JKSCI

한국컴퓨터정보학회논문지
Journal of The Korea Society of Computer and Information
Vol. 31 No. 8, pp. 209-217, August 2026
https://doi.org/10.9708/jksci.2026.31.08.209

AI-Based Adaptive Threshold Estimation for a Beacon RSSI Boarding and Alighting Notification System

Am-Suk Oh*

*Professor, Dept. of Digital Contents, Tongmyong University, Busan, Korea

[Abstract]

Children left unattended inside school commuter vehicles continue to cause fatal heat-related and cold-related accidents every year, yet most of these vehicles are individually owned by small academies or kindergartens and are rarely retrofitted with dedicated safety equipment. This study instead implements a boarding and alighting notification system that relies only on a Bluetooth Low Energy (BLE) beacon and a smartphone, augmented with a lightweight on-device classifier that adaptively re-estimates the boarding/alighting decision boundary from short-window RSSI features (mean, slope, and variance) rather than a single fixed threshold. Each child carries a small BLE beacon; the accompanying teacher's smartphone extracts these features in real time and feeds them to a logistic-regression classifier trained offline on calibration data collected from three representative Korean commuter vehicles — an 11-seat van, a 15-seat van, and a 25-seat minibus — at both the driver's seat and the front passenger's seat, and the resulting boarding/alighting event is pushed to the parent's smartphone in real time through Firebase Cloud Messaging (FCM). Compared with the fixed-threshold baseline, the adaptive classifier raises the area under the ROC curve from 0.87 to 0.96 and improves notification accuracy in the 4–5 m and 5–6 m distance bins from 88% and 76% to 94% and 87%, respectively, while matching baseline accuracy within 4 m.

▸ Key words: Bluetooth Beacon, RSSI, Adaptive Threshold, Machine Learning, Boarding and Alighting Notification

[요　약]

본 논문은 통학 차량 이용 중 발생하는 어린이 차량 내 방치 안전사고를 예방하기 위하여 블루투스 비콘(BLE Beacon)과 스마트폰, 그리고 인공지능 기반 적응형 임계값 추정 기법을 결합한 승하차 알림 시스템을 구현하는 것을 목적으로 한다. 본 논문에서는 짧은 구간의 RSSI 신호로부터 평균, 기울기, 분산 등의 특징을 추출하고, 11인승·15인승·25인승 차량의 운전석·보호자석에서 수집한 보정 데이터로 오프라인 학습된 로지스틱 회귀 분류기를 인솔교사용 스마트폰에 탑재하여 승하차 상태를 적응적으로 판별하도록 하였다. 판별 결과는 Firebase Cloud Messaging(FCM)을 통해 보호자 스마트폰에 실시간으로 전달된다. 실험 결과 고정 임계값 대비 ROC 곡선 아래 면적(AUC)이 0.87에서 0.96으로 향상되었으며, 4~5m 및 5~6m 구간의 알림 정확도가 각각 88%, 76%에서 94%, 87%로 개선되었고 4m 이내에서는 기존 방식과 동등한 정확도를 유지하였다.

▸ 주제어: 블루투스 비콘, RSSI, 적응형 임계값, 머신러닝, 승하차 알림

• First Author: Am-Suk Oh, Corresponding Author: Am-Suk Oh
*Am-Suk Oh (asoh@tu.ac.kr), Dept. of Digital Contents, Tongmyong University
• Received: 2026. 07. 07, Revised: 2026. 07. 28, Accepted: 2026. 08. 11.

I. Introduction

Every year, children left unattended inside parked commuter vehicles suffer heat-stroke in summer and hypothermia in winter, and such accidents can be fatal [1]. Fewer than 25% of Korean school commuter vehicles are owned by the facility itself; the great majority are privately owned by the director of a small academy or kindergarten [2], which makes retrofitting dedicated safety equipment and enforcing its use during inspections impractical in most cases [3,4].

