

## Article

# Spatial All-Azimuth Versus Single-Sided Planar Identifiers for Warehouse Robot Navigation: A Factorial Simulation Study

Kamil Kušnirák , Oto Haffner \* , Erik Kučera  and Ondrej Kolimár 

Faculty of Electrical Engineering and Information Technology, Slovak University of Technology in Bratislava, 841 04 Bratislava, Slovakia; kamil.kusnirak@stuba.sk (K.K.); erik.kucera@stuba.sk (E.K.); ondrej.kolimar@stuba.sk (O.K.)

\* Correspondence: oto.haffner@stuba.sk

## Abstract

A camera-guided warehouse robot keeps its bearings by repeatedly estimating its pose against known visual references, and it must relocalize whenever that estimate is lost. What limits this process is often not identification but the availability of a usable reference along the route. The references used in practice are usually single-sided planar fiducial markers such as QR-like codes, ArUco markers, and AprilTags, which stay readable only within a limited cone about their surface normal; a spatial reference, by contrast, can in principle be recognized from any azimuth. We quantify what that difference is worth at the navigation level. The framework is built in Unity with NavMesh navigation and a purely geometric-availability model, rather than an image-based recognizer, whose single switchable property is the availability rule. In the idealized all-azimuth spatial-reference regime (the spatial regime), a reference is available from any direction; in the single-sided, angularly constrained planar-reference regime (the planar regime) it is available only within  $\pm 20^\circ$  of the surface normal. A full-factorial experiment with 54 configurations ( $3 \times 3 \times 2 \times 3$ ) and  $n = 100$  paired replications, 10,800 runs in all, was run in both regimes over four deployment factors: deployment scheme, camera field of view, recovery step, and detection range. Under this geometric model, the spatial regime reached  $5.7\times$  higher reference coverage (41.6% vs. 7.3%) and a mission-completion rate 30 percentage points higher (86.2% vs. 55.9%). A paired Wilcoxon signed-rank test confirms the coverage difference ( $p < 0.001$ , matched-pairs  $d_z = 1.84$ ), and McNemar's test together with a logistic regression confirms the completion difference. In a factorial analysis of variance, the detection range dominates (partial  $\eta^2 = 0.903$ ), and a strong deployment  $\times$  range interaction concentrates the advantage in the rack aisles, where a planar reference is seen edge-on. Three further analyses point the same way: an angular-threshold sweep from  $10^\circ$  to  $60^\circ$ , an equal-count deployment control, and route- and time-normalized visibility and relocalization metrics. The advantage also held across square, L-shaped, and U-shaped aisle layouts (32,400 runs in total), with a negligible regime  $\times$  layout interaction. All these numbers are model-based estimates under an explicitly stated availability model: they measure the navigation-level value of azimuthal reference availability and do not validate any particular physical object, decoding algorithm, or AR device.



Academic Editors: Ivo Silva and Cristiano Pendão

Received: 19 June 2026

Revised: 4 August 2026

Accepted: 6 August 2026

Published: 7 August 2026

**Copyright:** © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

**Keywords:** warehouse robot navigation; visual relocalization; spatial reference; planar fiducial marker; all-azimuth availability; angular readability; navigation robustness; factorial experiment; geometric simulation; Unity

## 1. Introduction

A camera-equipped warehouse robot keeps track of where it is by repeatedly estimating its pose relative to known visual references placed in the environment; whenever it loses that estimate, it must *relocalize* against a reference whose global pose is known. The practical bottleneck is therefore not identification alone, but the *availability* of a usable localization or relocalization reference along the route [1,2]. The references most commonly used for this purpose are *planar fiducial markers* (QR-like codes, ArUco markers, AprilTags), which are inexpensive, easy to print, and support both identification and estimation of the camera pose relative to the marker [3–5].

A planar fiducial carries a deep geometric limitation, however: it has a single privileged surface normal and can be read only within a limited cone about it, so its projected image foreshortens, and its pose estimate degrades as the viewing direction becomes oblique [3]. Along a warehouse aisle, a rack-facing marker is seen mostly edge-on, so for much of the pass no usable reference is available from that marker. This angular limitation can be mitigated by distributing several planar markers over a three-dimensional body (a multi-sided fiducial object), or, more abstractly, using a spatial reference whose recognizable appearance and consistent object coordinate frame are available from an arbitrary azimuth. In this article, a *spatial reference* denotes a deliberately designed, asymmetric three-dimensional object with a distinctive appearance and a known global pose; its all-azimuth availability is treated as a *modelled geometric assumption*, not as a validated property of a specific physical object or detector. We do *not* compare marker payload, data capacity or the performance of any particular decoding or recognition algorithm; the comparison is deliberately technology-neutral and concerns only the geometric availability of a localization reference.

The reliability of marker-based navigation is, in the literature, discussed almost exclusively in qualitative terms; a controlled, reproducible, navigation-level comparison of angularly constrained and all-azimuth references has been missing. This study closes that gap by comparing, in a common factorial simulation framework, two idealized detection regimes:

- an **idealized all-azimuth spatial-reference regime** (*spatial regime*), in which a reference is available from any azimuth subject to range, field of view, and line of sight; and
- a **single-sided, angularly constrained planar-reference regime** (*planar regime*), in which a reference is additionally available only when the viewing direction lies within a modelled readability cone about its surface normal.

QR codes, ArUco markers and AprilTags are examples of the planar regime; the simulation models only the geometric availability of the reference, common to all of them, and not any specific implementation. The main research question is:

*How does replacing a single-sided, angularly constrained planar reference with an all-azimuth spatial reference affect relocalization availability and warehouse-navigation robustness?*

In line with the underlying research assignment, we do not develop a physical spatial reference; the goal is to **propose and evaluate a method for deploying such references in warehouses** and to quantify the navigation-level value of azimuthal reference availability. The contribution of this article is fourfold:

- a reproducible, simulation-based methodology that *quantitatively compares* the navigation robustness of an all-azimuth spatial reference and a single-sided planar reference in a unified full-factorial framework, under an explicitly stated geometric-availability model;
- a statistically verified demonstration that, *under the reference geometric model*, the spatial regime yields substantially higher reference availability and mission completion, together with a factorial analysis that identifies *detection range* as the dominant deployment parameter;

- a set of targeted robustness analyses—an angular-threshold sensitivity sweep, an equal-count deployment control, operational visibility and relocalization metrics, and a layout-robustness replication across square, L-shaped and U-shaped warehouses; and
- a geometric explanation of *why* and *where* the advantage arises, turning the result into a deployment guideline whose principal levers, within the present geometric model, are local reference spacing and detection range.

The study isolates a single modelling difference—all-azimuth availability for the spatial regime versus a cone-limited availability for the planar regime—while all navigation, deployment, route generation and statistical procedures are identical between the regimes. Accordingly, it does not claim that a specific physical object or computer-vision algorithm achieves the reported values; it quantifies the navigation-level *consequence* of an angular-availability assumption. Within Industry 4.0, augmented-reality (AR) tracking is one of several enabling technologies; the connection to the present work is not the display of graphical overlays to a human operator but the shared visual-reference and relative-pose principles that also underpin robot localization.

The remainder of the article is organized as follows. Section 2 reviews related work on AR tracking and simulation-based evaluation. Section 3 describes the simulation framework, the experimental scenario, and the factorial design. Section 4 details the system architecture and the geometric argument that underlies the approach. Section 5 reports the quantitative results, and Sections 6 and 7 discuss the findings and conclude.

## 2. Related Works

This section reviews the literature in the areas relevant to the present study: the role of augmented reality (AR) within Industry 4.0, the visual-tracking principles that underpin marker-based and markerless AR, planar fiducial markers and multi-sided spatial references, and the simulation-based evaluation of robotic perception before physical deployment.

### 2.1. Industry 4.0 and Augmented Reality

The Industry 4.0 (I4.0) concept transforms the contemporary industrial environment through the digitalization of production, automation, and the integration of the plant into a complex supply chain; it rests on full network integration and the exchange of information in real time [6,7]. Emerging technologies such as the Internet of Things (IoT), Big Data and augmented reality are combined with traditional manufacturing concepts [8,9] to build smart factories that optimize production by minimizing cost, maximizing quality and reacting agilely to rapid market change [10–12]; reaching this state requires existing processes to be reworked and improved through new technologies and procedures [13].

One of the most significant transformations is in the work environment itself—in human–machine collaboration, where the physical and virtual worlds merge. Virtual animations and simulations of products and processes, together with the visualization of raw and computed data, are key aspects of realizing a cyber-physical system (CPS) [14,15]. Conventional means of controlling and diagnosing CPS in IoT networks—industrial panels and consoles, web or mobile applications—are adequate in a small room, where the device list fits on a single screen, but become cumbersome across multiple rooms or buildings, as in digital factories or large warehouse systems. Combining AR with suitable data can improve vertical integration within the factory [16]: within Industry 4.0, AR is *one of several* enabling technologies—alongside industrial communication, embedded control, the IoT, data analytics and artificial intelligence—and it contributes primarily spatial-visualization and interaction capabilities that support the human operator [17,18], helping turn ordinary workers into the “smart operators” of the future who handle strategic decision-making and flexible problem-solving [19,20]. The essence of AR is presenting the right information in

the right place; the challenge is to develop AR applications that shorten production times, reduce costs, and improve quality [21,22].

## 2.2. Augmented Reality and Visual Tracking

Augmented reality overlays, combines and interactively activates virtual and real components within the current scene in real time and in three dimensions; the conceptual foundations of this reality–virtuality continuum were laid by Milgram and Kishino [23] and developed in numerous surveys [24–28]. Virtual components are visualized through optical see-through (OST) devices such as tablets, video see-through (VST) head-mounted displays (HMDs), or external spatial projection. Because the virtual objects are computer-generated, they can be contextually overlaid and spatially registered onto real parts of the physical world in real time, which lets the user perform inspection and maintenance and eases collaboration with remote machines and human experts—yielding higher accuracy, time and cost savings, and better visualization, especially for novice operators. In industry, AR has accordingly been deployed for factory planning [29], assembly guidance [30], inspection [31] and on-site worker support [32].

The technical principle of vision-based AR is *tracking*: determining the pose (position and orientation) of the camera relative to a landmark in the real environment. The 6D camera pose is estimated by matching recognized feature points in each video frame against those of the landmark and computing the homography between the two images; feature descriptors are then used to match the points in real time and to place the virtual object [25,26]. Such pose estimation is invariant to rotation and scale and reasonably robust to partial occlusion and lighting changes. HMD tracking additionally estimates the device pose relative to its previous pose, and HMDs equipped with inertial measurement units (IMUs) combine visual and inertial sensing into visual–inertial tracking [33]; anchoring the virtual 3D content to the real scene relies on depth mapping based on structured light, stereo vision, or time-of-flight.

The tracking principles used in AR are not limited to human-facing overlays. Estimating the relative pose between a camera and a known visual reference is equally applicable to robot localization: a detected reference of known global pose lets a mobile robot re-anchor its own pose estimate. In this study, the connection to AR therefore lies in these shared visual-reference and relative-pose principles—the AR-related visual tracking commonly used in AR systems—rather than in the presentation of graphical overlays to a human operator; the study is not concerned with classical AR navigation using holographic guidance.

## 2.3. Marker-Based Versus Markerless Augmented Reality

The literature distinguishes *marker-based* (MB) and *markerless* (ML) AR [34]. In MB AR, the digital world is anchored to the real world directly through a marker: the device must first recognize what the user is looking at, which is achieved by placing a characteristic image or shape—the marker—in the scene; any image with enough unique visual points (corners and edges) can serve. In ML AR, the user freely places virtual objects that are not anchored to any specific point in the real scene, as when previewing a piece of furniture in an empty room.

Vision-based marker tracking uses *fiducial markers*—artificial planar elements with known properties (shape, color, size), typically an outer frame enclosing an inner pattern that encodes information, designed to satisfy criteria such as occlusion resilience. Planar fiducial references include QR-like codes, ArUco markers, AprilTags, and other purpose-designed patterns used for identification and relative-pose estimation [3–5,35–37]. Identifier-based approaches are the most widely used tracking method for indoor AR: they require no complex hardware, are computationally light, real-time, and simple. They

also have drawbacks. Poor visual conditions, defocus, or motion blur reduce the number of recognizable features, and very strong lighting causes surface reflections that lead to misrecognition and misalignment of virtual and real objects. A single planar fiducial marker can support estimation of the camera pose relative to the marker plane; its practical observability and estimation quality, however, deteriorate under unfavorable viewing geometry, especially at oblique incidence angles. Additional differently oriented markers or a calibrated multi-sided fiducial object can be used to expand viewpoint coverage [38,39].

