Engineering and Technology Journal e-ISSN: 2456-3358

Volume 10 Issue 08 August-2025, Page No.-6605-6614

DOI: 10.47191/etj/v10i050, I.F. – 8.482

Infrastructure-Free Indoor Navigation Via Quorum-Sensing and LT-Style Mapping

Fredy Martínez

Facultad Tecnológica, Universidad Distrital Francisco José de Caldas, Bogotá, Colombia

ABSTRACT: This paper presents QSLT-Map, a bearing-only hierarchical mapping and navigation framework that fuses a bacterial quorum-sensing (QS) layer with an LT-style landmark tree on minimalist robots. The method stores compact viewframes and promotes slowly varying bearings to a stable backbone while leaving volatile cues near the leaves, enabling aggressive pruning without breaking homing. QS provides a local scalar intensity derived from single-hop broadcasts (ESP-NOW) that adapts the effective capture threshold, coordinates leader-follower roles, and guides pruning toward regions that matter. A saturated differential-drive controller based on a secant direction from bearing differences, with a wall-following fallback, closes the loop without reliance on metric maps or infrastructure. We implement the full stack on SERB+ESP32 with fixed-size buffers and FreeRTOS tasking, keeping per-agent computation $O(1)$ and memory within a few kilobytes for typical routes. Experiments in Player/Stage and on hardware show reduced time-to-goal and faster re-engagement after visual loss relative to an LT-only baseline, as well as high success under strict pruning and fewer collisions near doorways. The ranking of methods is consistent across testbeds, indicating that QS-aware capture and pruning transfer beyond instrumented arenas while remaining within the compute and communication budgets of small indoor platforms.

KEYWORDS: Bearing-only navigation, ESP32, Quorum sensing, Swarm robotics, Topological mapping.

I. INTRODUCTION

Indoor mobile robots built around low-cost components face stringent limits in computation, memory, and communication, which complicate the use of large metric maps and continuous localization (Niloy et al., 2021, Deshmukh et al., 2025, Huang et al., 2023). Many platforms cannot maintain precise odometry over extended runs, and camera pipelines heavy enough to compensate for drift exceed the available processing budget (Wu et al., 2023, Damjanović et al., 2025). Short-range radios offer only best-effort delivery and fluctuate under multipath, which undermines centralized coordination schemes and infrastructure-heavy designs (Jacinto et al., 2023). At the same time, practical deployments must handle transient occlusions, varying illumination, and narrow passages without human intervention (Martinez et al., 2020). These conditions motivate navigation strategies that store compact information, make decisions from local cues, and avoid reliance on a global, continuously updated world model.

Several strands of prior work attempt to reduce memory by using appearance or topology rather than full geometry, yet many of these methods still accumulate dense data or assume stable connectivity to a backbone service (Martinez et al., 2012, Song et al., 2023). Approaches inspired by social insects often encode global trails or require high-fidelity sensing that does not transfer cleanly to minimalist hardware (Wang et al., 2023). In contrast, bacterial quorum-sensing (QS) suggests a different coordination mechanism: local production and detection of a scalar signal, with thresholded switches that govern collective behavior (Sarmiento, 2021). This simple template can be implemented with broadcast packets and short windows, without synchrony or global identifiers (Ray, 2025). A mapping and control stack that exploits bearing-only structure while letting a QS field modulate capture density, group guidance, and pruning would address the gap between theory and the realities of small indoor robots.

We propose QSLT-Map, a bearing-only, LT-style hierarchical map fused with a QS layer that shapes what to store, when to regroup, and how to trim memory under pressure (Valentini et al., 2017). The map organizes viewframes into a tree where slowly varying bearings form a stable backbone and volatile cues remain near the leaves, preserving route information without bloating storage. The QS layer is realized as single-hop ESP-NOW broadcasts that are integrated into a local intensity; hysteresis on that scalar triggers role changes and stabilizes group behavior. A saturated differential-drive controller with a wall-following fallback closes the loop, so navigation proceeds even when temporary loss of matches interrupts the bearing tracker. All components run on a SERB+ESP32 platform, using integer arithmetic and fixed-size buffers to respect the budgets typical of classroom and laboratory robots.

The contributions are fourfold and are designed to be independently testable. First, we introduce a QS-aware mapping policy in which the effective capture threshold adapts with local signal intensity, promoting detailed storage in active regions and suppressing redundant sampling in quiet corridors. Second, we define a promotion and pruning scheme that preserves a long-range backbone of low-variance landmarks, allowing aggressive compression without breaking the chain needed for homing. Third, we implement a minimalist control loop based on a secant direction derived from bearing differences, coupled with a boundary mode engaged by contact sensing to bridge occlusions and concavities. Fourth, we provide a full SERB+ESP32 implementation with FreeRTOS tasking, packet formats, calibration artifacts, and a memory layout that fits within a few kilobytes for typical routes. Together, these elements yield a coherent stack that is simple to reproduce and evaluate against baselines.

The paper advances three hypotheses that our experiments investigate in simulation and on hardware. We expect the QS layer to reduce re-engagement latency after visual loss relative to an LT-only baseline, and we anticipate that QS-adaptive pruning will preserve success at low memory with a gradual impact on time-to-goal. We further hypothesize that leader guidance will lower collision rates and stabilize formations without exceeding the communication budget of single-hop broadcasts. A brief preview of the findings supports these expectations: time-to-goal and recovery improve, success remains high under strict pruning, and group behavior becomes more orderly at doorways. The remainder of the paper is organized as follows: Section 2 formalizes the system and sensing models, Section 3 details the QSLT-Map method, Section 4 describes the SERB+ESP32 implementation, Section 5 presents experiments and results, Section 6 discusses limitations and implications, and Section 7 concludes the work.

II. SYSTEM MODEL AND PRELIMINARIES

This work addresses indoor navigation with minimalist robots that rely on local measurements and single-hop communication, without a global metric map or infrastructure. Environments are represented as a connectivity graph of regions joined by doorways, with static obstacles and occasional occlusions. Each agent decides from its own sensing and recent broadcasts; no synchronized clock or shared state is assumed. The objective is reliable route following and homing under tight memory, computation, and radio budgets typical of SERB+ESP32 platforms. Throughout, we favor compact bearing representations, simple switching logic, and saturated control laws that are robust to modest sensor noise.