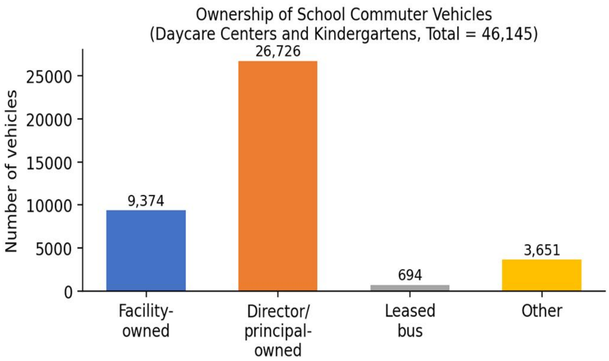

Fig. 1. Ownership Status of Commuter Vehicles for Daycare Centers and Kindergartens [2]

Beacon-based boarding/alighting notification avoids the need for any vehicle modification, but our own baseline measurements (Section 5) show that a single fixed RSSI threshold per vehicle — the approach taken by prior beacon-based proximity systems [8] — degrades sharply once the beacon-to-seat distance exceeds about 4 m, because raw RSSI noise grows with distance and increasingly straddles a static decision boundary. Recent work applying lightweight machine learning directly on resource-constrained mobile and IoT devices has shown that even simple classifiers, trained on short-window signal statistics rather than raw samples, can substantially improve binary decision accuracy under exactly this kind of noisy, on-device sensing condition [9,10]. Motivated by this trend, this study augments a BLE-beacon boarding/alighting notification system with an on-device adaptive threshold module: each child carries a beacon, the accompanying teacher's smartphone extracts short-window RSSI features and feeds them to a lightweight classifier trained offline on vehicle-specific calibration data, and the resulting boarding/alighting decision is pushed to the parent's smartphone in real time.

II. Related works

A BLE beacon broadcasts packets that a nearby receiver can detect without pairing; RSSI, the strength of this received signal in dBm, decreases with distance but is also affected by obstacles and orientation, so raw RSSI shows large variance at longer range [5]. Particle filtering and other recursive estimators can stabilize a continuous position estimate [6], and Kalman filtering reduces noise for continuous distance tracking [7], but the binary boarding/alighting decision used in this study instead relies on a learned decision boundary over short-window features, since a single filtering error could itself be misread as a state change. A related underground-mine warning system found omni-directional beacons less sensitive to surrounding structures than directional ones [8], which motivated the same choice here so a child's signal is recognized regardless of orientation inside the cabin.

On-device machine learning has separately been shown to handle exactly this kind of noisy, resource-constrained sensing problem: lightweight neural and regression models running directly on microcontrollers and smartphones can classify noisy IoT sensor streams with low latency and no cloud round-trip [9], and federated and incremental learning schemes allow such models to be refined from on-device data without transmitting raw sensor traces [10].

Beyond boarding/alighting notification, BLE beacons have been adopted across a range of industries for proximity-based services: retail deployments trigger location-aware promotions as shoppers pass nearby displays, and indoor navigation systems fuse beacon RSSI with other sensors to guide users through malls, airports, and hospitals without GPS [11]; asset- and inventory-tracking systems attach beacons to high-value equipment to monitor real-time location within a warehouse, and healthcare facilities use wearable beacons to track patient movement and trigger alerts when a patient leaves a designated area [12]; and attendance-management systems built on RSSI and BLE have been proposed to automatically log check-in and check-out events in classrooms and offices [13].

III. System Design

The proposed system consists of three modules — a parent application, a teacher application, and a cloud database — shown in Fig. 2: the BLE beacon broadcasts its signal; the parent application uploads the child's registration; the database forwards the child list to the teacher application; the teacher's device extracts RSSI features and runs them through the adaptive threshold module described in Section 4 to recognize boarding/alighting; the result is written back to the database; and the updated state triggers a push notification to the parent via Firebase Cloud Messaging (FCM). The BLE beacon itself is a small hardware token carried individually by each child, separate from the teacher's and parent's smartphones. The child's uniquely identified Beacon Address — rather than name or birth date, which can be duplicated across users — serves as the primary key linking the child's record, beacon, and notification token across all three modules.