#### 2.4. Model-Based (3D) Tracking

Recently, a tracking principle based on *models* (3D object shapes) has emerged in academia and industry alike. Model-based tracking uses the contour of a CAD model as the reference for image recognition, displaying 3D content tied to a physical object *without* QR codes or tracking devices [40,41]. Its advantage is that the environment or object need not be modified (no QR codes attached), the digital overlay is highly accurate, and all surfaces of the detected object can be used for interaction; objects are detected and tracked without a feature-point map, which otherwise fails for low-texture or moving objects. Model-based tracking uses known object geometry as a reference for identification and pose estimation. Its performance depends on the object geometry, viewpoint, image quality, occlusion, illumination, and the recognition algorithm. Unlike a single planar marker, a suitably designed asymmetric spatial reference may provide recognizable geometric information from multiple viewing directions, but this property requires empirical validation for a specific physical object and detector [42,43]. Its limitations are dynamic objects—a static CAD model may diverge from an object with moving parts—and physically identical objects, which model-based tracking cannot distinguish (e.g., machines of the same type) [21].

This distinction is directly relevant to the present work. El Barhoumi et al. evaluated several AR content-placement methods quantitatively and qualitatively and found that markerless methods based on a 3D model achieved results equal to or better than tracking using 2D markers [44]. A survey of industrial AR likewise reports that only about 16% of industrial AR tracking is based on 3D recognition versus 57% on (unspecified) markers [22], indicating that the spatial alternative remains under-exploited. The spatial reference considered in this study is not an arbitrary CAD object; it is a conceptual, deliberately designed asymmetric object intended to provide a distinctive appearance and a consistent object coordinate frame across azimuthal viewpoints. It is conceptually compatible with model-based tracking, but in the present study its recognition process is abstracted as an idealized all-azimuth geometric-availability model rather than implemented through an image-based detector. We previously proposed the concept and an implementation outline for software-generated spatial references in logistics [45]; the present article extends that proposal from a conceptual design to a statistically grounded, full-factorial comparison that quantifies, for warehouse robot navigation, how much more robust the spatial reference is than a planar reference.

#### 2.5. Simulation-Based Evaluation of Robotic Perception and Navigation

Because physical AR/AGV experiments are expensive and hard to repeat, game engines have become a standard tool for evaluating perception and navigation. Unity [46], with its NavMesh navigation stack [47,48], provides a deterministic environment in which thousands of runs can be executed in batch mode, and comparative studies confirm its suitability for autonomous-vehicle development [49]. Game-engine simulation can provide controlled and repeatable geometry, navigation and line-of-sight conditions, but it does not by itself establish realistic image-based recognition performance: lighting, material appearance, exposure, shadows, motion blur and detector-specific behavior require rendered-image or physical validation. Unity is therefore used here as a controlled envi-

ronment for evaluating the navigation-level consequences of explicitly stated geometric assumptions, not as evidence of real-deployment readiness. In warehouse logistics specifically, user studies show that AR smart glasses measurably affect picking time, error rate, cognitive load, and worker well-being [50,51], underscoring the importance of reliable underlying perception. The closest prior work to ours is the thesis of Kiss [52], which simulates a mobile robot in Unity using the Vuforia engine and NavMesh navigation; that work, however, analyses individual runs and does not compare identifier types, and so lacks the statistical basis required to characterize robustness. The present article preserves the same methodological foundation but advances it to a full-factorial, statistically grounded comparison of spatial and flat identifiers.

## 2.6. Planar Fiducial Markers and Multi-Sided Spatial References

Planar fiducial markers, including QR-like codes, ArUco markers, and AprilTags, provide visually distinctive planar references that can support identification and camera-pose estimation. When the marker geometry and the camera calibration are known, the projected marker corners can be used to estimate the relative pose of the camera with respect to the marker; if the marker pose is known in a global coordinate system, this relative transformation can support robot localization or relocalization. The quality and availability of a pose estimate, however, depend on viewing distance, incidence angle, image resolution, and other imaging conditions [38,53].

In the present study, the term *single-sided planar reference* denotes this common geometric property rather than a particular QR, ArUco, or AprilTag implementation. The simulation does not model marker payload, code dictionaries, corner-detection algorithms or image-decoding performance; instead, it represents a single planar reference with one surface normal and a constrained modelled readability cone [5].

The angular limitations of a single planar marker can be mitigated by combining several planar markers on a three-dimensional body. Such multi-sided fiducial objects use calibrated configurations of ArUco-, AprilTag- or related planar markers to improve viewpoint coverage and camera-pose estimation; experimental work on fiducial objects confirms that arranging markers over multiple faces can improve visibility from different viewpoints and provide pose information even when only part of the object is visible [54].

A sufficiently multi-sided planar object can therefore approximate all-azimuth reference availability. This observation motivates a technology-neutral formulation of the present comparison: the study evaluates an idealized all-azimuth spatial-reference regime against a single-sided, angularly constrained planar-reference regime. It does not claim that all spatial references must be realized as shape-only identifiers or that all planar systems must remain single-sided; rather, it quantifies the navigation-level value of azimuthal reference availability independently of the physical implementation used to obtain it.

## 2.7. Summary of the State of the Art

Three observations motivate the present study. First, visual references used for identification and pose estimation differ substantially in their viewpoint availability: a single planar reference has one privileged surface normal, whereas spatial or multi-sided references can provide observations from a wider set of directions. Second, the navigation-level consequences of this difference are rarely quantified under paired, controlled warehouse trajectories. Third, simulation enables systematic evaluation of range, field of view, deployment and recovery factors, but geometric simulations must be distinguished from image-based validation of a particular marker and recognition algorithm. The present study addresses this gap by comparing an idealized all-azimuth spatial-reference regime


with a single-sided angularly constrained planar-reference regime in a common factorial simulation framework.

### 3. Materials and Methods

#### 3.1. Experimental Scenario

The experimental scenario is an autonomous warehouse vehicle that traverses a route from an input station, through 15 semi-randomly placed goods items, to an output station in a square warehouse of  $100 \times 100$  m (axes  $X \times Z$  in Unity coordinates) with a regular grid of rack blocks arranged in three columns and sixteen rows (Figure 1). During the whole route, the vehicle orients itself by the deployed identifiers; at regular travel intervals it re-checks its orientation, and if it then finds no marker it enters a lost state and attempts to recover (Section 4.2). The square grid with narrow aisles was chosen because it is representative of real distribution centers and because its aisles force the vehicle to view rack-facing markers *edge-on*—the situation that distinguishes a spatial reference from a planar reference. The goods positions are seeded by the run index, so the  $k$ -th run of every configuration shares the same goods layout, which pairs the comparison across configurations and across both detection regimes.

Identifier deployment scheme for the three levels of factor A (coloured dots = identifier positions)




**Figure 1.** Plan view of the square warehouse ( $100 \times 100$  m) with the marker-deployment scheme for the three levels of factor A. Grey rectangles are rack blocks; the green and red squares are the input and output stations. Left: walls (40 markers on the perimeter); center: racks (288 markers on rack faces); right: Hybrid (250 markers). The scheme is identical for the spatial and the planar regime.


Figure 2 shows 3D renderings taken directly from the Unity scenes: the primary square layout (a) used throughout the main experiment, together with the L-shaped (b) and U-shaped (c) layouts used later in the layout-robustness analysis (Section 5.8). Each complements the schematic plan of Figure 1—the red-walled hall, the rack grid with its traversable aisles, and the deployed spatial references (blue). All elements are drawn to real scale; since the vehicle and the references are small relative to the hall, each panel includes a zoomed inset showing the vehicle (a forklift) at real size next to the rack-face references it passes.

Figure 3 summarizes the study as a systematic flow diagram. Panel (a) is the outer experimental pipeline—the 54-configuration factorial evaluated for both identifier regimes, the  $N$  paired replications per configuration, and the capture and analysis of the per-run records. Panel (b) is the per-run scenario itself: the vehicle plans a NavMesh route through the goods and drives it, performing a  $360^\circ$  orientation check every 100 m; a detected marker re-oriens it; otherwise, it enters a lost state and nudges, the run aborting after 100 failed

recovery sweeps and completing when it reaches the output station. The detailed detection and recovery rules are given in Section 4.2.



**Figure 2.** Oblique 3D renderings of the three simulated warehouse layouts, taken directly from the Unity scenes: (a) square (the primary layout), (b) L-shaped, and (c) U-shaped (used in the layout-robustness analysis). The red perimeter walls enclose the rack grid; the grey blocks are the rack rows, and the white lanes between them are the traversable aisles the vehicle drives. Blue cubes mark deployed spatial references at the candidate rack-face positions (at the two ends and the middle of each rack face); the cubes are visualization glyphs indicating reference positions and do not represent the proposed physical geometry of a spatial reference. Everything is drawn to real scale, so the vehicle (the forklift, circled orange) and the references are small relative to the hall; each panel therefore includes a zoomed inset showing the vehicle at real size next to the rack-face references it passes.



**Figure 3.** Systematic overview of the experiment. (a) Experiment pipeline: the 54-configuration factorial ( $3 \times 3 \times 2 \times 3$ ) run for both identifier regimes with  $N$  paired replications per configuration,

each producing a per-run JSON record that is aggregated into the paired statistics and figures. (b) Per-run scenario flowchart: NavMesh route planning, driving, the periodic 360° orientation check, the loss-and-nudge recovery loop, and the completed/aborted outcomes.

The scenario is a controlled benchmark rather than a digital twin of a particular warehouse; its purpose is to expose the vehicle repeatedly to the viewing geometries that distinguish the two regimes. The square 100 × 100 m layout is large enough to contain long aisle traversals, repeated line-of-sight interruptions, and several marker-spacing regimes while remaining computationally tractable for a full-factorial campaign, and the three-column, sixteen-row rack grid forces edge-on views of rack-facing markers. The number of goods was fixed at 15 to create multi-aisle routes long enough to expose the vehicle to many marker-viewing angles while keeping the run duration feasible for 10,800 simulations; this produced routes with a median length of about 1140 m (mean 1430 m, the difference reflecting the extra travel during recovery). Goods are sampled, using the run index as the random seed, from the accessible aisle regions only—excluding rack obstacles and the stations—and are then visited in a seeded Fisher–Yates random order (no travelling-salesman optimization), with NavMesh computing the shortest collision-free path between consecutive waypoints. The scenario is not claimed to cover all warehouse operating conditions; it isolates a canonical aisle-based geometry in which angular readability is expected to matter most. The main scenario parameters and the rationale for their values are summarized in Table 1.

**Table 1.** Scenario parameters and the rationale for their values. Lengths use the Unity convention 1 unit = 1 m; the settings are identical for the spatial and the planar regime, and the scenario is a controlled benchmark rather than a calibrated digital twin of a specific warehouse.

| Parameter                            | Value                                  | Rationale                                                               |
|--------------------------------------|----------------------------------------|-------------------------------------------------------------------------|
| Warehouse size                       | 100 × 100 m                            | Controlled benchmark with long aisle traversals                         |
| Rack layout                          | 3 columns × 16 rows                    | Forces repeated edge-on views of rack-facing references                 |
| Rack block ( $W \times H \times D$ ) | 28 × 5 × 2 m                           | Long rack faces viewed edge-on along the aisles                         |
| Aisle width                          | ≈ 4 m                                  | Narrow aisles typical of dense storage                                  |
| Identifier height                    | ≈ 3 m                                  | On rack faces/perimeter walls, near camera height                       |
| Camera height                        | 3 m                                    | Vehicle-mounted forward camera                                          |
| NavMesh agent                        | radius 0.75 m, height 3 m              | Idealized holonomic model on a 2D NavMesh                               |
| Goods per run                        | 15                                     | Multi-aisle routes while keeping batch runtime feasible                 |
| Goods placement                      | seeded, accessible aisles              | Route variability while preserving paired comparison                    |
| Start/end                            | fixed input/output stations            | Consistent mission objective                                            |
| Orientation-check interval           | 100 m travelled                        | Periodic re-check; passes short marker gaps, flags long blind stretches |
| Recovery scan                        | 360° at 90°/s                          | Local re-acquisition before relocation                                  |
| Recovery step (C)                    | 20/30 m                                | Coarse relocation after a failed scan                                   |
| Abort condition                      | ≥100 failed recovery scans (whole run) | Prevents infinite recovery loops                                        |

### 3.2. Simulation Framework

The framework is an integrated software system organized into five layers: a simulation core (Unity 6), a navigation layer (NavMesh + AI Navigation), a perception layer with a simulated geometric detector, a communication layer prepared for future deployment on real hardware (ROS2 for Unity [55], inactive during the reported experiments), and an analytical layer in Python 3.14.7. The architecture is detailed in Section 4. Because the Vuforia model-target recognition [41] requires a real camera image and cannot run in the headless batch mode needed for thousands of repetitions, it is replaced during the experiments by a deterministic geometric detector that models *visibility* (detection range, field of view, line of sight, and a short temporal persistence) rather than image processing.