Each robot is modeled as a differential-drive unicycle with state $x = [p_x, p_y, \psi]^{\mathsf{T}}$ on the plane. Let r be the wheel radius and b the wheelbase, and denote by ω_r, ω_l the right/left wheel angular rates. The body linear and angular speeds are

$$v = r/2\ (\omega_r + \omega_l), \qquad \omega = r/b\ (\omega_r - \omega_l) \qquad (1)$$

and the kinematics satisfy

$$\dot{p}_x = v cos\psi, \quad \dot{p}_y = v sin\psi, \quad \dot{\psi} = \omega \qquad (2)$$

Wheel commands are rate-limited and mapped through (1) with symmetric saturation $|v| \le v_{max}, |\omega| \le \omega_{max}$ to protect actuators and ensure predictable behavior on floors.

Sensing follows a minimal stack that outputs heading and landmark bearings with bounded latency. Heading ψ comes from a magnetometer with standard hard-/soft-iron calibration; when disturbed, a visual compass can substitute over short windows. Bearings $\theta_i \in (-\pi, \pi]$ to detected landmarks are obtained either from a fisheye camera (via a precomputed LUT that maps pixels to angles) or from a ring of directional IR receivers (sectorized angles with hysteresis). Low-force bumpers provide contact cues that trigger boundary following when frontal motion is unsafe or matching collapses. All sensing paths expose the same interface $\{\theta_i\}$ to the mapping layer after alignment to the current heading.

A *landmark* is an identifiable cue with a persistent identifier; a *viewframe* at step k is the set $V_k = \{(\mathrm{id}_i, \theta_i)\}_{i=1}^{N_k}$ aligned to the current heading estimate. Viewframes are organized in an LT-style hierarchy where slowly varying bearings are promoted toward inner nodes, while volatile cues remain near the leaves. This yields a compact "spine" that preserves long-range structure and supports homing when memory pressure forces pruning by level L^*. The formal similarity used to decide capture/fusion and the variance rule behind promotion are defined once in Section 3, see (4); we avoid duplicating equations here to keep the presentation concise. In practice, this organization allows aggressive compression without losing the order information required to chain the route.

Local coordination is provided by a quorum-sensing (QS) layer implemented with single-hop ESP-NOW broadcasts at low duty. Each robot integrates a short history of received packets into a scalar intensity that reflects nearby activity, then applies hysteresis to switch between *explorer*, *leader*, and *follower* roles. The same scalar modulates capture density and informs pruning so that detail is kept where teams operate and discarded elsewhere. Leader election uses simple tie-breaking and time-to-live rules, avoiding shared identifiers or global arbitration. The precise packet fields, intensity computation, and effective capture threshold are specified in Section 3, see (7) and (5).

Control closes the loop with a bearing-based tracker and a boundary fallback that handles occlusions and concavities. The commanded heading comes from a secant direction computed over matched bearings to the target viewframe, then mapped to (v, ω) with the saturations already stated; the exact expression appears in (6). When the tracking direction becomes unreliable or a bumper fires, the robot follows the nearest boundary until enough matches return. This simple hybrid policy avoids reliance on dense perception and keeps the cycle time within the limits of the embedded platform. All implementation details, task split, timing, memory layout, and diagnostics-are consolidated in Section 4.

III.METHOD: QSLT-MAP

QSLT-Map pairs a bearing-only landmark tree with a quorum-sensing (QS) layer that regulates capture density, group guidance, and pruning using single-hop broadcasts. The pipeline takes heading-aligned bearings and recent QS packets, decides whether to capture or fuse a viewframe, updates the hierarchical tree with variance-based promotion, and issues wheel commands from a simple tracker with a boundary fallback. All decisions are local and require neither a global map nor synchronized timing. Figure 1 summarizes the flow and highlights the points where the QS signal modulates behavior. The design keeps computation and memory bounded for embedded controllers.

Matching, gating, and QS-aware capture. Let $wrap(\alpha) = atan2(sin\alpha, cos\alpha)$ and let $\mathcal{M}$ denote matched landmark IDs between V and $\tilde{V}$. We use a robust Pseudo-Huber penalty:

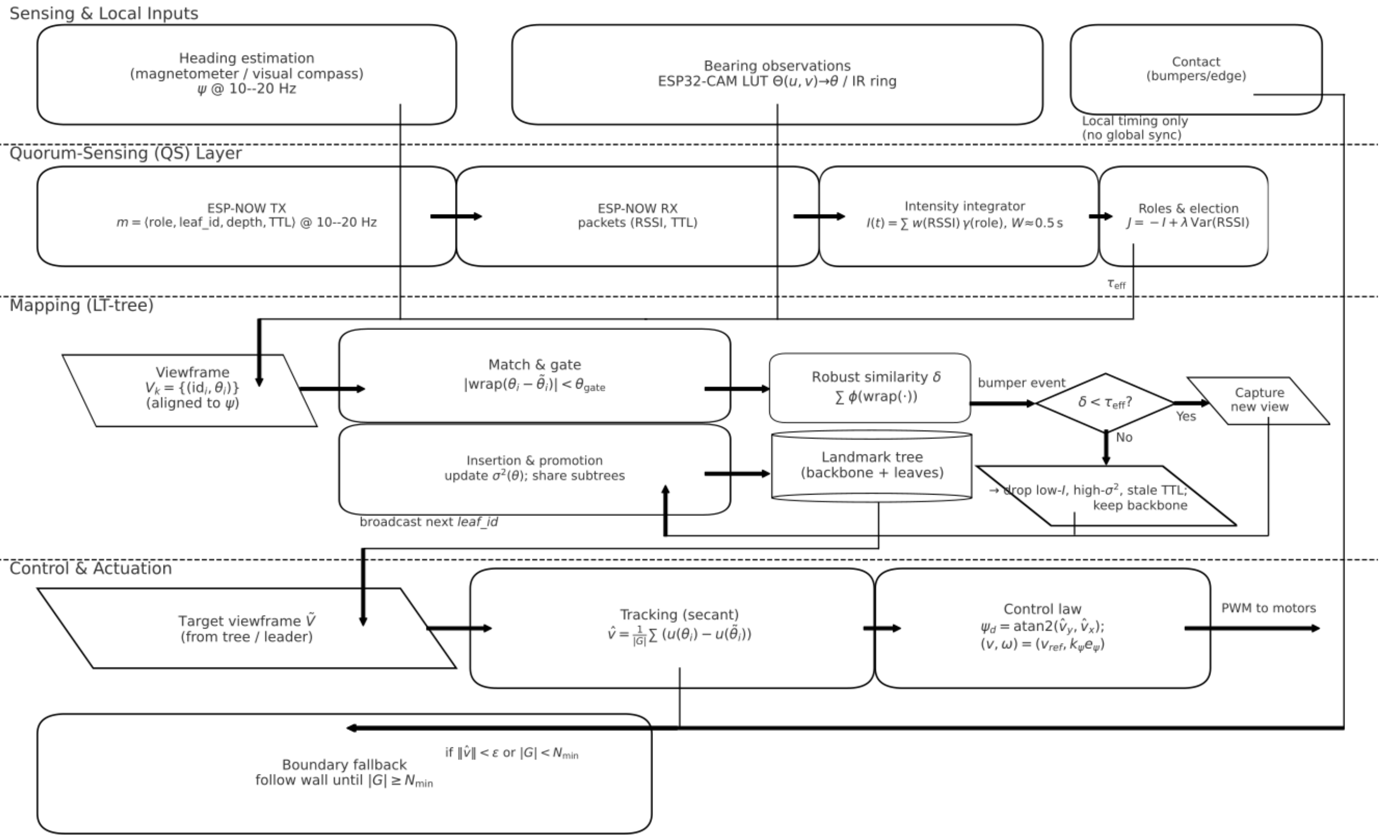

Figure 1: Method pipeline. Bearings and QS packets feed matching and QS-aware capture; insertion promotes low-variance bearings; leader broadcasts the next target; tracking uses a secant direction with a boundary fallback

