Perspective: From Orientation Effects to Trustworthy and Adaptive Bluetooth Channel Sounding

Dzenita Dzafic
Graz University of Technology
Graz, Austria
dzenita.dzafic@tugraz.at

Abstract

Bluetooth Channel Sounding (BLE-CS) promises fine-grained ranging on low-power devices, but our first experiments with commercial Bluetooth 6.0 platforms made clear that getting a distance number out of the radio is the easy part. Relative device orientation caused large, configuration-specific biases; antenna diversity made a real difference in robustness; and the quality indicators reported by the controller missed a good share of the bad estimates. Taken together, these results shift our research focus from asking "how accurate can BLE-CS be?" toward "when should an embedded device trust its own range estimate", and "what should it do when that trust breaks down?" We report these findings as motivation for a set of next steps we see as necessary to make BLE-CS robust and practically usable in real-world deployments: calibration that accounts for orientation and hardware variability, lightweight machine learning to catch and correct degraded estimates, and online adaptation of ranging parameters under energy, accuracy, and latency constraints. Getting BLE-CS to a trustworthy state will likely mean pairing signal-processing knowledge with small adaptive components, rather than chasing a single estimator tuned for a fixed environment.

CCS Concepts

• **Computer systems organization** → **Embedded systems**; **Sensor networks**; • **Hardware** → *Sensor devices and platforms*; *Wireless devices*; • **Information systems** → **Location based services**; **Sensor networks**; • **Computing methodologies** → *Machine learning approaches.*

Keywords

Bluetooth Channel Sounding, ranging, localization, calibration, orientation effects, reliability estimation, machine learning, embedded systems

1 What Did Our Measurements Change?

BLE-CS combines phase-based ranging (PBR) and round-trip-time (RTT) measurements to get fine-grained distance estimates. Recent evaluations on commercial hardware report promising accuracy, but also reveal sensitivity to algorithm choice, configuration, outliers, and the environment [1].

In our first BLE-CS study we collected more than 190 000 samples on Nordic nRF54L15 and Silicon Labs EFR32MG24 boards [2, 3], across several line-of-sight (LoS) environments, board orientations, rotations, and antenna configurations [4]. The clearest effect turned out to be surprisingly local: just how the two devices were positioned relative to one another. For single-antenna ranging, orientation shifted the calibration offset by up to 70 cm on Nordic and 138 cm on Silicon Labs, more than the shifts we measured from changing environments (Figure 1a). That alone turns calibration into its own research problem: a fixed offset removes the bias in the setup where it was measured, but breaks once a device is rotated or deployed elsewhere.

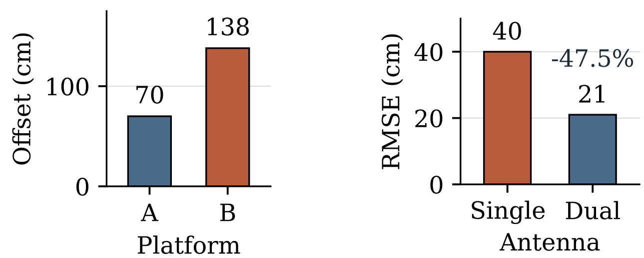

(a) Offset shift by orientation.

(b) Antenna-diversity effect on RMSE.

Figure 1: Orientation and antenna-diversity effects in the BLE-CS measurements.

Antenna diversity yielded a clear robustness gain: the tested dual-antenna configuration reduced RMSE by 47.5% (Figure 1b) and lowered outlier rates from 10.81–17.98% to 1.08–3.71%. We did not measure energy consumption in this campaign, so the cost of the additional antenna paths remains open. The radio's tone-quality indicator, meant to flag unreliable measurements, caught only 5.18–33.38% of the outliers, depending on orientation and environment. Put together, these results point to a broader goal: a deployable ranging system should provide not only a distance estimate, but also a calibrated measure of its reliability. Existing quality indicators are a good starting point, but their values must be calibrated against actual ranging errors.

2 Where Machine Learning Fits

Machine learning (ML) for radio ranging is not new, but work aimed specifically at BLE-CS is still determining what should be learned and what should remain model-based. Tsemko *et al.* apply a parametric neural network to Channel Sounding data and report better ranging accuracy than traditional super-resolution processing [5], while Chen *et al.* fuse several Bluetooth 6.0 ranging outputs with an attention-based network [6]. Earlier work on narrowband phase-based ranging trained a network to reconstruct missing or interfered tones before super-resolution ranging; for a small number of affected tones it approached reference accuracy but needed up to 60.5 kB of extra memory [7]. Our boards each provide only 256 kB of RAM [2, 3], so this already uses about 24% of the available memory, before accounting for the radio stack,

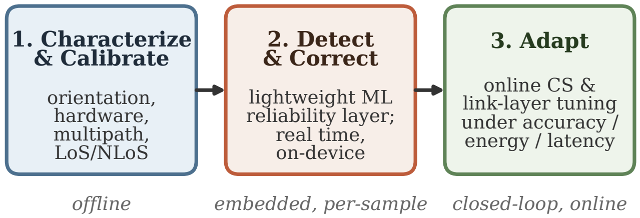

Figure 2: Progression from understanding BLE-CS failure modes to detecting unreliable measurements and adapting the ranging system online.

buffers, and application code. Whether such a method is practical therefore depends on how much accuracy it actually gains, and on its execution time and energy consumption on the target device. These results show that learned processing can help, but also raise a more practical question for embedded systems: what is the smallest learning problem that provides the largest robustness gain? Rather than replacing the entire distance estimator with a dense network, a promising route is a lightweight reliability layer on top of model-based processing. Its inputs could include agreement among IFFT- and phase-slope-based distance estimates – two established techniques that recover distance from the measured phase response – and RTT-based estimates, consistency across antenna paths, received-power or tone-quality readings, and features from the reconstructed delay response. The model could then flag likely LoS/NLoS or degraded measurements, estimate the probable ranging error, and correct the affected estimate instead of simply discarding it.