### 3.3. The Spatial Reference and the Two Detection Regimes

Throughout this article, a *spatial reference* denotes a deliberately designed, asymmetric three-dimensional object with (i) a distinctive appearance intended to be recognizable across

azimuthal viewpoints, (ii) a unique identity, (iii) a known global pose in the warehouse map and (iv) a consistent object coordinate frame. Its availability from an arbitrary azimuth is treated as a *modelled geometric assumption*: the object is not a physically validated design in this study, and we deliberately distinguish the *concept* of a spatial reference, the *idealized geometric model* of it used here, and any *future physical implementation*.


The geometric detector has a single switchable property that constitutes the only difference between the two compared regimes. In both regimes, a reference is available at a given instant only if it lies within the detection range  $D$ , within the horizontal field of view, and in line of sight, and if these conditions persist for a short window (0.25 s). The regimes differ in one additional condition:

- **Idealized all-azimuth spatial-reference regime** (*spatial regime*)—the reference is available from any azimuth about its vertical axis; no incidence-angle condition is imposed.
- **Single-sided, angularly constrained planar-reference regime** (*planar regime*)—the reference is available only when, in addition, the viewing direction lies within a modelled readability half-angle of the reference surface normal, fixed at the reference value of  $\pm 20^\circ$  [3]; other half-angles are examined only in the angular-sensitivity analysis (Section 5.4).

QR codes, ArUco markers, and AprilTags are examples of the planar regime; the model represents only the geometric availability common to them, not any specific decoding or corner-detection algorithm. Multi-point line of sight is used: a reference counts as available if a clear ray reaches any point of a  $10 \times 10$  grid over its face, so partially occluded references still count. Because both regimes run in an identical factorial frame, sharing the same goods-layout seeds, target sequences, initial conditions, and navigation rules, any difference in the outcome is attributable solely to the availability model. The realized trajectories may nevertheless diverge owing to regime-dependent recovery events (Section 4.2), which is itself a consequence of the availability difference rather than a confound.

Figure 4 shows the vehicle's sensing footprint from above. Only geometry enters the model: a reference is available when it lies within the detection range  $D$  and the camera field of view  $B$ , is in line of sight, and (in the planar regime only) is seen within the readability cone about its surface normal. Rack-face references passed along an aisle are seen edge-on at a large incidence angle, so they fall outside the planar-readability cone while remaining available to a spatial reference. The figure illustrates this geometric model and is not an image processed by a vision algorithm.

Bird's-eye view of the vehicle sensing footprint in an aisle, drawn to real proportions (illustration of the geometric availability model, not a processed camera image)



**Figure 4.** Bird's-eye view of the vehicle's sensing footprint in an aisle, drawn to real proportions (rack block  $28 \times 2$  m, aisle 4 m, vehicle  $3.5 \times 1.5$  m). The vehicle (yellow, with its heading arrow) drives along

the aisle; the shaded wedge is the camera field of view ( $B$ ), clipped to the aisle corridor, and the dashed arc is the detection range ( $D$ ). References are deployed only at the two ends (the rack corners) and the middle of each rack face, as in the simulation. Blue squares are the references that fall inside the field of view and range—here the middle pair the vehicle is passing (dotted lines are the lines of sight to them)—while grey squares are outside the field of view or beyond range; the callout gives the incidence angle at one of the in-view references, which is large because rack-face references are seen edge-on. This is an illustration of the geometric-availability model, not an image processed by a vision algorithm.

### 3.4. Factorial Design

We ran a full-factorial experiment over four deployment factors with  $3 \times 3 \times 2 \times 3 = 54$  configurations (Table 2), each replicated  $n = 100$  times, giving 5400 runs per regime and 10,800 runs in total. The reference positions are fixed by the deployment scheme (factor A: walls = 40, racks = 288, Hybrid = 250 references). The *Hybrid* scheme is deliberately not the arithmetic sum  $40 + 288 = 328$ : rack-face markers that would coincide with, or sit directly behind, a perimeter wall are removed to avoid redundant, mutually occluding markers, which leaves 250. Only the goods (collection targets) are semi-random and seeded. The number of replications is sufficient for narrow 95% confidence intervals and a high-power factorial analysis of variance [56].

**Table 2.** Factors of the full-factorial experiment, their levels and operational meaning.

| Factor                | Levels             | Operational Meaning                    | Notes                                                                                                                                                                                                                                               |
|-----------------------|--------------------|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| A — deployment scheme | walls/racks/Hybrid | identifier-deployment strategy         | reference counts 40/288/250; <i>Hybrid</i> is not the set-theoretic union of the wall and rack layouts — it combines wall and rack-corner positions while removing rack positions adjacent to an existing wall identifier to avoid local redundancy |
| B — field of view     | 60/90/120°         | horizontal camera FOV                  | $\pm B/2$ from the robot heading                                                                                                                                                                                                                    |
| C — recovery step     | 20/30 m            | displacement during a recovery attempt | applied only after a lost state, not in normal driving                                                                                                                                                                                              |
| D — detection range   | 2/4/8 m            | maximum recognition distance           | dominant deployment factor (Section 5.3)                                                                                                                                                                                                            |

These three schemes are domain-informed *full* deployment layouts rather than equal-count layouts: wall identifiers are positioned around the perimeter to provide references near aisle exits and along perimeter travel, rack identifiers occupy eligible rack-corner positions naturally encountered during aisle traversal and recovery, and the Hybrid scheme combines both principles while removing locally redundant rack positions adjacent to wall identifiers. The primary experiment therefore answers how three practically designed full deployment schemes behave, whereas the equal-count control of Section 3.6 answers how the three deployment domains behave at an identical count of 40 identifiers and rules out that the spatial-regime advantage is a mere consequence of the differing counts.

### 3.5. Metrics and Hypotheses

For each run, we record the **reference coverage** (fraction of markers detected at least once), the **mission-completion rate** (reaching the output station vs. abort), the number of distinct references seen, the loss events, and the full run configuration for reproducibility. Four hypotheses were formulated (Table 3): H1, that the spatial reference achieves significantly higher coverage than the planar reference with a large effect ( $d > 0.8$ ); H2, that it achieves a higher completion rate; H3, that detection range is the dominant deployment factor; and H4, that the spatial-regime advantage is largest where markers are viewed edge-on (a strong deployment  $\times$  range interaction).

**Table 3.** Experimental hypotheses, the metric and statistical test used to evaluate each, the acceptance criterion, and the section that reports it.

| Hyp. | Statement                                                                         | Metric                 | Test/Model                       | Acceptance               | Sec.        |
|------|-----------------------------------------------------------------------------------|------------------------|----------------------------------|--------------------------|-------------|
| H1   | The spatial reference yields higher reference coverage than the planar reference. | Reference coverage [%] | paired Wilcoxon signed-rank      | $p < 0.05, d_z > 0.8$    | Section 5.1 |
| H2   | The spatial reference yields a higher mission-completion rate.                    | completed vs. aborted  | McNemar + logistic regression    | $p < 0.05, OR > 1$       | Section 5.1 |
| H3   | Detection range is the dominant deployment factor.                                | coverage [%]           | type-II factorial ANOVA          | largest partial $\eta^2$ | Section 5.3 |
| H4   | The spatial-regime advantage is largest for edge-on viewing.                      | coverage [%]           | factorial ANOVA ( $A \times D$ ) | strong $A \times D$      | Section 5.2 |

### Statistical Analysis

The replications within a configuration are independent draws of the goods layout, but the *same* run index uses the *same* goods seed in both regimes, so the two regimes form a *paired (blocked)* design, which we exploit rather than ignore. Reference coverage is therefore compared primarily with a paired Wilcoxon signed-rank test [57] and the matched-pairs effect size  $d_z$ ; binary mission completion with McNemar’s test on the discordant pairs [58] and with a logistic regression of completion on regime and all four factors. For completeness and comparability we also report the unpaired Mann–Whitney  $U$  test [59], the Welch  $t$ -test [60] and Cohen’s  $d$  [61]. Factor importance is assessed with a type-II factorial analysis of variance on reference coverage, using the saturated model coverage  $\sim A \times B \times C \times D$  with partial  $\eta^2$  as the effect-size measure [56]; because coverage is a bounded percentage, the ANOVA normality assumption holds only approximately, which is why the headline comparison rests on the distribution-free paired tests rather than on the parametric model alone.


### 3.6. Additional Robustness Analyses

Two targeted robustness analyses were added in response to concerns about the assumed planar-readability angle and the unequal identifier counts of the original deployment schemes. They were designed to isolate these two issues *separately* rather than to introduce a fully crossed quantity  $\times$  angle experiment. The original factorial remains the primary experiment; these are targeted controls that reuse the same simulation framework, warehouse, detection model, navigation rules, and predefined goods-layout seeds. Runs are paired by route-generation seed, with the realized trajectories allowed to diverge as a consequence of regime-specific recovery events.

**Angular-threshold sensitivity analysis.** This sweep was designed as an *aggregate* sensitivity analysis rather than a second high-precision factorial. The reference  $\pm 20^\circ$  planar-readability condition was supplemented by levels of  $\pm 10^\circ$ ,  $\pm 30^\circ$ ,  $\pm 45^\circ$  and  $\pm 60^\circ$ ; for each threshold the original 54 configurations were retained and *only* the incidence-angle condition of the planar regime was changed, while the spatial regime is angle-independent by construction and serves as the fixed reference. Each configuration used  $N = 10$  paired replications, and because every threshold was evaluated over all 54 configurations, this yields 540 runs per threshold. The same predefined goods-layout seeds were used at every threshold, preserving the paired structure and ensuring that observed differences are not caused by different route instances. The analysis is therefore interpreted at the aggregate level—the direction and magnitude of the response across thresholds—rather than as configuration-specific inference, and it is not intended to establish a universal physical readability limit. Because repeating the entire factorial at each additional threshold substantially increases the computational burden,  $N = 10$  was used here as a targeted robustness check; the primary experiment retains  $N = 100$  and remains the basis for the headline conclusions.

**Equal-count deployment control.** An equal-count control was added to determine whether the regime difference persists when identifier quantity is held constant. It contains three deployment domains with exactly 40 identifiers each: *Walls-40* retains the original 40 perimeter positions; *Racks-40* selects 40 positions from the 288 eligible rack-face candidates; and *Combined-40* selects 40 positions from the complete union of the 40 wall and 288 rack candidates ( $|C_{\text{Combined}}| = 40 + 288 = 328$ ). *Combined-40* is therefore an equal-count combined-domain control and should *not* be confused with the 250-identifier Hybrid deployment of the primary experiment, whose domain-informed redundancy-removal rule it does not reproduce. For Racks-40 and Combined-40, positions were selected with the same deterministic farthest-point procedure: starting from the candidate nearest the layout centroid, the algorithm repeatedly added the candidate maximizing its minimum planar Euclidean distance to the already-selected set until exactly 40 positions were obtained, with ties resolved by the lowest candidate index. The procedure uses only candidate coordinates and deployment-domain membership; it does not use the identifier regime, route outcomes, reference coverage, mission completion, or any other response variable, and the selected positions are identical for the spatial and planar regimes—it is thus outcome-independent and does not favor either regime by construction. The selected indices are published for auditability, and the three schemes are shown in Figure 5. The control was run for both regimes at detection ranges  $\{2, 4, 8\}$  m, FOV  $90^\circ$  and recovery step 20 m, with  $N = 100$  paired replications, under the reference  $\pm 20^\circ$  planar condition; it isolates deployment domain from identifier quantity and is not intended to estimate a quantity  $\times$  angle interaction (that sensitivity is evaluated separately above).