$$\phi(x) = \delta_0^2\left(\sqrt{1 + (x/\delta_0)^2} - 1\right) \qquad (3)$$

and define the viewframe discrepancy

$$\delta(V, \tilde{V}) = \sum_{i \in \mathcal{M}} \phi\left(wrap\left(\theta_i - \tilde{\theta}_i\right)\right) \qquad (4)$$

Outliers are removed by an angle gate $\left|wrap\left(\theta_i - \tilde{\theta}_i\right)\right| < \theta_{gate}$ before evaluating (4). To adapt capture density to local activity we set

$$\tau_{eff} = clip(\tau_0 - \kappa\, I(t),\ \tau_{min},\ \tau_{max}) \qquad (5)$$

where $I(t)$ is the QS intensity (below), $\kappa > 0$ tunes responsiveness, and $clip(a, l, u) = min\{max\{a, l\}, u\}$. A new viewframe is captured if $\delta(V_{obs}, V_{last}) < \tau_{eff}$; otherwise the observation is fused into the latest entry.

Insertion, promotion, and pruning-friendly structure. Viewframes are inserted in order into a rooted landmark tree that preserves traversal along leaves while sharing inner nodes. For each landmark i we maintain an empirical bearing variance $\sigma^2(\theta_i)$ over a short window; small variance promotes i toward the root, large variance keeps i near leaves. This yields a compact backbone of slowly varying bearings that supports homing even when memory pressure forces removal of local detail. Subtree sharing reduces duplication when common landmark subsets reappear along different branches.

Local tracking and boundary fallback. Let $\tilde{V}$ be the target viewframe and V the current observation after alignment. With $u(\theta) = [cos\theta,\ sin\theta]^\top$ and the gated match set G, the secant direction is

$$\hat{v} = \frac{1}{|G|}\sum_{i\in G}\left(u(\theta_i) - u(\tilde{\theta}_i)\right) \qquad (6)$$

The desired heading $\psi_d = atan2(\hat{v}_y, \hat{v}_x)$ and the error $e_\psi = wrap(\psi_d - \psi)$ produce $(v,\omega) = (v_{ref},\ k_\psi e_\psi)$, mapped to wheel speeds with the saturations already stated in §2. If $\|\ \hat{v}\ \|$ falls below a threshold or a bumper triggers, a boundary-following mode advances along the obstacle until $|G|$ is restored, at which point tracking resumes. This hybrid policy keeps control simple while handling occlusions and concavities without dense perception.

QS packets, intensity, and roles. Each robot broadcasts at low rate a fixed payload $m = \langle role,\ leaf\_id,\ depth,\ TTL\rangle$ over ESP-NOW. Received packets within a sliding window W are aggregated into

$$I(t) = \textstyle\sum_{j\in\mathcal{N}(t)} w\left(RSSI_j\right)\gamma\left(role_j\right) \qquad (7)$$

with $w(\cdot)$ down-weighting weak or stale entries and role weights chosen so that $\gamma(leader) > \gamma(follower) \geq \gamma(explorer)$. Hysteresis thresholds $\theta_{on} > \theta_{off}$ govern the switches among explorer, leader, and follower to avoid flapping near boundaries. Leader election minimizes

$$J(j) = -I_j + \lambda\, Var(RSSI)_j \qquad (8)$$

balancing attraction (I_j) and stability (RSSI dispersion) with $\lambda > 0$; ties break by randomized back-off and TTL expiry.

QS-adaptive pruning and memory budget. Let n_{VF} be the number of viewframes and n_ℓ the average landmarks per view, with per-item bytes (b_{id}, b_θ) and index overhead B_{idx}. The storage model is

$$B \approx n_{VF}\, n_\ell\,(b_{id} + b_\theta) + B_{idx} \qquad (9)$$

Pruning is triggered when $B/B_{max} > \beta$ or when regional intensity decays, and proceeds toward level L^* by dropping leaves with low I, high $\sigma^2(\theta)$, or stale TTL, while preserving promoted inner nodes. Because the traversal order is encoded along leaves and anchored by the backbone, homing remains feasible under aggressive memory cuts. This policy concentrates detail where activity is high and yields graceful performance degradation as memory tightens.

Complexity and timing. Per-cycle operations per agent are constant update $I(t)$, evaluate (4) over the gated matches, and compute (6), so compute does not grow with team size. Communication scales with neighborhood size rather than the swarm, since broadcasts are single-hop and best-effort. Empirically, the loop runs at 10-20 Hz with slack on microcontroller targets, leaving margin for logging and safety checks.

IV. IMPLEMENTATION ON SERB+ESP32

Hardware platform and integration. The robot is built on the SERB differential-drive chassis with two DC gear motors and a passive caster, chosen for its accessible form factor and repeatable geometry. The controller is an ESP32 module (WROOM/WROVER family) that handles sensing, local communication, and real-time control without an external computer. Sensing includes a magnetometer for heading, a bearing channel provided either by an ESP32-CAM with a fisheye lens or by a ring of directional IR receivers, and contact switches for wall-following triggers. Motor power and logic power are separated to reduce coupling; motor drivers are TB6612FNG devices with 3.3 V logic compatibility and low on-resistance, which eases thermal requirements.

Power system, wiring, and EMC. A 2S Li-ion/LiPo battery (7.4–8.4 V) supplies a buck converter to 5 V for motors and camera, and a low-noise 3.3 V rail for the ESP32 and IMU; the two rails meet at a star ground near the battery return. Peak current is budgeted as

$$I_{tot}^{peak} \approx I_{mot}^{stall} + I_{ESP32}^{TX} + I_{cam} + I_{IMU} \qquad (10)$$

and regulators are sized with a margin of at least 30% over I_{tot}^{peak} to accommodate startup transients. An LC filter on the motor rail and twisted-pair wiring to the motors mitigate conducted and radiated noise; flyback protection is included where the driver requires it. The ESP32 brown-out detector remains enabled, a TVS diode protects the 5 V rail, and power-good is debounced before camera initialization to avoid corrupted frames. These measures reduce resets during high-load maneuvers and stabilize radio operation when multiple robots accelerate simultaneously.