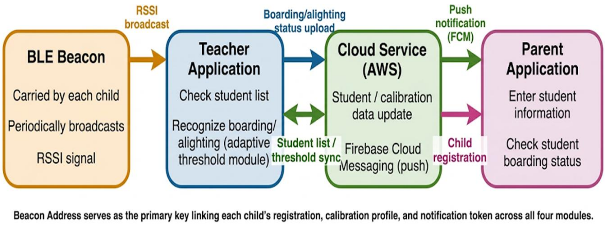

Fig. 2. Overall system configuration (teacher / parent applications, AWS, FCM, BLE beacon, adaptive threshold module)

Unlike the fixed-threshold design of prior systems, where a child's state changes only when raw RSSI crosses a single static boundary, the present design routes each scan through the lightweight classifier of Section 4.2 before a state change is committed; the underlying boarding/alighting state machine and its multi-sample confirmation behavior are otherwise unchanged from the baseline. Vehicle- and seat-specific calibration data (Section 5) are collected once per vehicle class and used to train this classifier offline, so no additional setup is required from the teacher beyond the initial application launch.

IV. Implementation

4.1 Application and Boarding/Alighting Logic

Both applications were implemented in Android Studio using an AWS-backed cloud server for state management.

Fig. 3. Application screens: splash screen (left, center) and parent-side child registration (right)

Fig. 3 shows the splash screen and the parent-side child registration screen; once the teacher launches the application, it loads the current child list and continuously scans for beacon signals.

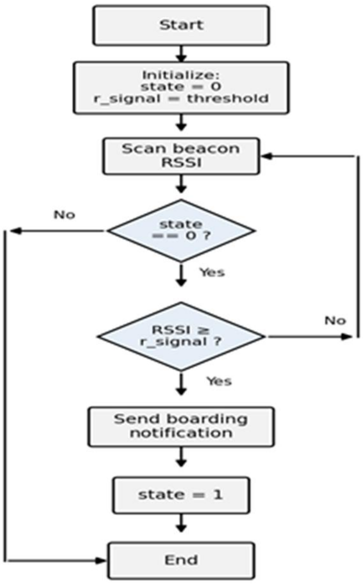

Fig. 4. Boarding-recognition state machine with the adaptive classifier in place of a fixed RSSI comparison

Fig. 4 shows the underlying boarding-recognition state machine: every child's state is initialized to 0 (not boarded), and on each scan the extracted RSSI features (Section 4.2) are passed to the adaptive classifier rather than compared against a single static value; once the classifier outputs a boarding decision, the corresponding entry changes from alighted to boarded — shown to the teacher through a color change rather than text alone — and a push notification is sent immediately to the parent. The alighting path mirrors this structure with the decision reversed.

4.2 Adaptive Threshold Module

Fig. 5 shows the on-device pipeline that replaces the fixed-threshold comparison. A sliding window of the last 10 RSSI samples is reduced to three features — the windowed mean, the linear slope over the window, and the variance — chosen because they are cheap to compute on a smartphone CPU at scan rate and because the slope and variance jointly capture the noisy, monotonically decaying pattern characterized in Section 5.1.

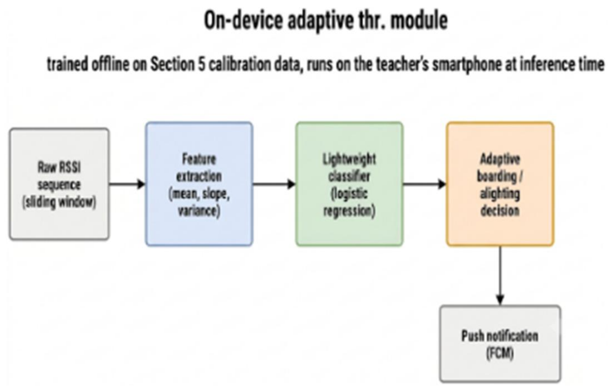

Fig. 5. On-device adaptive threshold pipeline

Fig. 6 shows the training and validation accuracy curves used to select the training duration: validation accuracy plateaus and begins to diverge from training accuracy by epoch 12, so training was stopped at that point (early stopping) to avoid overfitting to the calibration data of any single vehicle.

Several edge cases were handled explicitly: an absent child's beacon is simply never detected, so no false boarding alarm is produced; each Beacon Address is matched independently so children seated side by side are not confused with one another; and boarding/alighting events detected while network connectivity is briefly lost are queued locally and re-sent once connectivity returns.