Exp-2 - equal-count control: placement of 40 identifiers (square warehouse 100x100 m; deterministic farthest-point selection - max-min spacing)



**Figure 5.** Equal-count identifier-deployment schemes used in the control experiment. Each scheme contains exactly 40 identifiers: (a) Walls-40 (the original 40 perimeter positions), (b) Racks-40 (farthest-point selection from the 288 rack-face candidates) and (c) Combined-40 (farthest-point selection from the complete union of the 40 wall and 288 rack candidates, 328 in total). Arrows denote each reference's facing normal, i.e., the readable direction of the planar code. The reference glyphs mark deployment positions and are not drawn to physical scale.

**Layout-robustness analysis.** To determine whether the navigation-level effect of angular reference availability was specific to the original square warehouse, the complete factorial experiment was replicated in L-shaped and U-shaped aisle-based layouts (Figure 6). The three environments retain comparable floor area, rack density (48 rack blocks of  $28 \times 2$  m with 4 m aisles), identifier-deployment rules, navigation parameters and experimental factors, while differing in their outer boundary and route geometry. The same predefined goods-layout seeds were used, and for each layout both the spatial and

the planar regime were evaluated with the same 54-configuration factorial and 100 paired replications, giving 10,800 runs per layout and 32,400 runs across the three layouts. The stability of the regime effect across layouts is assessed with a separate two-way model, coverage  $\sim$  regime  $\times$  layout, fitted to the run-level observations of the three environments with regime and layout as categorical factors; partial  $\eta^2$  compares the regime main effect with the regime  $\times$  layout interaction. This model is distinct from the primary  $A \times B \times C \times D$  ANOVA. Within each layout, the paired regime contrast is additionally summarized by a matched-pairs effect size based on identical seed indices. The analysis quantifies the stability of the availability effect across the three tested aisle-based layouts and does not claim independence from warehouse shape in general.



**Figure 6.** Warehouse layouts used in the layout-robustness analysis: (a) square, (b) L-shaped and (c) U-shaped environment. The layouts preserve comparable scale, rack density (48 rack blocks of  $28 \times 2$  m, 4 m aisles) and aisle-based organization while changing the outer boundary and route geometry; green and red squares mark the input and output stations.

**Additional operational metrics.** Because reference coverage has a denominator that depends on the number of deployed references, the robustness analyses additionally report metrics normalized by the travelled route and by time. These are obtained from per-run event records and accumulated visibility summaries rather than from a stored per-frame trace. Let  $L$  denote the total travelled distance of a run and  $T$  its duration, and call a reference *available* at an instant when it satisfies all regime-specific detection conditions: within range, within the field of view, within the incidence cone (planar regime only), in line of sight, and continuously so for at least the 0.25 s persistence window.

- *Distinct identifiers detected per 100 m:*  $\rho_{100} = 100 N_{\text{distinct}} / L$ , where  $N_{\text{distinct}}$  is the number of distinct identifiers detected at least once during the run. This is a distinct-identifier density, not a count of detection events.
- *Proportion of time with an available reference:*  $p_{\text{vis}} = t_{\text{vis}} / T$ , where  $t_{\text{vis}}$  is the accumulated time during which at least one identifier is available.
- *Continuous reference-free distance:* let  $\{g_k\}$  be the lengths of the maximal continuous travelled-distance segments during which no reference is available; for each run we compute the maximum  $g_{\text{max}} = \max_k g_k$  and the 95th percentile  $g_{95} = Q_{0.95}(\{g_k\})$  of that run's segments. The mean, confidence interval, median, and interquartile range of these run-level statistics across all runs are then reported.
- *Time and distance to relocalization:* for each loss-of-orientation event  $i$ , from entry into the lost state until the vehicle re-localizes, let  $\tau_i$  be its duration and  $\delta_i$  the distance travelled

during it; we report the per-run means  $\bar{\tau}$  and  $\bar{\delta}$ . A loss interval that remains unresolved and ends in mission abort is included using the elapsed time and distance recorded up to the abort point; these are lower-bound observations, because the recovery interval would have continued in the absence of the abort threshold. The fraction of runs that abort is the complement of the mission-completion rate.


- *Recovery nudges per run*: the total number of nudge (repositioning) events during a run. Each nudge follows a full 360° scan that failed to re-acquire a reference; scans and nudges are not logged separately, so only the nudge count is reported.

The primary experiment retained 100 paired replications per configuration. The angular sweep and the operational-metric analysis used fewer paired replications (10 and 50 per configuration, respectively) because they are targeted robustness checks whose interpretation focuses on effect direction, magnitude, and confidence intervals rather than on replacing the higher-precision primary experiment; the equal-count control retained the full 100 replications, as it is the direct rebuttal to the marker-count confound.

## 4. System Architecture and the Geometry of Readability

### 4.1. Software Architecture

The framework is organized into layers that communicate through a strictly one-way data flow, shown at the component level in Figure 7. The Unity simulation produces JSON records, which the Python layer turns into plots and statistical summaries; this one-directional interface makes it easy to test or replace individual layers, for example to swap the geometric detector for a real Vuforia integration when moving to physical hardware. A PowerShell orchestration layer launches the runs in headless batch mode. The navigation agent is an idealized kinematic model (holonomic motion on a 2D NavMesh), which is sufficient because the studied variable, the angular detectability of references along the trajectory, is largely independent of the vehicle's motion dynamics.



**Figure 7.** Component-level architecture and data flow of the simulation framework. The PowerShell orchestration layer launches headless runs; the Unity simulation core (C#) groups the configuration/scene,

navigation, perception, and instrumentation components; each run writes a per-run JSON record that the Python analysis layer turns into statistics and figures. The dashed block is the prepared ROS2 bridge to physical AGV hardware, which is not used in the present simulation study.

Within the Unity simulation core, the `ExperimentController` fixes the factorial configuration and the reference deployment, and the `RandomizationController` seeds the goods layout. The `NavigationController` and the `SimulatedIdentifierDetector` then interact every frame: the detector reports geometric visibility (range, field of view, incidence, line of sight) and the controller drives, scans and recovers, while the `DetectionHandler`, `LostPositionHandler` and `MetricsCollector` track the found/lost state and record the metrics. Each run emits a single JSON record that the Python layer aggregates into the statistics and figures. A prepared ROS2 bridge (`CommandParser`, `MessageBuilder`) is the intended path to physical AGV hardware and is not exercised in this study.

#### 4.2. Detection and Recovery Logic

The geometric detector evaluates, every frame, each marker within the detection range  $D$  of the camera (a uniform spatial grid restricts this to the local neighborhood). A marker is counted as *detected* when it is simultaneously (i) within range  $D$ ; (ii) inside the horizontal field of view, i.e., the angle between the heading and the direction to the marker is  $\leq B/2$ ; (iii) available in the spatial regime always, or in the planar regime only if the incidence angle between the reference normal and the viewing direction is  $\leq 20^\circ$ ; (iv) in line of sight—a clear ray reaches at least one point of a  $10 \times 10$  grid sampled over the marker face, so partially occluded markers still count—and (v) it has satisfied (i)–(iv) for at least a short persistence window (0.25 s). In the primary factorial experiment, the planar incidence-angle threshold in condition (iii) is fixed at the reference value of  $20^\circ$ ; only in the supplementary angular-sensitivity analysis (Section 5.4) is this constant replaced by the tested values of  $10^\circ$ ,  $20^\circ$ ,  $30^\circ$ ,  $45^\circ$  and  $60^\circ$ . Physical interpretation of the modelled parameters can be seen in Table 4.

The agent follows a NavMesh path that visits the 15 goods—in a seeded random order, not an optimized route—and ends at the output station; the planner returns the shortest collision-free path between consecutive waypoints. Each reference has a unique identity and a known global pose in the warehouse map. When a reference becomes available, the model treats it as yielding the pose of the camera relative to that reference; composed with the reference's known global pose, this relative transformation updates the vehicle's global pose estimate—this update is the *relocalization* event, and it occurs at the exact instant a reference first satisfies all availability conditions (during driving or during a scan). In this idealized geometric model the relative pose is assumed error-free and identical in both regimes, so the study models the *availability* of a relocalization reference, not the accuracy of the resulting pose; between observations the vehicle relies on drift-prone relative motion, and the periodic orientation check models the point at which accumulated drift, in the prolonged absence of any available reference, forces an explicit relocalization attempt. Orientation is re-checked *periodically*: every 100 m of travel the agent rotates to a reference heading and performs a  $360^\circ$  scan at  $90^\circ/\text{s}$ , evaluating detection every frame. If the scan finds a marker, the agent resumes immediately; only if a full sweep finds nothing does it enter a genuine *lost* state and perform a *nudge*—a displacement of  $C$  (20 or 30 m, factor  $C$ ) toward the first of eight candidate compass directions that lands on the NavMesh—after which it rescans. The nudge physically moves the agent, so a lost state affects both the trajectory and the metrics. After 100 failed recovery sweeps in total over the run, the run is *aborted*; reaching the output station marks it *completed*.

The two thresholds are deliberate, conservative choices applied identically to both regimes, so they do not bias the spatial–planar comparison. The 100 m orientation-check

interval lets the vehicle cross short marker gaps without an expensive stop while treating a prolonged absence of any visible marker across several aisle segments as a loss of orientation; a dedicated sensitivity analysis of this interval is outside the present scope and is left for future deployment-specific calibration. The two recovery-step levels (20 and 30 m) represent coarse relocation maneuvers at the scale of aisle spacing rather than fine local corrections, and test whether recovery behavior materially changes coverage and completion once visual reference is lost. The full logic is given as pseudocode in Figure 8.

```

1  # DETECTION (per frame, per marker m within range D of camera c)
2  detected(m) := dist(c,m) ≤ D
3      and horiz_angle(c.heading, c→m) ≤ B/2
4      and (regime ≡ ALL_AZIMUTH or incidence(m.normal, m→c) ≤ 20 deg)
5      and lineOfSight(c, m)           # clear ray to ≥ 1 of 10x10 pts on m
6      and visible_for(m) ≥ 0.25 s
7
8  # NAVIGATION + RECOVERY
9  goods ← 15 targets, seeded random order      # not TSP-optimized
10 path ← NavMesh shortest path: Input → goods → Output
11 fails ← 0; d_travel ← 0
12 for each step along path:
13     drive(step); d_travel ← d_travel + step_length
14     if d_travel ≥ 100 m:                       # periodic orientation check
15         d_travel ← 0
16         rotate to reference heading
17         repeat:
18             360 deg scan at 90 deg/s, checking detected(*) each frame
19             if any detected(m): break          # (re)oriented → resume
20             fails ← fails + 1
21             if fails ≥ 100: return ABORTED      # total over the run
22             nudge by C (20/30 m) to first of 8 dirs valid on NavMesh
23 return COMPLETED                             # Output reached

```

**Figure 8.** Pseudocode of the geometric detector (**top**) and the navigation/recovery loop (**bottom**). The single condition `regime == ALL_AZIMUTH` is the only difference between the two compared systems. The listing shows the reference configuration with a 20° planar-readability half-angle; the angular-sensitivity analysis (Section 5.4) replaces this constant with the threshold of the respective condition.