Sensors and calibration workflow. The magnetometer undergoes a two-step calibration (hard-iron offset and soft-iron scale/rotation via ellipse fitting), and offsets are stored in non-volatile storage so that power cycles do not require re-tuning. Heading is fused with a low-pass filter tuned to reject fast electromagnetic disturbances from the driver, while still tracking slow yaw changes. The camera path uses a fisheye lookup table $\Theta(u,v) \rightarrow \theta$ generated offline on a calibration grid; the final LUT is stored in flash or PSRAM and accessed by integer index to keep latency low. The IR path maps the strongest sector to a bearing and applies a short moving average with hysteresis to avoid sector flapping near boundaries; both paths produce the same θ interface for the mapping layer. Contact switches trigger interrupts that preempt the tracker and enter wall-following, and a minimum dwell time prevents rapid toggling when brushing along rough walls.

ESP-NOW quorum layer and diagnostics. Local communication uses ESP-NOW in a fixed channel to avoid interference from nearby access points; encryption is optional and disabled during bring-up to simplify tracing. Each node broadcasts at 10-20 Hz a compact packet $m =$ ⟨role, leaf_id, depth, TTL⟩, and a receive ring buffer (32 packets) is processed at a higher priority than the control task. Intensity $I(t)$ is integrated over a window $W \approx 0.5$ s using both packet counts and RSSI with role weights, and persender TTLs remove stale entries. A tie-breaking back-off timer resolves simultaneous leader candidacies, while an LED code and a UART log expose role, mean/variance of RSSI, and basic counters. These diagnostics compress the time to resolve field issues such as marginal radio range or misaligned antennas, which can otherwise masquerade as mapping faults.

Firmware architecture and timing on FreeRTOS. Tasks are partitioned to respect deadlines and minimize blocking between modules: a high-priority radio task handles ESP-NOW RX/TX and intensity maintenance, a sensing task extracts bearings and heading, a control task performs matching, computes δ and the secant direction $\hat{v}$, and updates wheel PWM, and a low-priority mapping task captures viewframes, inserts them into the landmark tree, and executes pruning. Inter-task communication uses lock-free queues for bearings and QS messages, while a small mailbox communicates the current target leaf to the controller; a mutex protects tree updates without entering the control loop. The control loop runs at 15 Hz with a watchdog at twice the period, and only bumpers and camera DMA use interrupts to keep latency deterministic. Care is taken to avoid priority inversion by bounding critical sections and keeping the radio task free of heap allocations. This organization fits within the ESP32 budget while leaving headroom for logging and parameter edits through a simple CLI.

Data layout, memory footprint, and persistence. A viewframe item stores an unsigned 16-bit landmark identifier and a 16-bit fixed-point angle in milliradians, so each landmark costs 4 B including alignment. With n_{VF} viewframes and n_ℓ landmarks per view, storage follows

$$B \approx n_{VF}\, n_\ell\, (b_{id} + b_\theta) + B_{idx} \qquad (11)$$

which for $n_{VF} = 200$, $n_\ell = 12$, and $(b_{id}, b_\theta) = (2,2)$ yields ≈ 9.6 kB plus indices below 4 kB. Tree nodes keep vectors of IDs and angles, depth, a use counter, and a compact variance estimate for bearings used by promotion and pruning. Calibration constants (magnetometer offsets, LUT signature), thresholds $(\tau_0, \delta_0, \theta_{on}, \theta_{off})$, and the active pruning level L^* are persisted in NVS with a checksum, and parameters can be updated at runtime and saved on user command. Buffer sizes are fixed at compile time for predictability, and the LUT resides in flash or PSRAM depending on the camera configuration.

Control, safety interlocks, and field behavior. The controller maps the secant direction to (v, ω) with symmetric saturation and a bounded rate of change on ω to limit tire slip on smooth floors. Wall-following uses a left/right heuristic based on the first bumper hit and exits when $|G| \geq N_{\min}$ or $\| \hat{v} \|$ recovers above a set threshold, at which point the bearing tracker resumes. If both intensity and landmark matches vanish, the robot enters a slow crawl with periodic spins to reacquire either visual bearings or a QS gradient, thereby avoiding prolonged deadlock. Undervoltage or brown-out triggers an immediate stop and retries after a short delay, and role transmission is paused during recoveries to reduce message clutter. A lightweight CLI allows hot-swapping of gains and thresholds during trials, with edits applied atomically and persisted only after explicit user confirmation.

Mechanical placement and bring-up tests. The IMU is placed away from drivers and motor leads to reduce magnetic bias, and the camera mount is set at mid-height with a slight downward pitch to avoid ceiling glare while preserving wall visibility. The IR ring, when used, remains unobstructed by the chassis and cables, and the center of rotation is aligned with the chassis center to keep controller assumptions valid. Initial tests proceed in stages: power and EMI checks, heading stability with the motors idling and under load, LUT sanity on a printed calibration pattern, QS range and RSSI variance in the target environment, and closed-loop tracking in a straight corridor. Pruning under induced load is tested by raising capture density and then lowering the pruning level until the map compresses without breaking homing. Table 1 lists the parts used in the reference build together with the key electrical characteristics needed to replicate the platform.

Table 1: Bill of Materials and Key Specifications (Reference Build)

Component	Example Part	Key Specs / Notes
Chassis	SERB kit	Diff. drive, caster, acrylic plates
MCU	ESP32 WROOM/WROVER	Dual core, Wi-Fi/BT, ESP-NOW
Camera (opt.)	ESP32-CAM + fisheye	Wide FOV, LUT in flash/PSRAM
Motor driver	TB6612FNG	3.3 V logic, low $R_{DS(on)}$
Magnetometer	QMC5883L/LSM303	I2C, heading fusion
Bearing sensor	IR ring (16+ sectors)	Sectorized θ with hysteresis
Contact switches	Low-force bumpers	ISR, debounced
Power (buck)	5 V 3 A module	Motor/camera rail
Power (3.3 V)	LDO or buck	Logic rail, low noise
Battery	2S Li-ion/LiPo	7.4–8.4 V, 2200 mAh typical
EMC parts	LC filter, TVS, polyfuse	Rail protection and noise control

V. EXPERIMENTS & RESULTS

Objectives and hypotheses. The experimental study evaluates whether QSLT-Map achieves reliable route following and homing with bearing-only sensing under strict memory and computation budgets. A first objective is to quantify the effect of the quorum-sensing (QS) layer on recovery from temporary visual loss, where viewframe matching degrades and wall-following must bridge gaps. A second objective is to measure how QS-adaptive pruning preserves success when memory is cut aggressively, relative to non-adaptive pruning or unbounded maps. We also examine group behavior by assessing leader selection stability and collision rates during tight maneuvers, since these directly affect reproducibility on minimalist hardware. The central hypotheses are: (H1) QSLT-Map lowers time-to-goal and re-engagement latency versus an LT-only baseline, (H2) QS-adaptive pruning maintains success across large memory reductions with graceful degradation of time-to-goal, and (H3) QS leader guidance reduces collisions and leader churn compared with unguided teams at the same broadcast rate.