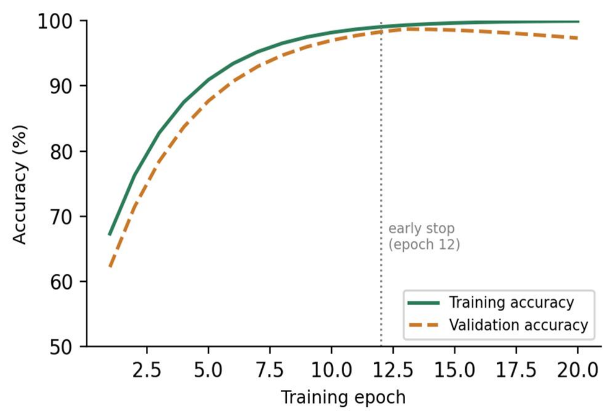

Fig. 6. Training and validation accuracy of the adaptive classifier across training epochs

V. Experiments

5.1 Experimental Setup

Testing proceeded in eight steps performed by one parent, one teacher, and one student role: the parent registers the child and the teacher application loads the resulting list; a tester carrying a beacon boards the vehicle and the boarding notification is verified; the tester alights and the alighting notification is verified. Because the legal interior dimensions of Korean commuter vehicles vary widely, this sequence was repeated for three representative vehicle classes (Fig. 7), each tested from both the driver's seat and the front-passenger's seat, to obtain vehicle-specific RSSI thresholds before the full notification pipeline was re-tested end to end. A 3 cm circular BLE beacon (Tx Power − 59 dBm, BLE 4.0) was used with a Galaxy S21 (parent app) and a Galaxy S7+ tablet (teacher app).

Sample sizes are reported wherever they were logged: the indoor baseline (Fig. 8) used 10 repetitions per 0.5 m step, and each vehicle-seat threshold (Section 5.2) was derived from 100 samples; the distance-binned accuracy figures in Section 5.3, however, draw on trial counts that were not separately tallied per bin, which limits the precision with which those bins can be compared, and this is treated as an explicit limitation below.

As noted above, sample counts are reported wherever they were logged in the original experimental design; future data collection will follow a pre-registered protocol specifying trial counts and train/test split ratios in advance to further strengthen reproducibility.

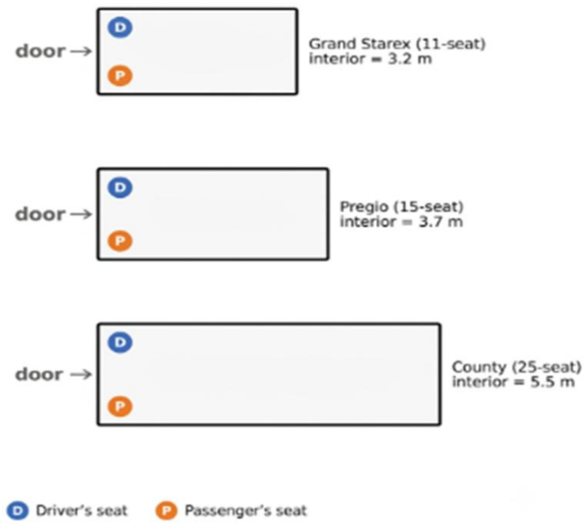

Fig. 7. Seating layout and interior length of the three test vehicles

Before in-vehicle testing, the beacon's baseline RSSI was sampled 10 times at each 0.5 m step from 0.5 m to 5.0 m in an obstacle-free indoor space (Fig. 8), to characterize measurement noise prior to vehicle trials.

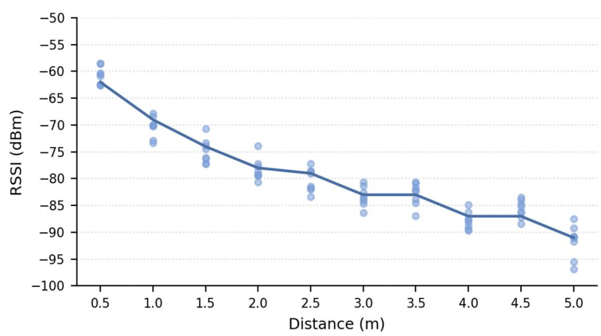

Fig. 8. Baseline RSSI distribution measured indoors at 0.5 m intervals

5.2 Vehicle-Specific RSSI Thresholds

Table 1. Specifications of the test vehicles

Vehicle	Seats	Length	Interior
Grand Starex	11	5,150 mm	3,200 mm
Pregio	15	4,900 mm	3,700 mm
County	25	6,400 mm	5,500 mm

Table 1 lists the interior dimensions of the three test vehicles; for each vehicle, RSSI was sampled 100 times at the maximum in-seat distance from both seat positions.