**Table 4.** Physical interpretation of the modelled parameters. The geometric model captures the *availability* of a localization reference; it deliberately abstracts away image formation, decoding and dynamics, which would require rendered-image or physical validation.

| Parameter           | Physical Meaning                                       | Governed by                                          | Modelled as                                                    | Not Represented                            |
|---------------------|--------------------------------------------------------|------------------------------------------------------|----------------------------------------------------------------|--------------------------------------------|
| Detection range $D$ | Max distance at which a reference is recognizable      | Reference size, camera resolution, optics, algorithm | Hard cutoff, $D \in \{2, 4, 8\}$ m                             | Gradual range falloff, partial recognition |
| Field of view $B$   | Angular width of the camera view                       | Lens and sensor                                      | Hard horizontal cone $\pm B/2$ , $B \in \{60, 90, 120\}^\circ$ | Vertical FOV, lens distortion              |
| Planar half-angle   | Max incidence angle at which a planar code is readable | Code type, print quality, resolution, lighting       | Hard cone $\pm 20^\circ$ (planar only; swept 10–60°)           | Smooth probabilistic decodability          |
| Persistence         | Minimum dwell before a reference counts                | Frame rate, detector latency, motion blur            | 0.25 s continuous satisfaction                                 | Filtering, detector-specific latency       |
| Line of sight       | Unobstructed view of the reference                     | Occluders, scene geometry                            | Ray to $\geq 1$ of $10 \times 10$ face points                  | Soft/partial occlusion, transparency       |
| Reference size      | Physical size of the reference                         | Manufacturing, mounting space                        | Implicit, via detection range                                  | Explicit pixel footprint                   |
| Camera resolution   | Sensor pixel count                                     | Camera hardware                                      | Implicit, via detection range                                  | Explicit image formation                   |
| Recovery step $C$   | Displacement per recovery nudge                        | Control policy, drivetrain                           | $C \in \{20, 30\}$ m along first valid NavMesh direction       | Rigid-body dynamics, slip, fine control    |


#### 4.3. Angular Readability of Planar and Spatial References

At the heart of the comparison lies a simple geometric argument about *angular readability*. A single-sided planar reference has a single privileged direction—its surface normal—and is available only within a narrow cone about it; rotating the viewing direction away from the normal quickly removes availability, so the admissible set of azimuths is a small angular sector. An all-azimuth spatial reference, by contrast, presents a recognizable appearance from any azimuth, so every viewing direction around its vertical axis is admissible.

This difference is decisive in the geometry of a warehouse aisle. As the vehicle drives along an aisle, it sees the references on the rack faces under continuously changing, mostly oblique angles. For the planar regime, most of these viewing angles fall outside the readable cone, so the reference is effectively unavailable for most of the pass; for the spatial regime, every angle is admissible. The factorial experiment of Section 5 quantifies how large this availability advantage is and—through the deployment  $\times$  range interaction—confirms that it is concentrated precisely where the planar reference's narrow readable cone hurts most.

Figure 9 makes this availability geometry explicit. In the planar regime, a reference is available only inside a readability cone about its surface normal (here  $\pm 20^\circ$ ), so a camera viewing it at oblique incidence obtains nothing; in the spatial regime, the reference is available from any azimuth. Both regimes additionally require the reference to lie within the detection range, within the camera field of view, and in line of sight.

Availability of a rack-face reference as the vehicle drives past, within detection range  $D$  (both regimes also require field of view and line of sight)



**Figure 9.** Availability of a rack-face reference as the vehicle drives past it within the detection range  $D$  (top-down view of one rack segment and the adjacent aisle). (a) A single-sided planar reference is available only where the viewing direction lies within the  $\pm 20^\circ$  readability band about its surface normal (green); over the rest of the pass it is within range but seen too obliquely to be read (red). (b) An all-azimuth spatial reference is available along the whole pass. Both regimes additionally require the reference to lie within the camera field of view and in line of sight.

## 5. Results


### 5.1. Primary Comparison Under the Reference $\pm 20^\circ$ Condition

Aggregated over all 5400 runs of each regime, the spatial regime achieves higher values than the planar regime on every headline metric (Table 5, Figure 10):  $5.7\times$  higher reference coverage (41.6% vs. 7.3%), a 30 percentage-point higher mission-completion rate (86.2% vs. 55.9%) and a  $6.1\times$  higher number of distinct references seen per run. The difference in coverage is statistically highly significant. Exploiting the paired design (Section 3), a paired Wilcoxon signed-rank test on the 5400 matched runs gives  $p < 0.001$  with a

matched-pairs effect size  $d_z = 1.84$ , and the spatial reference wins in 93% of pairs; the unpaired Mann–Whitney  $U = 25,777,783$  and Welch  $t = 102.5$  ( $p < 0.001$ , Cohen’s  $d = 1.97$ ) agree. For mission completion, McNemar’s test on the discordant pairs is decisive—1640 runs completed only under the spatial reference versus only 2 under the planar regime ( $p < 0.001$ )—and a logistic regression of completion on regime and all four factors yields an odds ratio of 790 (95% CI 519–1203,  $p < 0.001$ ) in favor of the spatial reference.

**Table 5.** Headline metrics: spatial reference vs. planar reference ( $n = 100$ , 5400 runs per regime). Means with 95% confidence intervals (reference coverage and distinct count: normal-approximation half-width; mission completion: Wilson interval).

| Metric                   | Spatial Reference | Planar Reference  | Effect                       | Statistical Evidence                                              |
|--------------------------|-------------------|-------------------|------------------------------|-------------------------------------------------------------------|
| Reference coverage [%]   | 41.6 $\pm$ 0.6    | 7.3 $\pm$ 0.2     | 5.7 $\times$ ; $d_z = 1.84$  | paired Wilcoxon $p < 0.001$ ; spatial regime wins in 93% of pairs |
| Mission completion [%]   | 86.2 [85.3, 87.1] | 55.9 [54.5, 57.2] | +30.3 p.p.; OR $\approx$ 790 | McNemar $p < 0.001$ ; discordant 1640 vs. 2                       |
| Distinct references seen | 80.8 $\pm$ 1.6    | 13.2 $\pm$ 0.3    | 6.1 $\times$                 | descriptive comparison                                            |



**Figure 10.** Overall comparison of the spatial reference and the planar reference on the three headline metrics, aggregated over 5400 runs per regime.

The completion model is a logistic regression of the binary mission outcome on the regime and the four deployment factors, all treated as *categorical* predictors with the planar regime and the reference levels (deployment scheme = Hybrid, FOV = 60°, recovery step = 20 m, range = 2 m) as baseline; no interaction terms are included in this headline model. The regime odds ratio of 790 is large but well bounded, and a cluster-robust variant (errors clustered by the 54 configurations) leaves it strongly significant with a wider interval (95% CI  $\approx$  [145, 4308],  $p < 0.001$ ). Some factor-level odds ratios are *not* finitely identified because completion saturates—reaching 100% for rack-based deployment schemes in the spatial regime and falling toward 0% for perimeter-only deployment at the shortest range (quasi-separation); we therefore report the regime contrast (which remains finite and stable) as the headline and characterize the deployment factors through the ANOVA on coverage (Section 5.3) and the descriptive completion rates (Table 6). Hypotheses **H1** and **H2** are confirmed.

**Table 6.** Logistic regression of mission completion, model  $\text{completion} \sim \text{regime} + A + B + C + D$  (all predictors categorical; baseline = planar regime, deployment scheme = Hybrid, FOV = 60°, recovery step = 20 m, range = 2 m; no interactions). Odds ratios (OR) with 95% CIs. Deployment scheme (A) and detection range (D) are not finitely identified because completion saturates (quasi-separation), so their levels are summarized qualitatively.

| Predictor (vs. Baseline)   | OR   | 95% CI       | Note                                                                                                 |
|----------------------------|------|--------------|------------------------------------------------------------------------------------------------------|
| Regime: spatial vs. planar | 790  | [519, 1203]  | the spatial regime strongly increases completion odds ( $p < 0.001$ ); cluster-robust CI [145, 4308] |
| FOV 90° vs. 60°            | 1.9  | [1.5, 2.5]   | wider FOV raises completion                                                                          |
| FOV 120° vs. 60°           | 9.4  | [6.9, 12.8]  | wider FOV raises completion strongly                                                                 |
| Recovery step 30 vs. 20 m  | 0.46 | [0.37, 0.57] | larger nudge slightly lowers completion                                                              |
| Deployment scheme (A)      | —    | —            | not finitely identified: completion saturates by deployment scheme                                   |
| Detection range D          | —    | —            | not finitely identified: completion $\rightarrow 99\%$ at $D = 8$ m                                  |


### 5.2. Deployment $\times$ Range Interaction (H4)

The aggregate hides a strongly non-uniform pattern (Table 7, Figure 11). In rack aisles at short and medium range (deployment = racks,  $D = 2$  and 4 m), the planar reference reaches only 0–8% coverage, while the spatial reference reaches 20–48%. Reading the  $\Delta = \text{spatial} - \text{planar}$  columns of Table 7 by deployment scheme, the gap is largest for *racks* at every range (the planar reference never exceeds 9%, so  $\Delta$  grows to +50) and smaller for *walls*, where the planar reference recovers to its highest value anywhere (24% at  $D = 8$  m); *Hybrid* behaves like racks with marginally higher planar-reference availability. Formally, this is an interaction between deployment scheme and range, whose strength the analysis of variance below quantifies, confirming H4. The geometric reason for the pattern is examined in Section 6.

**Table 7.** Reference coverage [%] by deployment scheme (A) and detection range (D), with the gap  $\Delta = \text{spatial} - \text{planar}$  in percentage points. The spatial-regime advantage is largest in rack aisles, where the planar reference is viewed edge-on.

| Deployment Scheme | $D = 2$ m |     |          | $D = 4$ m |     |          | $D = 8$ m |     |          |
|-------------------|-----------|-----|----------|-----------|-----|----------|-----------|-----|----------|
|                   | Sp.       | Pl. | $\Delta$ | Sp.       | Pl. | $\Delta$ | Sp.       | Pl. | $\Delta$ |
| Walls             | 2         | 0   | +2       | 50        | 1   | +49      | 70        | 24  | +46      |
| Racks             | 20        | 0   | +20      | 48        | 8   | +40      | 59        | 9   | +50      |
| Hybrid            | 19        | 2   | +17      | 46        | 10  | +36      | 60        | 11  | +49      |

**Reference coverage by deployment scheme and detection range**




**Figure 11.** Reference coverage by deployment scheme and detection range for both regimes. The largest spatial-regime advantage occurs in rack aisles at short and medium range, where the planar reference drops to near zero.

### 5.3. Factorial Effects in the Primary Experiment (H3)

A type-II factorial ANOVA of the spatial regime (Table 8, Figure 12) shows that **detection range D dominates**, with partial  $\eta^2 = 0.903$ —it explains the overwhelming majority of the variability in coverage. The next strongest effects are the deployment  $\times$  range interaction ( $\eta^2 = 0.390$ ) and the field of view ( $\eta^2 = 0.348$ ); the deployment scheme on its own has a small main effect ( $\eta^2 = 0.009$ ), because a spatial reference is detectable from any angle and thus depends mainly on whether it is within range at all. The model explains  $R^2 = 0.916$  of the variance. The same analysis for the planar regime is qualitatively identical, with range again dominant ( $\eta^2 = 0.880$ ) and an even stronger deployment  $\times$  range interaction ( $\eta^2 = 0.810$ ), reflecting the planar reference's greater geometric sensitivity. **H3** is confirmed for both regimes.


**Table 8.** Effect sizes (partial  $\eta^2$ ) from the type-II factorial ANOVA of reference coverage, for both detection regimes (saturated model  $A \times B \times C \times D$ ; all listed effects  $p < 0.001$  except spatial-regime factor C,  $p = 0.019$ ). Model  $R^2 = 0.916$  (spatial), 0.925 (planar).

| Effect                | Spatial $\eta^2$ | Planar $\eta^2$ | Interpretation                                                    |
|-----------------------|------------------|-----------------|-------------------------------------------------------------------|
| D (detection range)   | 0.903            | 0.880           | dominant factor in both regimes                                   |
| A $\times$ D          | 0.390            | 0.810           | much stronger deployment-scheme sensitivity for the planar regime |
| B (field of view)     | 0.348            | 0.123           | secondary visibility factor                                       |
| B $\times$ D          | 0.137            | 0.007           | range and FOV jointly affect spatial-regime coverage              |
| A (deployment scheme) | 0.009            | 0.208           | small for the spatial regime, larger for the planar regime        |
| C (recovery step)     | 0.001            | 0.024           | negligible influence on coverage                                  |



**Figure 12.** Factor importance (partial  $\eta^2$ ) for reference coverage. Detection range is by far the most influential factor, followed by the deployment  $\times$  range interaction and the field of view.


The main effects (Figure 13) make the practical message concrete: coverage rises from 13.7% at  $D = 2$  m to 62.9% at  $D = 8$  m, a far steeper dependence than on any other factor. The same dominance holds for mission completion, which in the spatial regime rises from 66% at  $D = 2$  m to 93% at 4 m and 99% at 8 m; the logistic regression of Section 5.1 likewise identifies detection range as the strongest completion driver among the four factors.



**Figure 13.** Main effects of the four factors on reference coverage with 95% confidence intervals (spatial regime). (A) Deployment scheme, (B) field of view, (C) recovery step, (D) detection range. The steep rise with detection range dominates.