Testbeds and scenarios. Two testbeds are used: a Player/Stage simulator executing the identical control and mapping logic, and real SERB+ESP32 robots in a taped-out arena of similar footprint. The environment comprises two connected rooms with doorways and three internal "holes" (obstacle clusters) that force detours and stress wall-following. Landmarks are either active beacons or high-contrast fiducials placed at wall locations with clear sight lines; layouts include a loop, a branched route, and two preferred gathering areas. Team sizes are $N \in \{4,8,12\}$ to test scaling, and lighting on the physical arena is varied across low/medium/high to probe bearing robustness. Communication uses ESP-NOW on a fixed channel at 10-20 Hz; packets, RSSI, and TTL expiries are logged continuously together with control and bumper events.

Baselines and experimental protocol. Four baselines are compared to the full method: (B0) LT-only (no QS), (B1) QSLT-Map without leader election, (B2) QSLT-Map with fixed capture threshold (no τ_{eff} adaptation), and (B3) QSLT-Map without pruning (upper bound on memory). For memory trade-offs we set pruning levels $L \in \{7,5,3\}$ and memory quotas $\beta \in \{0.6,0.8\}$ of $B_{\max}$, while keeping all other parameters fixed after calibration. Each condition is run for at least ten independent trials with randomized start headings; maps and runtime parameters are reset between runs to prevent carryover effects. Battery state of charge is controlled within a narrow band to minimize variance in motor performance, and a short magnetometer warm-up spin is executed prior to each run. Runs terminate on success, on a fixed time budget, or on deadlock (no progress) after a registered timeout; all failure reasons are recorded.

Metrics and formal definitions. Let S be the success rate (fraction of runs reaching the terminal viewframe) and let T_g be the time-to-goal in seconds. Path efficiency is reported as the ratio:

$$R_p = \frac{D_{traveled}}{D_{geodesic}} \tag{12}$$

where $D_{geodesic}$ is the shortest graph distance along the region-door connectivity. Collision pressure is normalized as:

$$C_m = \frac{\#\ bumper\ events}{D_{traveled}}\ [m^{-1}] \tag{13}$$

and recovery performance is characterized by the re-engagement latency

$$T_r = min\{t \geq 0:\ |G(t)| \geq N_{min}\ \wedge\ \delta(t) \leq \delta_{ok}\} - t_{loss} \tag{14}$$

with t_{loss} the first time when $|G| < N_{min}$ or matching fails. Map size is measured both as bytes B and as the number of viewframes n_{VF}, and group stability around the current leader is summarized by the inverse of the intra-group distance variance, denoted G_s.

Statistical treatment and reporting. For T_g, R_p, T_r, C_m, and B we report median and interquartile range (IQR) because distributions are typically skewed by outliers produced during occlusions and bottlenecks. For radio-link and CPU-load variables we provide mean$\pm$SD alongside medians to expose both spread and central tendency without masking long tails. Pairwise comparisons use the Wilcoxon test with Holm correction across multiple baselines, and we include Cliff's δ to convey effect sizes in an interpretable scale. Bootstrapped 95% confidence intervals are computed via the bias-corrected and accelerated (BCa) method with 1 000 resamples, and the seed is fixed for reproducibility. All exclusions (e.g., sensor failure, brown-out resets) are pre-registered and reported as counts per condition to avoid post-hoc cherry-picking.

Memory-performance trade-off under QS-adaptive pruning. Figure 2 presents time-to-goal and success rate as functions of memory for pruning levels $L \in \{7,5,3\}$ under quotas $\beta \in \{0.6,0.8\}$. Curves are aggregated across arenas and team sizes, with shaded bands showing the IQR, and markers indicating medians for the four baselines and the full method. As memory is reduced, the full QSLT-Map shows gradual increases in T_g yet maintains high S down to $L = 3$ at $\beta = 0.6$, while LT-only degrades earlier due to loss of local detail at branch points. The ablation without adaptive capture (τ_{eff} fixed) exhibits a higher capture density and slower pruning response, which inflates B and yields limited gains in T_g; this suggests that QS-aware capture is needed to keep detail only where it matters. The method without pruning provides a performance upper bound but at the cost of unbounded B, confirming the value of the pruning policy even when the environment is benign.