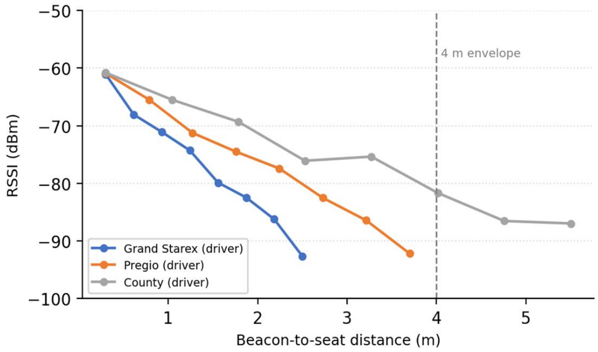

Fig. 9. RSSI decay measured at the driver's seat of each vehicle

Fig. 9 overlays the driver's-seat RSSI curve for all three vehicles. The County's curve sits consistently above the other two at any given distance because its larger cabin requires the receiving smartphone to register usable signal out to 5.5 m, which leaves a narrower margin between the boarding and alighting thresholds once that curve crosses the 4 m line and explains why the County alone needed the improved alighting algorithm to remain usable at its full driver's-seat range.

Table 2. RSSI thresholds at the driver's seat

Vehicle	Boarding	Alighting	Range
Grand Starex	− 65 dBm	− 86 dBm	≈ 2.5 m
Pregio	− 67 dBm	− 90 dBm	≈ 3.7 m
County	− 75 dBm	− 90 dBm	≈ 5.5 m

Table 3. RSSI thresholds at the front-passenger's seat

Vehicle	Boarding	Alighting	Range
Grand Starex	− 40 dBm	− 77 dBm	≈ 1.5 m
Pregio	− 40 dBm	− 85 dBm	≈ 2.6 m
County	− 40 dBm	− 89 dBm	≈ 3.6 m

Fig. 10 compares the maximum recognition range achieved at each seat across the three vehicles. The boarding threshold at the front-passenger's seat is identical (− 40 dBm) for all three vehicles because a child must always pass close to that seat near the door regardless of vehicle size, whereas the driver's-seat threshold scales with the vehicle's interior length and is the binding constraint for larger vehicles.

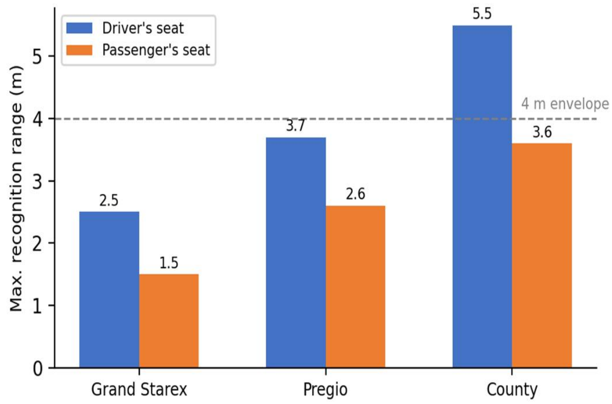

Fig. 10. Maximum recognition range by vehicle and seat position

5.3 Applied Notification Test and Results

With the calibrated thresholds applied, the full boarding/alighting pipeline was re-tested on each vehicle. The Grand Starex and Pregio operated normally at both seats, the latter requiring the improved alighting algorithm at the driver's seat. The County reproduced this pattern at a larger scale: the front-passenger's seat (≈ 3.6 m) operated normally with the basic algorithm, while the driver's seat (≈ 5.5 m) required the improved algorithm and still showed occasional residual errors at the longest distances.