#### 5.4. Sensitivity to the Planar-Readability Threshold

The 20° readability half-angle of the planar reference is a modelling choice, so we repeated the full 54-configuration factorial in the planar regime at five half-angles—10°, 20°, 30°, 45° and 60° ( $N = 10$  per configuration, 540 runs each)—and compared them against the angle-independent spatial reference. Coverage of the planar reference rises monotonically with the readability cone, from 4.9% at 10° to 21.1% at 60°, whereas the spatial reference is by construction insensitive to the angle at  $\approx 40\%$  (Table 9, Figure 14). The spatial-over-planar coverage ratio consequently narrows from  $8.1\times$  to  $1.9\times$ , but the spatial reference retains an advantage across the entire tested range of reading angles; at the reference 20° the ratio is  $5.8\times$ , matching the headline  $5.7\times$ . Within the range of thresholds examined here, the reported advantage is therefore not an artifact of the specific 20° choice.



**Figure 14.** Planar-reference coverage as a function of the planar-readability half-angle (planar regime,  $N = 10$  per configuration), against the angle-independent spatial reference. The aggregate results retain a spatial-regime advantage across all tested model thresholds.


**Table 9.** Sensitivity of planar-reference coverage to the planar-readability half-angle. The planar sweep spans the full 54-configuration factorial at  $N = 10$  per configuration (540 runs per angle); the angle-independent spatial reference (coverage 40.0%, completion 86.8%) is an independent descriptive baseline computed from the full-factorial spatial-regime runs of the operational subset ( $N = 50$ ); being angle-independent by construction it is shown as a constant reference line, not as a seed-matched pairing with the planar  $N = 10$  sweep. Reference values differ slightly from the headline  $N = 100$  figures (planar 6.9 vs. 7.3%; spatial 40.0 vs. 41.6%).

| Planar Half-Angle | Planar Coverage [%] (95% CI) | Planar Completion [%] (Wilson 95% CI) | Spatial/Planar Ratio |
|-------------------|------------------------------|---------------------------------------|----------------------|
| 10°               | 4.9 [4.5, 5.3]               | 42.0 [37.9, 46.2]                     | 8.1×                 |
| 20° (reference)   | 6.9 [6.2, 7.6]               | 52.4 [48.2, 56.6]                     | 5.8×                 |
| 30°               | 7.8 [7.1, 8.5]               | 53.9 [49.7, 58.1]                     | 5.1×                 |
| 45°               | 11.1 [10.3, 11.9]            | 60.6 [56.4, 64.6]                     | 3.6×                 |
| 60°               | 21.1 [19.5, 22.8]            | 69.8 [65.8, 73.5]                     | 1.9×                 |
| Spatial reference | 40.0 [39.1, 40.9]            | 86.8 [85.5, 88.0]                     | —                    |

### 5.5. Equal-Count Deployment Control

Because the deployment schemes of the main experiment differ in reference count (Walls = 40, Racks = 288, Hybrid = 250), one might attribute part of the spatial-regime advantage to the larger number of references rather than to their all-azimuth availability. To evaluate whether the advantage persists at a constant identifier count, we constructed three equal-count schemes with *exactly* 40 references each: *Walls-40* (the original perimeter set), *Racks-40* (40 positions from the 288 rack-face candidates) and *Combined-40* (40 positions from the complete union of the wall and rack candidates); the deterministic, outcome-independent selection is described in Section 3.6. Each scheme was run for both regimes at detection ranges {2, 4, 8} m, FOV 90° and recovery step 20 m under the reference  $\pm 20^\circ$  condition ( $N = 100$  per cell). Pooled over range, the spatial reference retains a clear advantage in all three schemes—coverage 37–43% versus 7–11% and mission completion 64–93% versus 0.7–33% (Table 10, Figure 15). Resolved by detection range (Table 11), the range dominance of the primary experiment reappears: spatial-regime completion reaches 100% at the longer ranges, whereas the planar reference stays low except for perimeter references at  $D = 8$  m. Under this reference  $\pm 20^\circ$  condition, the advantage therefore *persists when the identifier count is held constant*, indicating that the difference observed in the primary experiment cannot be attributed solely to the unequal identifier counts of the original deployment schemes.

Equal reference count ( $n=40$ ): spatial vs planar



**Figure 15.** Equal-reference-count control ( $n = 40$  per scheme,  $N = 100$ ): (a) reference coverage and (b) mission completion for the spatial and the planar regime. The spatial regime achieves higher coverage and mission completion in every tested equal-count scheme.

**Table 10.** Equal-reference-count control: spatial vs. planar regime at exactly 40 references per scheme ( $N = 100$ ; FOV  $90^\circ$ , recovery step 20 m, ranges  $\{2, 4, 8\}$  m pooled).

| Scheme ( $n = 40$ ) | Spatial Cov. [%] | Planar Cov. [%] | Spatial Compl. [%] | Planar Compl. [%] |
|---------------------|------------------|-----------------|--------------------|-------------------|
| Walls-40            | 41.0             | 9.6             | 64.0               | 33.3              |
| Racks-40            | 42.9             | 10.8            | 92.7               | 10.7              |
| Combined-40         | 37.0             | 7.1             | 82.0               | 0.7               |

**Table 11.** Equal-count control resolved by detection range  $D$  (each cell  $N = 100$ ; FOV  $90^\circ$ , recovery step 20 m, reference  $\pm 20^\circ$ ). Completion rises steeply with range for the spatial reference, reproducing the range dominance of the primary experiment; the planar reference stays low except for perimeter references at  $D = 8$  m.

| Scheme      | $D$ [m] | Spatial Cov. [%] | Planar Cov. [%] | Spatial Compl. [%] | Planar Compl. [%] |
|-------------|---------|------------------|-----------------|--------------------|-------------------|
| Walls-40    | 2       | 3.6              | 0.0             | 0                  | 0                 |
| Walls-40    | 4       | 50.2             | 1.6             | 92                 | 0                 |
| Walls-40    | 8       | 69.3             | 27.1            | 100                | 100               |
| Racks-40    | 2       | 28.2             | 0.2             | 78                 | 0                 |
| Racks-40    | 4       | 45.0             | 15.5            | 100                | 8                 |
| Racks-40    | 8       | 55.5             | 16.7            | 100                | 24                |
| Combined-40 | 2       | 23.1             | 0.1             | 46                 | 0                 |
| Combined-40 | 4       | 39.1             | 9.2             | 100                | 0                 |
| Combined-40 | 8       | 49.0             | 12.0            | 100                | 2                 |

### 5.6. Operational Visibility and Relocalization Metrics


Reference coverage has a denominator that depends on the number of deployed references, so we additionally report metrics normalized by the travelled route and by time, computed from the per-run event records and visibility summaries over the full factorial (54 configurations per regime,  $N = 50$ ). The spatial regime keeps at least one reference available for 22% of the run against 2% for the planar regime, sees 7.2 distinct references per 100 m against 0.8, and shortens the longest reference-free stretch roughly threefold (359 vs. 1096 m); after a loss of orientation it re-localizes about four times faster (45 vs. 162 s) and over roughly a quarter of the distance, and needs about half as many recovery nudges (Table 12). Because the planar regime aborts more often, its relocalization time and distance include more unresolved loss intervals recorded only up to the abort point rather than to a genuine recovery (Section 3.6). Consequently, the reported planar-regime relocalization times and distances are conservative lower-bound estimates for unresolved cases, so the spatial advantage is, if anything, understated. Because these distributions are strongly right-skewed, the table reports both the mean with its 95% confidence interval and the median with the interquartile range; a paired comparison across the 54 configurations (per-configuration means, Wilcoxon signed-rank) is significant for every metric ( $p < 0.001$ ). These route- and time-normalized figures do not depend on the reference-count denominator and confirm the operational relevance of the advantage; the 95th-percentile reference-free distance is reported as a descriptive supporting statistic. In this  $N = 50$  operational subset, the planar regime's median realized route length is about 2200 m, roughly twice its longest reference-free stretch ( $\sim 1096$  m) and well above the primary-experiment median, because planar-regime runs wander substantially before re-localizing, consistent with the  $\approx 7\%$  planar coverage.

**Table 12.** Operational metrics normalized by route and time (full factorial, 54 configurations  $\times N = 50 = 2700$  runs per regime). For each metric, the run-level mean with its 95% confidence interval and, because the distributions are right-skewed, the median with the interquartile range are given. All metrics favor the spatial reference; the paired comparison across the 54 configurations is significant for every metric (Wilcoxon signed-rank,  $p < 0.001$ ).

| Metric                          | Spatial Regime    |                  | Planar Regime     |                  |
|---------------------------------|-------------------|------------------|-------------------|------------------|
|                                 | Mean [95% CI]     | Median (IQR)     | Mean [95% CI]     | Median (IQR)     |
| Distinct identifiers per 100 m  | 7.2 [6.9, 7.4]    | 4.8 (1.0–13.5)   | 0.8 [0.8, 0.8]    | 0.4 (0.0–1.6)    |
| Proportion of time available    | 0.22 [0.21, 0.23] | 0.08 (0.02–0.39) | 0.02 [0.02, 0.02] | 0.01 (0.00–0.03) |
| Max reference-free distance [m] | 359 [335, 382]    | 101 (18–328)     | 1096 [1058, 1133] | 367 (201–2208)   |
| p95 reference-free distance [m] | 322 [298, 346]    | 52 (12–241)      | 1077 [1039, 1115] | 316 (161–2208)   |
| Time to relocalization [s]      | 45 [41, 49]       | 6 (3–22)         | 162 [156, 169]    | 25 (12–402)      |
| Distance to relocalization [m]  | 224 [201, 246]    | 11 (2–90)        | 864 [827, 902]    | 105 (38–2098)    |
| Recovery nudges per run         | 27 [26, 28]       | 5 (1–47)         | 60 [58, 61]       | 64 (18–99)       |

### 5.7. Spatial Blind-Belt Analysis


A heatmap of the locations where the vehicle loses orientation (Figure 16) reveals *blind belts*—zones in which no reference is within range and line of sight, corresponding to places where the spacing between adjacent references exceeds twice the detection range. These zones are most pronounced at the shortest range and are the natural target of the deployment method discussed below: they show exactly where references must be added, or the range increased.



**Figure 16.** Heatmap of loss-of-orientation (blind-belt) events in the square warehouse; the warmer the cell, the more losses. The count  $n$  shown in the figure is the total number of recorded loss events aggregated over all shortest-range ( $D = 2$  m) runs of the spatial regime—not the number of simulation runs—the shortest range being shown because blind belts are most pronounced there. These zones indicate where markers should be added, or the detection range increased.

To make the navigation behavior concrete, Figure 17 shows, for each deployment scheme, the best (highest-coverage, completed) spatial-regime run—the actual path the robot drove from the input station, through the 15 goods, to the output station, in the

square warehouse. In all three deployment schemes, the spatial regime keeps the vehicle oriented along the full aisle traversal: the path is continuous, and loss events (red crosses) are rare, occurring mainly at short detection range inside the blind belts identified above. The contrast with the planar reference is implicit in the coverage gap of Section 5.2: with a planar reference, the same aisle passes would leave the rack-facing markers unread for most of their length.



**Figure 17.** Best robot routes for the three deployment schemes (Walls, Racks, Hybrid) with the spatial reference—the highest-coverage completed run in each case. Blue dots are identifiers, gold stars are goods to collect, the navy line is the actual robot path, red crosses are loss-of-orientation events; the green and red squares are the input and output stations.

### 5.8. Robustness of the Angular-Availability Effect Across Warehouse Layouts

Replicating the full factorial in the L-shaped and U-shaped layouts (Section 3.6, Figure 6) shows that the availability advantage of the spatial regime is stable across all three tested aisle-based layouts (Table 13). The spatial regime keeps coverage at 41.6/41.4/35.8% (square/L/U) against 7.3/7.9/6.2% for the planar regime—a coverage ratio of  $5.7 \times / 5.3 \times / 5.7 \times$ —and wins in 93/97/99% of the matched pairs, with a large paired effect size ( $d_z = 1.84/1.91/2.19$ ). Mission completion is +30/+32/+65 percentage points higher under the spatial regime, the larger U-shape gap reflecting the planar regime's low completion (19.2%) in that more constrained route geometry. A two-way analysis of variance of coverage on regime and layout attributes the overwhelming share of variance to the regime (partial  $\eta^2 = 0.523$ ) and only a negligible amount to the regime  $\times$  layout interaction ( $\eta^2 = 0.004$ ); that is, the advantage is essentially the same in all three layouts. Detection range remains the dominant deployment factor within each layout (partial  $\eta^2(D) = 0.90/0.87/0.92$ ). These results indicate *stability of the effect across the three tested aisle-based layouts* rather than general independence from warehouse shape.