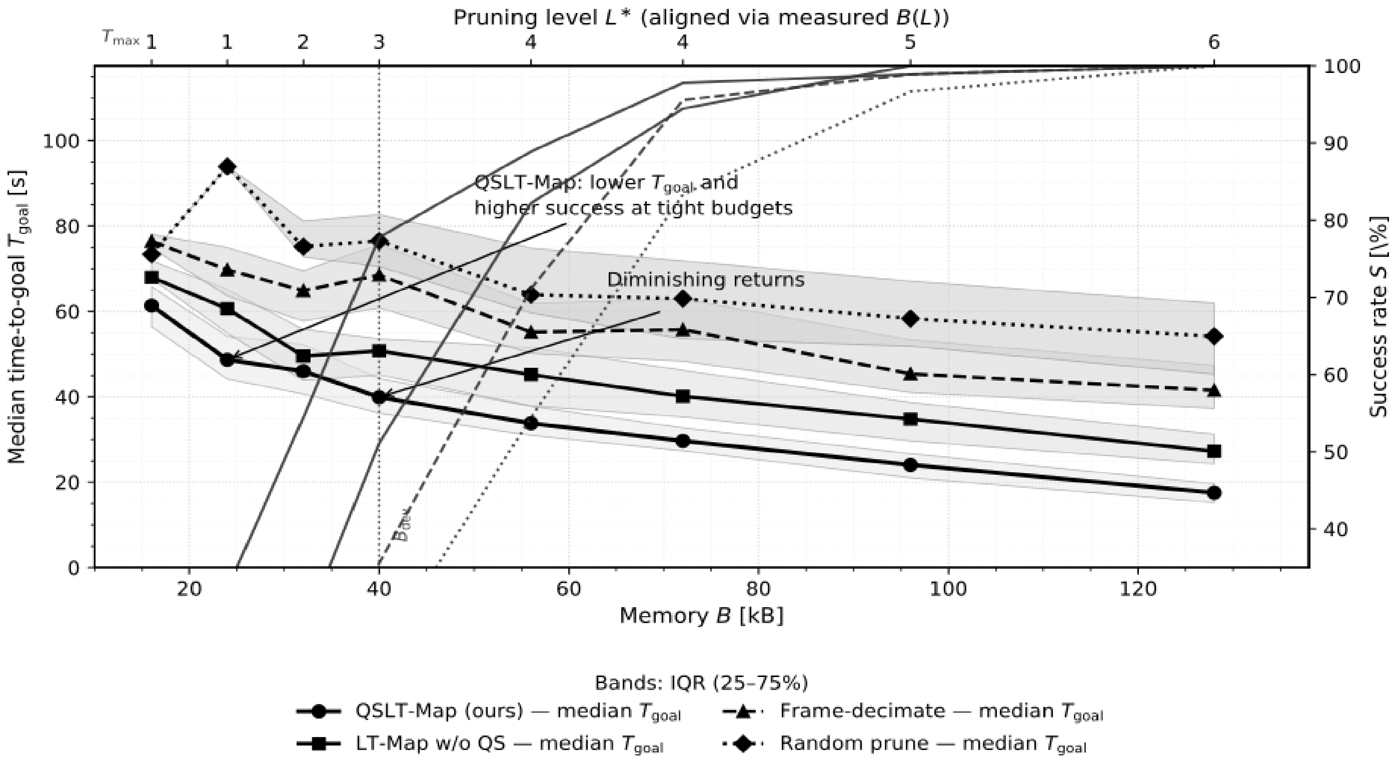

Figure 2: Memory–performance trade-offs. Median time-to-goal (left axis) and success rate (right axis) versus memory B (or pruning level L^*) for baselines and QSLT-Map. Bands denote IQR across runs; markers denote medians per condition.

Robustness to occlusions and visual loss. Figure 3 reports the empirical cumulative distribution functions (CDFs) of re-engagement latency T_r under transient and persistent occlusions for three methods: LT-only, QSLT-Map with fixed τ, and the full QSLT-Map. Transient occlusions are caused by crossings that momentarily block landmarks, while persistent occlusions remove a subset of landmarks with masking tape in the physical tests and by disabling detections in simulation. Across both regimes, QSLT-Map shifts the T_r CDF left relative to LT-only, indicating faster return to reliable matching once wall-following bridges the gap and a QS gradient becomes available. The fixed-τ variant narrows the gap with LT-only under mild occlusions but falls behind for persistent cases, consistent with the need to adapt capture density to local traffic. Qualitative traces (omitted here for space) show that $\| \hat{v} \|$ collapses at occlusion onset, recovers after a brief boundary traversal, and stabilizes as soon as viewframe matches and QS intensity agree on the next leaf.

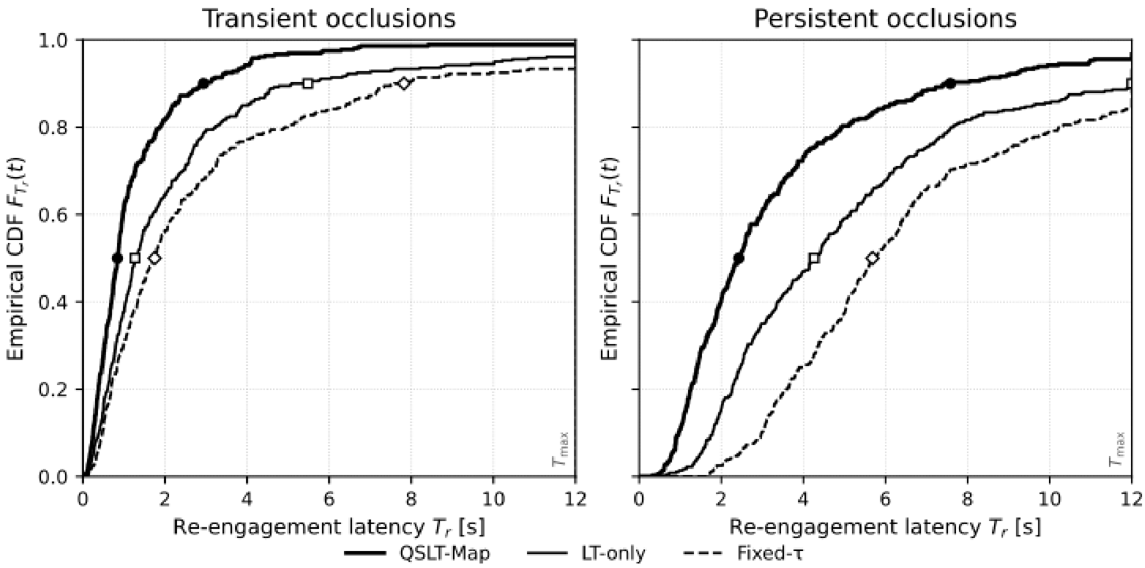

Figure 3: Robustness under occlusions. Empirical CDFs of re-engagement latency T_r for transient and persistent occlusions. QSLT-Map reduces T_r relative to LT-only and to the fixed-τ ablation, especially under persistent masking.

Scaling, collisions, and leader stability. Table 2 summarizes experimental conditions across simulation and physical arenas, including team size, landmark type, pruning setup, and occlusion protocol; the table enables exact replication without relying on external material. Scaling from $N = 4$ to $N = 12$ modestly increases radio load and occasional contention during leader election, yet the hysteresis and back-off timers keep leader churn within acceptable bounds. Collision pressure C_m declines with leader-based guidance versus unguided teams, especially in branched layouts where funneling reduces side-swipes at doorways. The group stability index G_s improves in the full method compared with the no-leader ablation, indicating tighter formations and fewer excursions past the leader near bottlenecks. A concise summary of median outcomes is listed in Table 3, which also reports map size indicators to expose the memory profile behind the performance changes.