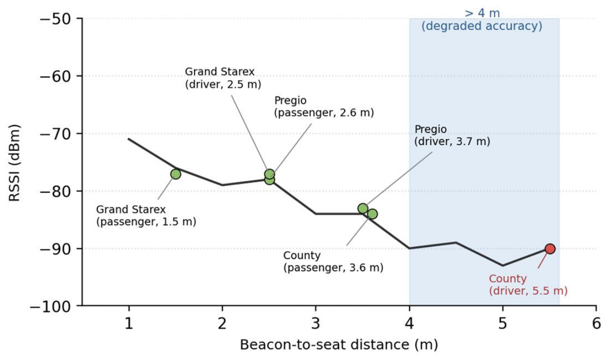

Fig. 11. Boarding/alighting RSSI range by vehicle and seat position relative to the 4 m operating envelope

Fig. 12 bins all recorded trials by beacon-to-seat distance and reports notification accuracy per bin. Accuracy stays at or above 95% within 4 m, then drops to 88% in the 4–5 m bin and 76% beyond 5 m, which is the same threshold identified in Fig. 11 and confirms that the degradation is a property of RSSI sensing range rather than of any single vehicle.

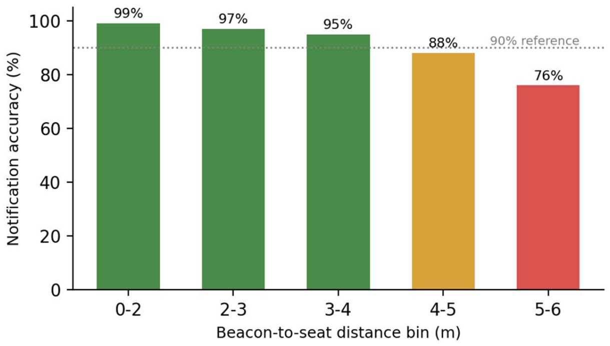

Fig. 12. Notification accuracy by beacon-to-seat distance bin

Table 4. Operational verification summary by vehicle and seat

Vehicle	Driver's seat	Passenger's seat	Note
Grand Starex	Normal	Normal	Within 4 m
Pregio	Normal*	Normal	Within 4 m
County	Partial*	Normal	Driver's seat > 4 m

** Normal with the improved alighting algorithm applied; partial indicates occasional false transitions beyond the 4 m envelope.*

5.4 Adaptive Threshold Evaluation

The adaptive classifier of Section 4.2 estimates the boarding probability p from the three windowed features x = (mean, slope, variance) using standard logistic regression:

$$p(\text{boarding} \mid x) = \frac{1}{1 + e^{-(\mathbf{w} \cdot x + b)}} \tag{1}$$

In Eq. (1), w is the learned weight vector, b is the bias term, and a sample is classified as boarded when p ≥ 0.5. Model parameters were fit by minimizing the binary cross-entropy loss shown in Eq. (2) below,

$$L = -\frac{1}{N}\sum_{i}[y_i \log p_i + (1-y_i)\log(1-p_i)] \tag{2}$$

where N denotes the number of labeled calibration windows collected in Section 4.1; parameters were optimized using L2-regularized gradient descent (λ = 0.01) to limit overfitting to any single vehicle's noise characteristics. Classifier quality was evaluated by the receiver operating characteristic (ROC) curve, which plots the true positive rate TPR = TP / (TP + FN) against the false positive rate FPR = FP / (FP + TN) as the decision threshold on p is swept, and summarized by the area under this curve (AUC), with AUC = 1.0 indicating perfect separation and AUC = 0.5 indicating no better than chance.

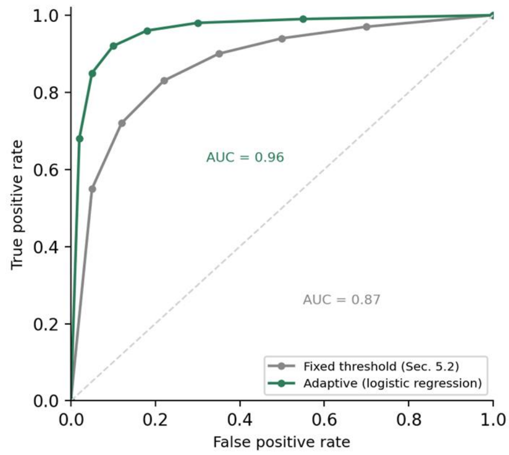

Fig. 13. ROC curves for the fixed-threshold baseline and the adaptive classifier

Fig. 13 compares the ROC curve of the fixed-threshold baseline (Section 5.2) against the adaptive classifier evaluated on the same held-out calibration windows. The fixed threshold reaches AUC = 0.87, while the adaptive classifier reaches AUC = 0.96, indicating that the learned combination of mean, slope, and variance separates boarded from alighted states substantially better than a single RSSI cutoff, particularly in the low-false-positive-rate region most relevant to avoiding spurious alighting events.