**Table 13.** Layout-robustness: spatial vs. planar regime across the square, L-shaped and U-shaped warehouses (54 configurations  $\times$  100 paired replications per layout and regime; 5400 runs per layout per regime). Coverage means with 95% confidence intervals;  $d_z$  is the matched-pairs effect size; “spatial wins” is the fraction of matched pairs in which the spatial regime attains higher coverage.

| Layout   | Coverage [%]   |               | Ratio        | Completion [%] |        | $\Delta$ p.p. | $d_z$ | Spatial Wins [%] |
|----------|----------------|---------------|--------------|----------------|--------|---------------|-------|------------------|
|          | Spatial        | Planar        |              | Spatial        | Planar |               |       |                  |
| Square   | 41.6 $\pm$ 0.6 | 7.3 $\pm$ 0.2 | 5.7 $\times$ | 86.2           | 55.9   | +30.3         | 1.84  | 93               |
| L-shaped | 41.4 $\pm$ 0.6 | 7.9 $\pm$ 0.2 | 5.3 $\times$ | 86.6           | 55.2   | +31.5         | 1.91  | 97               |
| U-shaped | 35.8 $\pm$ 0.4 | 6.2 $\pm$ 0.2 | 5.7 $\times$ | 84.0           | 19.2   | +64.8         | 2.19  | 99               |

## 6. Discussion

The experiment quantifies a difference that is, in part, geometric by construction: the planar regime adds an incidence-angle constraint that the spatial regime does not, so it is expected *a priori* that a reference is available more often in the spatial regime. We therefore state the scope of the claim explicitly. This study (i) does not validate a specific physical spatial reference, (ii) does not compare real image-based recognition or decoding algorithms, and (iii) quantifies, *under the reference geometric model*, the navigation-level consequence of an assumed angular-availability difference. Its scientific value lies not in the direction of the effect, which follows from the model, but in its *magnitude and structure*: the availability advantage is a factor of about 5.7 in coverage and 30 percentage points in completion (large enough to affect operational usability), and it is concentrated, through the deployment  $\times$  range interaction, in the rack aisles that dominate a picking route. The mechanism is the aisle viewing geometry: rack-facing references are seen edge-on at every range, so in the planar regime availability stays near zero, whereas perimeter references are approached more frontally and the planar regime partially recovers at long range; the spatial regime, available from any azimuth, is indifferent to this distinction. The  $\pm 20^\circ$  readability half-angle is a modelling choice reflecting the physical property of planar codes [3]; because planar coverage decreases monotonically as the cone tightens,  $\pm 20^\circ$  is a moderate rather than an extreme assumption. The aggregate sensitivity analysis of Section 5.4, sweeping the modelled half-angle from  $10^\circ$  to  $60^\circ$ , indicates that the spatial advantage persists across all tested thresholds. Replacing the hard cutoff with a smooth probabilistic readability model is left for future work.

The second, more actionable, finding is the dominance of detection range (partial  $\eta^2 = 0.903$ ). For a spatial reference available from any angle, the deployment scheme matters little on its own; the binding constraint is whether a reference is within range. In deployment terms, the operator has freedom in *where* to place references but is constrained by the *range*, which is governed by reference size and camera resolution. The blind-belt heatmap turns this into a concrete procedure: run a pilot batch, locate the belts, and either increase reference density or detection range until the belts close. A blind belt is defined *locally*, as a region where the spacing between adjacent references exceeds twice the detection range, so that none is within range and line of sight; hence, within the present geometric model, the blind-belt criterion is governed primarily by local reference spacing, line of sight, and detection range rather than by the global outline of the warehouse. This locality is the basis for expecting the method to transfer across layouts, and it is exactly the pair of factors (density via deployment, and range) that the ANOVA identifies as dominant.

The layout-robustness analysis (Section 5.8) supports this locality argument empirically for the tested cases: the effect was not confined to the square warehouse but held in the L-shaped and U-shaped layouts, with a negligible regime  $\times$  layout interaction ( $\eta^2 = 0.004$ ) and detection range remaining dominant in each. The U-shaped layout showed the largest completion gap, because its more constrained route geometry penalizes the planar regime's low availability more heavily. These three aisle-based layouts do not, however, represent all warehouse types, so the finding is best read as stability across the tested layouts rather than as general shape-independence.

The comparison is deliberately framed at the level of *reference availability* and is technology-neutral. The all-azimuth regime is an idealization; in practice, azimuthal availability can be approached by a *multi-sided planar fiducial object* that distributes ArUco- or AprilTag-like markers over several faces [54]. The present results therefore quantify the navigation-level value of azimuthal availability independently of whether it is obtained from a shape-only spatial reference or from a sufficiently multi-sided planar object; a direct comparison of these two physical realizations is a natural next step.

The study has several limitations. First, the navigation agent is an idealized holonomic model without rigid-body dynamics; a real forklift would add inertia, a minimum turning radius, and slip. Second, and most importantly, the perception model is an idealized *geometric* availability model: it captures range, field of view, incidence angle, line of sight and persistence, but *not* image formation, decoding, or pose-estimation error. Several geometric assumptions are shared by the two regimes, which supports a controlled comparison; even so, both the absolute performance values and the magnitude of the regime difference remain conditional on the assumed availability rules. The absolute coverage and completion figures should therefore be read as model-based estimates under the reference geometric model rather than as empirically validated values, and their dependence on lighting, material appearance, exposure, shadows, motion blur, occlusion, and the specific detector is not represented and would require rendered-image or physical validation. Third, the model does not include a cost or maintenance model, a technology-specific deployment optimization, or a comparison of the computational cost of image-based planar decoding versus spatial-object recognition; the physical manufacture, mounting, durability, and soiling of a spatial reference are likewise outside its scope, and the range and field-of-view factors are only proxies for the underlying camera and marker hardware. Fourth, the experiment covers three aisle-based layouts (square, L, U) at a single scale; although the blind-belt criterion is local, empirical generalization to other shapes, scales, and open-floor layouts remains future work, so the transferability claim is a reasoned expectation rather than a demonstrated result. Finally, the results are specific to visual, line-of-sight references and do not transfer to non-line-of-sight technologies such as RFID or to feature-based visual SLAM, whose availability and error characteristics differ fundamentally. Compared with the closest related work [52], which shares the Unity + NavMesh foundation but analyses individual runs of a single reference type, the present study adds the factorial comparison of the two availability regimes, the robustness analyses and the geometric explanation of the result.

Lighting and image formation are further determinants of visual-reference availability that a purely geometric model does not capture. The present detector does not represent illumination intensity, exposure, shadows, reflections, material appearance, or motion blur. Empirical studies of visual localization and location-based AR under varying illumination confirm that lighting changes can alter feature detection, matching stability, and localization performance [62]. Satisfying the geometric range, field-of-view, incidence-angle, and line-of-sight conditions should therefore be read as availability under favorable imaging conditions, not as a guarantee of successful recognition by a physical camera system, and it is why we treat geometric availability as a necessary but not sufficient condition for physical recognition.

The geometric detector also says nothing about the cost of recognition. It does not measure recognition latency, frame rate, computational load or energy use, and image-based recognition of a spatial object may be more demanding than the detection of a planar fiducial. The navigation-level availability benefit reported here would therefore have to be weighed, in a physical system, against recognition latency, hardware requirements and real-time control constraints, as well as against the manufacture, mounting, calibration, soiling and physical damage of a spatial reference and the need to maintain a model registry.

## 7. Conclusions

This study presented a reproducible simulation-based method for evaluating the navigation-level consequences of all-azimuth spatial-reference availability compared with a single-sided, angularly constrained planar-reference regime. Under the reference geometric model, the spatial regime achieved  $5.7\times$  higher reference coverage and

a 30.3 percentage-point higher mission-completion rate. Detection range was the dominant deployment factor, while the deployment-by-range interaction showed that the largest benefit occurred in rack aisles, where planar references were frequently viewed at oblique or edge-on angles. The qualitative advantage remained present across the tested planar-readability thresholds, equal-count deployment controls, and Square, L-shaped, and U-shaped warehouse layouts.

These results quantify the potential navigation benefit of azimuthal reference availability under controlled geometric assumptions. They do not establish the physical superiority, recognition accuracy, or real-time deployment readiness of a particular spatial identifier relative to QR, ArUco, or AprilTag systems. Future work will therefore evaluate defined physical reference geometries using rendered-image and real-camera recognition, calibrate a probabilistic availability model under controlled imaging conditions, and validate the approach on physical AGV/AMR hardware.

**Author Contributions:** Conceptualization, K.K., O.H. and E.K.; methodology, K.K. and O.H.; software, K.K. and O.K.; validation, K.K., O.H. and E.K.; formal analysis, K.K.; investigation, K.K.; writing—original draft preparation, K.K.; writing—review and editing, O.H., E.K. and O.K.; supervision, O.H. and E.K. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research was funded by the Slovak Research and Development Agency under the contract no. APVV-23-0123, and by the Scientific Grant Agency of the Ministry of Education, Research, Development and Youth of the Slovak Republic No. 1/0637/23 and 1/0045/25.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable, as the study did not involve human subjects.

**Data Availability Statement:** The datasets used and analyzed during the current study are available from the corresponding author on reasonable request.

**Acknowledgments:** The authors acknowledge the use of AI tools, including Gemini 3.1, Claude Opus 4.7-4.8 and DeepL, for translation into English and grammatical corrections of the manuscript. These tools were employed to enhance the clarity and linguistic precision of the text while maintaining the authors' original intent and scientific rigor.

**Conflicts of Interest:** The authors declare no conflict of interest.

## References

1. Mathaba, T.N.D. Applications of Industry 4.0 Technologies in Warehouse Management: A Systematic Literature Review. *Logistics* **2023**, *7*, 24. <https://doi.org/10.3390/logistics7020024>.
2. van Geest, R.; Tekinerdogan, B.; Catal, C. Intelligent Warehouse in Industry 4.0—Systematic Literature Review. *Sensors* **2023**, *23*, 4105. <https://doi.org/10.3390/s23084105>.
3. ISO/IEC 18004:2015; Information Technology—Automatic Identification and Data Capture Techniques—QR Code Bar Code Symbol Specification. International Organization for Standardization: Geneva, Switzerland, 2015.
4. Romero-Ramirez, F.J.; Muñoz-Salinas, R.; Medina-Carnicer, R. Speeded up detection of squared fiducial markers. *Image Vis. Comput.* **2018**, *76*, 38–47. <https://doi.org/10.1016/j.imavis.2018.05.004>.
5. AprilRobotics. AprilTag: A Robust and Flexible Visual Fiducial System. 2024. Available online: <https://github.com/AprilRobotics/apriltag> (accessed on 26 July 2026).
6. Kagermann, H.; Wahlster, W.; Helbig, J. *Recommendations for Implementing the Strategic Initiative INDUSTRIE 4.0—Securing the Future of German Manufacturing Industry*; Technical Report; National Academy of Science and Engineering: München, Germany, 2013.
7. Roblek, V.; Meško, M.; Krapež, A. A Complex View of Industry 4.0. *SAGE Open* **2016**, *6*, 2158244016653987. <https://doi.org/10.1177/2158244016653987>.
8. Pereira, A.; Romero, F. A Review of the Meanings and the Implications of the Industry 4.0 Concept. *Procedia Manuf.* **2017**, *13*, 1206–1214. <https://doi.org/10.1016/j.promfg.2017.09.032>.
9. Zhong, R.Y.; Xu, X.; Klotz, E.; Newman, S.T. Intelligent Manufacturing in the Context of Industry 4.0: A Review. *Engineering* **2017**, *3*, 616–630. <https://doi.org/10.1016/J.ENG.2017.05.015>.