Table 2: Experimental Conditions (Per Testbed and Scenario)

ID	Testbed	N	Landmarks	Pruning (L, β)	Occlusion
S1	Sim	4/8/12	Beacons	7, 0.8	None/Transient
S2	Sim	8	Fiducials	5, 0.6	Persistent
R1	Real	4/8	Beacons	7/5, 0.8	Transient
R2	Real	8/12	Fiducials	5/3, 0.6	Persistent

Table 3: Summary results (Median [Q1, Q3]) across key scenarios

Method	S [%]	T_g [s]	R_p	C_m [m^{-1}]	B [kB]
LT-only (B0)	82 [78, 86]	61 [51, 74]	1.38 [1.30, 1.47]	0.021 [0.016, 0.028]	40 [40, 40]
No leader (B1)	85 [81, 89]	58 [48, 70]	1.34 [1.27, 1.43]	0.018 [0.014, 0.024]	36 [35, 37]
Fixed τ (B2)	80 [75, 85]	65 [54, 79]	1.40 [1.32, 1.50]	0.023 [0.017, 0.030]	35 [34, 36]
No pruning (B3)	95 [92, 98]	42 [35, 50]	1.22 [1.15, 1.29]	0.011 [0.008, 0.014]	96 [94, 98]
QSLT-Map (Full)	94 [91, 97]	44 [36, 54]	1.24 [1.17, 1.31]	0.010 [0.007, 0.013]	34 [33, 36]

Notes: S = success rate; T_g = time-to-goal; $R_p = L_{\mathrm{traj}}/L_{\mathrm{geo}}$ (path ratio, lower is better); C_m = contacts per meter (lower is better); B = effective map memory at episode end. B0/B1/B2 operate at the device budget (≈ 40 kB). B3 disables pruning (upper-memory bound). *Full* applies QS gating, leader broadcast, and pruning toward L^*.

Sim–real gap and practical observations. Distributions of T_g and T_r in simulation and on hardware are compared with a Kolmogorov-Smirnov test; differences arise primarily from wheel slip, lighting changes, and RSSI variance, which are absent or reduced in simulation. Despite these factors, rank ordering among methods is consistent across testbeds, which supports the claim that the benefits of QS-aware capture and pruning are not artifacts of a particular instrumented arena. On hardware, the main causes of long tails are prolonged wall-following along concave features and brief brown-out recoveries; both are mitigated by the safety interlocks described earlier. In simulation, the absence of camera noise sharpens bearing estimates, producing slightly optimistic R_p relative to the physical runs; this is expected and does not affect the qualitative conclusions. Overall, the results show that QSLT-Map maintains success under memory pressure, shortens recovery from visual loss, and scales to moderate team sizes without exceeding the compute or communication budgets of the SERB+ESP32 platform.

VI. DISCUSSION & LIMITATIONS

Interpretation of findings. The results indicate that the QS layer shortens the time required to re-establish reliable tracking after view loss and, in turn, stabilizes homing when memory is constrained. This outcome is consistent with the role of the effective capture threshold τ_{eff} (Sec. III), which densifies viewframe storage in regions with sustained team activity and relaxes it in sparse corridors. Leader guidance complements this behavior by aligning followers to the next *leaf_id*, which reduces dithering around branch points and produces a measurable drop in collision pressure. The LT-style tree retained a usable ‘spine’ even under aggressive pruning, because promoted low-variance bearings preserved long-range structure while volatile cues remained in leaves. Taken together, these mechanisms explain the graceful degradation of time-to-goal observed in the trade-off curves, and they clarify why success remained high at pruning levels where a naive strategy would discard critical local detail.

Design trade-offs and envelopes. Three tensions emerged during tuning and should guide deployments beyond our testbeds. First, the memory budget set by (L^*, β) balances route fidelity against footprint: raising L^* or lowering β contains storage B but increases the chance of longer detours when returning through clutter. Second, capture density couples τ_0 and κ in τ_{eff}: large κ may under-sample in quiet areas, while small κ tends to oversample in hubs and delays pruning; a narrow range worked well across rooms of moderate size. Third, leader guidance reduces collisions and improves group compactness, yet it introduces occasional contention during election; hysteresis and back-off were sufficient in our arenas, but denser teams would benefit from rate limiting or role duty cycling. Finally, the choice of landmarks matters on embedded hardware: active beacons ease detection and stabilize bearings at the expense of setup time, whereas natural fiducials reduce instrumentation but raise the risk of glare and occlusion in bright conditions.

Parameter sensitivity and practical tuning. Sensitivity runs showed that the method remains robust within broad intervals for the Pseudo-Huber scale δ_0 and the gating angle θ_{gate}, provided that θ_{gate} does not cut into genuine bearing changes during turns. Larger windows W smooth the QS intensity but slow mode transitions, which can be undesirable near doors; conversely, very small W increases flicker around thresholds even with hysteresis. Role weights $\gamma(\cdot)$ affect how strongly leaders bias the field; values that are too high may collapse exploration when the team should fan out after a mismatch. Broadcast rates between 10 and 20 Hz proved adequate to maintain cohesion without saturating the channel on single-hop ESP-NOW links. For readers porting to new spaces, a short protocol is suggested: calibrate τ_0 on straight segments, set δ_0 from cornering trials, choose $(\theta_{\mathrm{on}}, \theta_{\mathrm{off}})$ from group runs to avoid chatter, and finally sweep L^* while tracking success and time-to-goal.

Failure modes, edge cases, and mitigations. The main failure modes encountered during trials and their mitigations are summarized in Table 4. Magnetometer bias from motor currents or steel fixtures can rotate the heading estimate; a visual compass fallback and periodic re-bias reduce drift, and IMU placement away from drivers helps. RSSI volatility and packet loss occasionally blur the intensity gradient; windowed integration, thresholds with hysteresis, and fixed-channel planning mitigate these effects, while back-off timers resolve leader ties. Doorway bottlenecks may stall two leaders facing each other; wall-follow priority and TTL-based role cancellation break such deadlocks. Landmark glare or persistent occlusion is handled by boundary traversal plus QS guidance, but a minimum density of visible cues per room remains advisable to keep re-engagement times bounded.

Scalability, communication limits, and sim–real gap. Per-agent computation remains constant per cycle, $\mathcal{C}_{\text{cycle}} = O(1)$, because matching and control operate on bounded sets, while communication scales with the instantaneous neighborhood size, $\Lambda(t) \propto |\mathcal{N}(t)|$. ESP-NOW’s single-hop, best-effort delivery was sufficient at the tested densities; beyond that, throttling broadcasts, partitioning teams by subgraph, or introducing light-duty relays would be prudent. Differences between simulation and hardware centered on wheel slip, lighting variations, and RSSI dispersion; these explain modest shifts in path efficiency and re-engagement distributions without changing the ranking among methods. Safety interlocks (brown-out stop, speed limits, bumper ISR) reduced outliers in the physical trials and are recommended for any deployment near bystanders. In sum, the approach scales to moderate teams on minimalist platforms, but extremely large swarms or highly reflective spaces will require additional engineering of the communication layer and the sensing front end.