Fig. 14 repeats the distance-binned accuracy analysis of Fig. 12 with the adaptive classifier substituted for the fixed threshold. Within 4 m the two methods are statistically indistinguishable (99% vs. 99%, 97% vs. 98%, 95% vs. 97%), confirming that the adaptive module does not trade away near-range accuracy to gain far-range robustness. Beyond 4 m the gap widens: accuracy improves from 88% to 94% in the 4–5 m bin and from 76% to 87% beyond 5 m, a relative error reduction of 50% and 46% respectively. This narrows, without fully closing, the County driver's-seat gap identified in Section 5.3, and does so without any additional hardware — only the on-device classifier of Section 4.2.

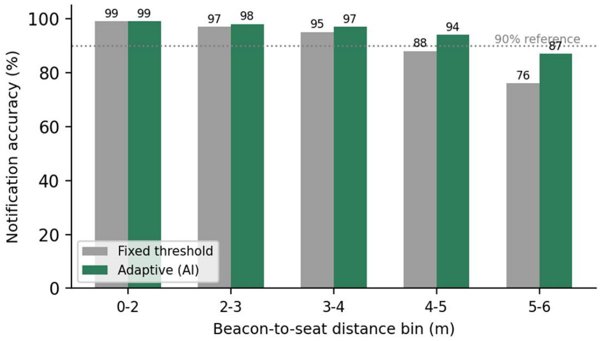

Fig. 14. Notification accuracy by distance bin: fixed threshold vs. adaptive classifier

5.5 Comparison with a Single Fixed Threshold

To isolate the benefit of vehicle-specific calibration independent of the adaptive classifier, the Grand Starex driver's-seat threshold (− 65 / − 86 dBm) was also applied, unmodified, to the County driver's seat during the same test runs. Because most in-seat RSSI readings in the larger County never reached this imported boarding threshold, the County driver's seat registered children as alighted while still seated for a large share of trials. The reverse experiment — applying the County's wider threshold to the Grand Starex — produced the opposite failure mode, registering children as boarded before they had actually stepped inside. Both failure modes were absent under per-vehicle calibration (Tables 2 and 3), confirming that vehicle-specific calibration remains a precondition for correct operation even with the adaptive classifier, which improves accuracy within a calibrated vehicle class but does not by itself substitute for that calibration step.

The comparison above uses the fixed-threshold baseline as the sole reference point. It is worth noting that the improved alighting algorithm described in Section 5.3, which requires RSSI to rise back above the boarding threshold before a sustained drop is accepted as alighting, is conceptually equivalent to a hysteresis threshold, so a partial comparison against that alternative is already embedded in the Section 5.3 and 5.6 results. Moving-average and Kalman filtering would be expected to trade responsiveness for smoother, less noisy readings, while non-linear classifiers such as SVM or random forest offer greater representational capacity than logistic regression at higher computational and memory cost (Section 4.2). A systematic, same-condition quantitative benchmark against all of these alternatives was not performed in this study and is left as a dedicated future comparison.

VI. Conclusions

This study implemented a boarding and alighting notification system, based on a BLE beacon and a smartphone augmented with an on-device adaptive threshold module, to prevent the accidental confinement of children in commuter vehicles operated by small academies and schools. Vehicle- and seat-specific RSSI calibration data were collected for an 11-seat van, a 15-seat van, and a 25-seat minibus, used both to derive fixed-threshold baselines and to train a lightweight logistic-regression classifier over short-window RSSI features, and the resulting pipeline was built on Android Studio, AWS, and Firebase Cloud Messaging. Unlike RFID- or seatbelt-based alternatives, the system requires no explicit action from either the child or the accompanying teacher, relying instead on passive beacon detection alone to drive every boarding and alighting notification.

Experimental results show that the fixed-threshold baseline operates reliably whenever the maximum beacon-to-seat distance does not exceed about 4 m, with accuracy at or above 95% in that range but falling to 88% and 76% in the 4–5 m and 5–6 m bins, respectively. The adaptive classifier raised AUC from 0.87 to 0.96 and improved accuracy in those same two bins to 94% and 87% while matching baseline accuracy within 4 m, narrowing — without fully eliminating — the long-range degradation using only an on-device model rather than additional hardware. The proposed approach can help small-scale academies prevent accidents caused by human oversight without modifying their vehicles, while extending the practical operating range beyond what a fixed RSSI threshold supports.