Work outside BLE-CS points in the same direction. Supervised learning has been used to detect LoS/NLoS conditions from Wi-Fi and BLE received-signal strength measurements [8]. In UWB ranging, transfer learning has been applied to NLoS classification and error correction across environments [9]. Related work jointly estimates distance and environmental labels with a variational model, illustrating how closely the two tasks can be coupled [10]. For BLE-CS, the central issue is therefore not simply whether a model transfers to unseen data, but which inputs remain calibrated when orientation, environment, interference, or hardware changes. The next step is to evaluate the approach across different deployment conditions, and assess both accuracy and whether unreliable estimates can be detected at acceptable memory, latency, and energy cost.

3 From Detection to Adaptation

The next step is moving from detecting unreliable measurements to actually reacting to them. We see this as a set of next steps necessary to make BLE-CS robust and practically usable in real-world deployments, following three connected stages, shown in Figure 2: characterizing and calibrating for how orientation, antenna configuration, multipath, LoS/NLoS propagation, interference, and hardware affect CS measurements; detecting degraded measurements and correcting or quality-weighting them in real time on energy-constrained devices; and adapting CS procedure and link-layer parameters online, targeting an accuracy–energy–latency trade-off rather than accuracy alone.

The third stage comes later in the trajectory on purpose. Online reinforcement learning is tempting here, since the best ranging configuration genuinely depends on environment and application, but it brings sample efficiency, stability, and exploration costs that are hard to justify on an embedded radio. Before a control policy can be learned at all, the system needs a state representation it can trust and a cost model it can actually measure for CS procedures.

4 A Different Evaluation Target

This perspective also changes what a BLE-CS evaluation should actually report. Median error and RMSE still matter, but neither tells you whether a method knows when it is failing or whether it holds up under distribution shift.

Future methods should be judged on three things instead: *accuracy, confidence or failure detection, and resource cost*. Whether a method holds up across orientations, environments, and eventually different devices should count as a real generalization test.

The broader lesson from this first study is that trustworthy ranging probably will not come from a single universally best estimator. It will come from a system that combines physical structure, diversity, confidence estimation, and adaptation, while staying small enough to actually run where BLE is useful in the first place: on resource-constrained devices.

5 Questions for the Community

Three questions feel worth raising at the EWSN community event. First, should calibration for orientation and hardware differences be standardized, or handled separately by each vendor? Second, how should the community fairly compare reliability and failure-detection methods across different platforms? Should applications state their accuracy, energy, and latency needs and have the radio translate that into settings automatically, instead of developers tuning them by hand for each deployment?

References

[1] Jorg Wieme et al. Performance Evaluation of Bluetooth Channel Sounding on Commercial Hardware. *IEEE Access*, 13:180862–180876, 2025.
[2] Nordic Semiconductor. *nRF54L15, nRF54L10, and nRF54L05 Wireless SoCs: Datasheet*, 2025.
[3] Silicon Labs. *EFR32MG24 Wireless SoC Family Data Sheet*, 2025.
[4] Dzenita Dzafic et al. Orientation Matters: An Empirical Study on Ranging Performance of Bluetooth 6.0 Channel Sounding. In *Proceedings of the 23rd International Conference on Embedded Wireless Systems and Networks (EWSN)*, 2026.
[5] Andrii Tsemko et al. Data-Driven Processing Using Parametric Neural Network for Improved Bluetooth Channel Sounding Distance Estimation. In *2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2025.
[6] Yutao Chen et al. High-Precision Ranging Fusion Using Neural Network for Bluetooth Channel Sounding. In *Wireless Artificial Intelligent Computing Systems and Applications*, volume 15686 of *Lecture Notes in Computer Science*, pages 131–142. Springer Nature Singapore, 2025.
[7] Alireza Sheikh et al. Phase-Based Ranging in Narrowband Systems With Missing/Interfered Tones. *IEEE Internet of Things Journal*, 10(17):15171–15185, 2023.
[8] Fernando J. Aranda et al. A Machine Learning-Based Approach to NLOS Detection for Wi-Fi and BLE RSS Measurements. In *2023 IEEE International Conference on Omni-Layer Intelligent Systems (COINS)*, pages 22–27. IEEE, 2023.
[9] Jaron Fontaine et al. Transfer Learning for UWB Error Correction and (N)LOS Classification in Multiple Environments. *IEEE Internet of Things Journal*, 11(3):4085–4101, 2024.
[10] Yuxiao Li et al. A Variational Learning Approach for Concurrent Distance Estimation and Environmental Identification. *IEEE Transactions on Wireless Communications*, 22(9):6252–6266, 2023.