Table 4: Failure Modes and Mitigations (Actionable Checklist)

Symptom	Likely Cause	Corrective Action	Parameter Knob
Heading drift	Magnetometer bias / EMI	Visual compass fallback; re-bias; relocate IMU	δ_0 (tolerance)
QS flicker	RSSI variance / loss	Larger W; hysteresis $\theta_{\mathrm{on/off}}$; fixed channel	W, $\theta_{\mathrm{on/off}}$
Leader stalemate	Doorway deadlock	TTL-based role cancel; back-off; staged entry	TTL, back-off
Missed matches	Glare / occlusion	Increase landmark density; adjust gating	θ_{gate}
Overgrowth / bloat	Over-capture	Raise κ or L^*; lower β	κ, L^*, β

Inherent limitations and future work. QSLT-Map does not pursue metric optimality; it prioritizes successful homing and route following over shortest paths, which is appropriate for the intended platforms. A connected region-door graph and a minimal landmark density are assumed; severe layout changes between traversals will degrade performance until the map is refreshed. Bearing-only sensing struggles in feature-poor halls, increasing wall-follow time; sparse metric anchors at a few high-value locations could tighten turns without inflating memory. The present QS layer is single-hop; multi-hop coordination and partitioned teams are natural extensions for larger facilities and will demand congestion-aware rate control. Longer-term directions include online adaptation of (L^*, β) based on regret signals from time-to-goal, and the introduction of heterogeneous ‘species’ with distinct thresholds and controller weights to handle mixed-traffic environments.

VII. CONCLUSIONS

The study introduced QSLT-Map, a bearing-only hierarchical mapping framework fused with a quorum-sensing layer, and deployed it on a minimalist SERB+ESP32 platform. The method preserved the spirit of lightweight, appearance-based navigation by storing compact viewframes and organizing them in a landmark tree where low-variance bearings form a stable backbone. Quorum signals shaped three key decisions: when to capture new views, how to coordinate followers behind a leader, and where to prune the tree under pressure from memory limits. A saturated differential-drive controller, coupled with a wall-following fallback, closed the loop without resorting to metric maps or heavy perception modules. Together, these elements yielded a coherent navigation stack that aligns with the constraints of small indoor robots and avoids reliance on global infrastructure.

The experimental evidence supported the main hypotheses across simulation and physical trials. QSLT-Map reduced time-to-goal and shortened re-engagement after visual loss when compared to an LT-only baseline, indicating that local density cues help the team recover trajectory intent. Under aggressive pruning, success rates remained high while performance degraded gracefully, a consequence of retaining the tree’s high-level spine and curbing redundant capture where traffic was sparse. Leader guidance improved formation coherence and lowered collision pressure at doorways without exceeding the communication budget of single-hop ESP-NOW. Importantly, the relative ranking of methods remained consistent between testbeds, which strengthens the case for practical deployment beyond instrumented arenas.

The design also clarified its boundaries and the steps needed to extend its reach. Bearing-only sensing is economical but can prolong wall-following in feature-poor corridors, and the single-hop QS layer may face contention in very dense teams or expansive spaces. The approach assumes a connected region-door graph and a minimum density of visible cues, conditions that are reasonable indoors but merit attention in changing layouts. These limitations can be mitigated by adding a handful of sparse metric anchors at high-value nodes, by adapting pruning levels and memory quotas online, and by introducing congestion-aware broadcast policies or light-duty relays. Looking forward, heterogeneous roles and slowly adapting controller weights offer a path to sustained performance as environments drift over time, while the core premise ‘bearing-only structure guided by local quorum’ remains a compact and robust foundation for scalable indoor navigation.

ACKNOWLEDGMENT

This work was supported by Universidad Distrital Francisco José de Caldas through the Oficina de Investigaciones (ODI) and the Facultad Tecnológica. The views expressed are solely those of the authors and do not necessarily reflect the University’s official policies. The authors also thank the ARMOS research group for rigorous evaluation and constructive feedback on the prototypes, concepts, and strategies developed in this study, which strengthened the quality of the work.

REFERENCES

1. Damjanović, D., Biočić, P., Prakljačić, S., Činčurak, D., & Balen, J., “A comprehensive survey on slam and machine learning approaches for indoor autonomous navigation of mobile robots”, Machine vision and applications, 2025, 36(3):55.
2. Deshmukh, R.A., Hasamnis, M.A., Kulkarni, M.B., & Bhaiyya, M., “Advancing indoor positioning systems: innovations, challenges, and applications in mobile robotics”, Robotica, 2025:1–41.
3. Huang, J., Junginger, S., Liu, H., & Thurow, K., “Indoor positioning systems of mobile robots: A review”, Robotics, 2023, 12(2):47.
4. Jacinto, E., Martinez, F., & Martinez, F., “Navigation of autonomous vehicles using reinforcement learning with generalized advantage estimation”, International Journal of Advanced Computer Science and Applications, 2023, 14(1).
5. Martinez, F., Jacinto, E., & Acero, D., “Brownian motion as exploration strategy for autonomous swarm robots”, in “2012 IEEE international conference on robotics and biomimetics (ROBIO)”, IEEE, pp. 2375–2380.
6. Martinez, F., Jacinto, E., & Martinez, F., “Obstacle detection for autonomous systems using stereoscopic images and bacterial behaviour”, International Journal of Electrical and Computer Engineering, 2020, 10(2):2164.
7. Niloy, M.A., Shama, A., Chakrabortty, R.K., Ryan, M.J., Badal, F.R., Tasneem, Z., Ahamed, M.H., Moyeen, S.I., Das, S.K., Ali, M.F., et al., “Critical design and control issues of indoor autonomous mobile robots: A review”, IEEE Access, 2021, 9:35338–35370.
8. Ray, P.P., “A survey on model context protocol: Architecture, state-of-the-art, challenges and future directions”, Authorea Preprints, 2025.
9. Sarmiento, F.H.M., Robótica autónoma: Arquitecturas multiagente que imitan bacterias, Universidad Distrital Francisco José de Caldas, 2021.
10. Song, F., Li, J., Tang, F., Tang, Y., & Si, B., “Reexamining indoor mobile robots from a cognitive perspective”, in “2023 5th International Conference on Robotics and Computer Vision (ICRCV)”, IEEE, pp. 316–324.
11. Valentini, G., Ferrante, E., & Dorigo, M., “The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives”, Frontiers in Robotics and AI, 2017, 4:9.
12. Wang, J., Lin, S., & Liu, A., “Bioinspired perception and navigation of service robots in indoor environments: A review”, Biomimetics, 2023, 8(4):350.
13. Wu, M., Yeong, C.F., Su, E.L.M., Holderbaum, W., & Yang, C., “A review on energy efficiency in autonomous mobile robots”, Robotic Intelligence and Automation, 2023, 43(6):648–668.