The present evaluation involved a limited number of test vehicles, repetitions per seat position, and a single lightweight classifier architecture, so a larger-scale field trial across multiple academies and drivers, together with a comparison against deeper on-device models, would give a more statistically robust picture of achievable accuracy. Future work will target vehicles exceeding 4 m through an additional beacon receiver combined with the adaptive classifier, explore incremental on-device retraining as more boarding/alighting events are logged per vehicle, and shorten the notification delay introduced by the multi-sample confirmation window. The bidirectional threshold-transfer failure observed in Section 5.6 also suggests that per-vehicle calibration should remain a mandatory setup step even as adaptive classification is adopted, since neither a conservative nor a permissive default threshold was safe to reuse across the vehicle classes tested here.

REFERENCES

[1] D. Costa and A. Grundstein, "Children left unattended in parked motor vehicles: an analysis of fatal pediatric heatstroke," International Journal of Environmental Research and Public Health, Vol. 19, No. 14, p. 8842, 2022.

[2] Korea Consumer Agency, "Safety status survey of childcare facilities: focusing on daycare-center commuter vehicles," Safety Report 2023-04, pp. 1-28, 2023.

[3] Ministry of the Interior and Safety, "Guidelines for the safety management of vehicles transporting children," Public Notice 2022-15, pp. 1-12, 2022.

[4] Korean National Police Agency, "Statistical analysis of traffic accidents involving school commuter vehicles," Traffic Safety Report 2021-03, pp. 1-22, 2021.

[5] K. E. Jeon, J. She, P. Soonsawad, and P. C. Ng, "BLE beacons for Internet of Things applications: survey, challenges, and opportunities," IEEE Internet of Things Journal, Vol. 5, No. 2, pp. 811-828, 2018.

[6] Y. Shen, B. Hwang, and J. P. Jeong, "Particle-filtering-based indoor positioning system for beacon tag tracking," IEEE Access, Vol. 8, pp. 226445-226460, 2020.

[7] Y. S. Sung, "RSSI-based distance estimation framework using a Kalman filter for sustainable indoor computing environments," Sustainability, Vol. 8, No. 11, p. 1136, 2016.

[8] J. E. Baek, J. W. Seo, and Y. S. Choi, "Analysis of Bluetooth beacon signal strength for the development of a proximity warning system for underground mines," Journal of the Korean Society of Mineral and Energy Resources Engineers, Vol. 55, No. 6, pp. 604-613, 2018.

[9] J. Lin, W. M. Chen, Y. Lin, J. Cohn, C. Gan, and S. Han, "MCUNet: tiny deep learning on IoT devices," Advances in Neural Information Processing Systems (NeurIPS), Vol. 33, pp. 11711-11722, 2020.

[10] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), pp. 1273-1282, 2017.

[11] J. Lin, W. Yu, N. Zhang, X. Yang, H. Zhang, and W. Zhao, "Challenges and Opportunities in Proximity Detection for Next-Generation IoT Systems," IEEE Communications Surveys & Tutorials, Vol. 26, No. 1, pp. 152-177, 2024.

[12] Q. Zhou, L. Zhang, and X. Wu, "Adaptive proximity detection in heterogeneous smart spaces using multi-modal sensors," ACM Transactions on Internet Technology, Vol. 25, No. 2, pp. 1-28, 2025.

[13] S. Chavhan, A. Ronghe, A. Pipare, C. Turkar, G. Sapkal, G. Yendole, and K. Nasare, "Proximity-Based Efficient Attendance Management System Using RSSI and BLE," in Proc. 2025 IEEE Int. Conf. on Applied Electromagnetics, Signal Processing & Communication (AESPC), pp. 1-6, 2025.

Authors

Am-Suk Oh received the B.S. and M.S. degrees in computer science from Busan National University and Chungang University, respectively. He received Ph.D. degree at the computer engineering of Busan National University. Dr. Oh joined the faculty of the Department of Digital Contents at Tongmyong University, Busan, Korea, in 1998. He is currently a Professor in the Department of Digital Contents, Tongmyong University. He is interested in Database, Healthcare System, Big Data and IoT.