.

[46] Jie Yang, Simon Sidhom, Gayathri Chandrasekaran, Tam Vu, Hongbo Liu, Nicolae Cecan, Yingying Chen, Marco Gruteser, and Richard P Martin. 2011. Detecting driver phone use leveraging car speakers. In *Proceedings of the 17th annual international conference on Mobile computing and networking*. 97–108.

[47] Fusang Zhang, Zhaoxin Chang, Jie Xiong, Junqi Ma, Jiazhi Ni, Wenbo Zhang, Beihong Jin, and Daqing Zhang. 2023. Embracing consumer-level UWB-equipped devices for fine-grained wireless sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 4 (2023), 1–27.

[48] Minghui Zhao, Tyler Chang, Aditya Arun, Roshan Ayyalasomayajula, Chi Zhang, and Dinesh Bharadia. 2021. Uloc: Low-power, scalable and cm-accurate uwb-tag localization and tracking for indoor applications. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 3 (2021), 1–31.

[49] Zhenyun Zhuang, Kyu-Han Kim, and Jatinder Pal Singh. 2010. Improving energy efficiency of location sensing on smartphones. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*. 315–330.