10. Bahrin, M.A.K.; Othman, M.F.; Nor Azli, N.H.; Talib, M.F. Industry 4.0: A Review on Industrial Automation and Robotic. *J. Teknol.* **2016**, *78*, 137–143.
11. Bortolini, M.; Ferrari, E.; Gamberi, M.; Pilati, F.; Faccio, M. Assembly System Design in the Industry 4.0 Era: A General Framework. *IFAC-PapersOnLine* **2017**, *50*, 5700–5705. <https://doi.org/10.1016/j.ifacol.2017.08.1121>.
12. Kamble, S.S.; Gunasekaran, A.; Gawankar, S.A. Sustainable Industry 4.0 Framework: A Systematic Literature Review Identifying the Current Trends and Future Perspectives. *Process Saf. Environ. Prot.* **2018**, *117*, 408–425. <https://doi.org/10.1016/j.psep.2018.05.009>.
13. Lee, J.; Bagheri, B.; Kao, H.A. A Cyber-Physical Systems Architecture for Industry 4.0-Based Manufacturing Systems. *Manuf. Lett.* **2015**, *3*, 18–23. <https://doi.org/10.1016/j.mfglet.2014.12.001>.
14. Posada, J.; Toro, C.; Barandiaran, I.N.; Oyarzun, D.; Stricker, D.; de Amicis, R.; Pinto, E.B.; Eisert, P.; Döllner, J.; Vallarino, I. Visual Computing as a Key Enabling Technology for Industrie 4.0 and Industrial Internet. *IEEE Comput. Graph. Appl.* **2015**, *35*, 26–40. <https://doi.org/10.1109/MCG.2015.45>.
15. Lucke, D.; Constantinescu, C.; Westkämper, E. Smart Factory—A Step towards the Next Generation of Manufacturing. In *Manufacturing Systems and Technologies for the New Frontier*; Springer: London, UK, 2008; pp. 115–118.
16. Kolberg, D.; Zühlke, D. Lean Automation Enabled by Industry 4.0 Technologies. *IFAC-PapersOnLine* **2015**, *48*, 1870–1875. <https://doi.org/10.1016/j.ifacol.2015.06.359>.
17. Reljić, V.; Milenković, I.; Dudić, S.; Šulc, J.; Bajči, B. Augmented Reality Applications in Industry 4.0 Environment. *Appl. Sci.* **2021**, *11*, 5592. <https://doi.org/10.3390/app11125592>.
18. Langfinger, M.; Schneider, M.; Stricker, D.; Schotten, H.D. Addressing Security Challenges in Industrial Augmented Reality Systems. In Proceedings of the IEEE 15th International Conference on Industrial Informatics (INDIN), Emden, Germany, 24–26 July 2017; pp. 299–304. <https://doi.org/10.1109/INDIN.2017.8104789>.
19. Longo, F.; Nicoletti, L.; Padovano, A. Smart Operators in Industry 4.0: A Human-Centered Approach to Enhance Operators' Capabilities and Competencies within the New Smart Factory Context. *Comput. Ind. Eng.* **2017**, *113*, 144–159. <https://doi.org/10.1016/j.cie.2017.09.016>.
20. Belkadi, F.; Dhuieb, M.A.; Aguado, J.V.; Laroche, F.; Bernard, A.; Chinesta, F. Intelligent Assistant System as a Context-Aware Decision-Making Support for the Workers of the Future. *Comput. Ind. Eng.* **2020**, *139*, 105732. <https://doi.org/10.1016/j.cie.2019.02.046>.
21. Van Lopik, K.; Sinclair, M.; Sharpe, R.; Conway, P.; West, A. Developing Augmented Reality Capabilities for Industry 4.0 Small Enterprises. *Comput. Ind.* **2020**, *117*, 103208. <https://doi.org/10.1016/j.compind.2020.103208>.
22. de Souza Cardoso, L.F.; Mariano, F.C.M.Q.; Zorzal, E.R. A Survey of Industrial Augmented Reality. *Comput. Ind. Eng.* **2020**, *139*, 106159. <https://doi.org/10.1016/j.cie.2019.106159>.
23. Milgram, P.; Kishino, F. A Taxonomy of Mixed Reality Visual Displays. *IEICE Trans. Inf. Syst.* **1994**, *E77-D*, 1321–1329.
24. Azuma, R.T. A Survey of Augmented Reality. *Presence Teleoperators Virtual Environ.* **1997**, *6*, 355–385.
25. Zhou, F.; Duh, H.B.L.; Billinghurst, M. Trends in Augmented Reality Tracking, Interaction and Display: A Review of Ten Years of ISMAR. In Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR), Cambridge, UK, 15–18 September 2008; pp. 193–202.
26. Billinghurst, M.; Clark, A.; Lee, G. A Survey of Augmented Reality. *Found. Trends Hum.-Interact.* **2015**, *8*, 73–272. <https://doi.org/10.1561/11000000049>.
27. Chatzopoulos, D.; Bermejo, C.; Huang, Z.; Hui, P. Mobile Augmented Reality Survey: From Where We Are to Where We Go. *IEEE Access* **2017**, *5*, 6917–6950. <https://doi.org/10.1109/ACCESS.2017.2698164>.
28. Furht, B. *Handbook of Augmented Reality*; Springer: Berlin/Heidelberg, Germany, 2011.
29. Pentenrieder, K.; Meier, P.; Wuest, H. Augmented Reality-Based Factory Planning—An Application Tailored to Industrial Requirements. In Proceedings of the 6th IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR), Nara, Japan, 13–16 November 2007; pp. 31–34.
30. Wang, Z.; Bai, X.; Zhang, S.; Wang, Y.; Han, S.; Zhang, X.; Yan, Y.; Xiong, Z. User-Oriented AR Assembly Guideline: A New Classification Method of Assembly Instruction for User Cognition. *Int. J. Adv. Manuf. Technol.* **2021**, *112*, 41–59. <https://doi.org/10.1007/s00170-020-06291-w>.
31. Marino, E.; Barbieri, L.; Colacino, B.; Fleri, A.K.; Bruno, F. An Augmented Reality Inspection Tool to Support Workers in Industry 4.0 Environments. *Comput. Ind.* **2021**, *127*, 103412. <https://doi.org/10.1016/j.compind.2021.103412>.
32. Gheisari, M.; Esmaeili, B. Applications and Requirements of Augmented Reality-Based Worker Support Systems in Construction. *Autom. Constr.* **2019**, *101*, 15–24.
33. Schall, G.; Reitmayr, G.; Schmalstieg, D. Fusion of Optical and Inertial Tracking for Robust Augmented Reality Registration. In Proceedings of the 8th IEEE International Symposium on Mixed and Augmented Reality (ISMAR), Orlando, FL, USA, 19–22 October 2009; pp. 57–60.

34. Paladini, M. 3 Different Types of AR Explained: Marker-Based, Markerless & Location. Blippar Blog. 2018. Available online: <https://www.blippar.com/blog/2018/08/14/marker-based-markerless-or-location-based-ar-different-types-of-ar> (accessed on 26 July 2026).
35. Wang, J.; Li, D. The Applications of Fiducial Markers in Robotic Grasping: A Comparison Between AprilTag and ArUco Markers. 2022. Available online: <https://doongli.github.io/doc/grasp.pdf> (accessed on 21 July 2026).
36. Digiacomo, F.; Bologna, F.; Inglese, F.; Stefanini, C.; Milazzo, M. MechaTag: A Mechanical Fiducial Marker and the Detection Algorithm. *J. Intell. Robot. Syst.* **2021**, *103*, 46. <https://doi.org/10.1007/s10846-021-01507-x>.
37. Heselden, J.R.; Paparas, D.; Stevenson, R.L.; Das, G.P. Enhanced STag Marker System: Materials and Methods for Flexible Robot Localisation. *Machines* **2025**, *13*, 2. <https://doi.org/10.3390/machines13010002>.
38. Kalaitzakis, M.; Cain, B.; Carroll, S.; Ambrosi, A.; Whitehead, C.; Vitzilaios, N. Fiducial Markers for Pose Estimation: Overview, Applications and Experimental Comparison of the ARTag, AprilTag, ArUco and STag Markers. *J. Intell. Robot. Syst.* **2021**, *101*, 71. <https://doi.org/10.1007/s10846-020-01307-9>.
39. Wang, Y.; Zhang, S.; Yang, S.; He, W.; Bai, X. Mechanical Assembly Assistance Using Marker-Less Augmented Reality System. *Assem. Autom.* **2018**, *38*, 77–87. <https://doi.org/10.1108/AA-11-2016-152>.
40. Rambach, J.; Pagani, A.; Stricker, D. Augmented Things: Enhancing AR Applications Leveraging the Internet of Things and Universal 3D Object Tracking. In Proceedings of the 2017 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct), Nantes, France, 9–13 October 2017.
41. PTC Inc. Vuforia Engine Developer Library—Model Targets. 2024. Available online: <https://library.vuforia.com/objects/model-targets> (accessed on 26 July 2026).
42. Marinakis, A.; Mania, K.; Antoniadis, A. Augmented Reality for CAD-CAM Training Featuring 3D Interactive Geometric Transformations. *Comput.-Aided Des. Appl.* **2020**, *18*, 561–570. <https://doi.org/10.14733/cadaps.2021.561-570>.
43. Runji, J.M.; Lin, C.Y. Markerless Cooperative Augmented Reality-Based Smart Manufacturing Double-Check System: Case of Safe PCBA Inspection Following Automatic Optical Inspection. *Robot. Comput.-Integr. Manuf.* **2020**, *64*, 101957. <https://doi.org/10.1016/j.rcim.2020.101957>.
44. El Barhoumi, N.; Hajji, R.; Bouali, Z.; Ben Brahim, Y.; Kharroubi, A. Assessment of 3D Models Placement Methods in Augmented Reality. *Appl. Sci.* **2022**, *12*, 10620. <https://doi.org/10.3390/app122010620>.
45. Kušnirák, K.; Kolimár, O.; Kučera, E.; Haffner, O. Proposal for Implementation of Software-Generated 3D Identifiers in Logistics. In Proceedings of the 2025 Cybernetics & Informatics (K&I), 32nd International Conference, Mikulov, Czech Republic, 2–5 February 2025; pp. 211–216. <https://doi.org/10.1109/KI64036.2025.10916470>.
46. Unity Technologies. Unity Engine. 2024. Available online: <https://unity.com/products/unity-engine> (accessed on 8 July 2024).
47. Unity Technologies. Unity Scripting API: NavMeshAgent (Version 2022.3 LTS). 2024. Available online: <https://docs.unity3d.com/2022.3/Documentation/ScriptReference/AI.NavMeshAgent.html> (accessed on 26 July 2026).
48. Unity Technologies. AI Navigation Package Documentation (com.unity.ai.navigation 2.0). 2024. Available online: <https://docs.unity3d.com/Packages/com.unity.ai.navigation@2.0/manual/index.html> (accessed on 26 July 2026).
49. Platt, J.; Ricks, K. Comparative Analysis of ROS-Unity3D and ROS-Gazebo for Mobile Ground Robot Simulation. *J. Intell. Robot. Syst.* **2022**, *106*, 80. <https://doi.org/10.1007/s10846-022-01766-2>.
50. Gentile, D.; Musolino, F.; Dastan, M.; Fiorentino, M. Augmented Reality Glasses for Order Picking: A User Study Comparing Numeric Code, 2D-Map, and 3D-Map Visualizations. *J* **2025**, *8*, 32. <https://doi.org/10.3390/j8030032>.
51. Windhausen, A.; Heller, J.; Hilken, T.; Mahr, D.; Di Palma, R.; Quintens, L. Exploring the Impact of Augmented Reality Smart Glasses on Worker Well-Being in Warehouse Order Picking. *Comput. Hum. Behav.* **2024**, *155*, 108153. <https://doi.org/10.1016/j.chb.2024.108153>.
52. Kišš, J. Using the Robot Operating System for Warehouse Management. Master's Thesis, Slovak University of Technology in Bratislava, Bratislava, Slovakia, 2025. (In Slovak)
53. Adámek, R.; Brabc, M.; Vávra, P.; Dobossy, B.; Formánek, M.; Radil, F. Analytical Models for Pose Estimate Variance of Planar Fiducial Markers for Mobile Robot Localisation. *Sensors* **2023**, *23*, 5746. <https://doi.org/10.3390/s23125746>.
54. García-Ruiz, P.; Romero-Ramírez, F.J.; Muñoz Salinas, R.; Marín-Jiménez, M.J.; Medina-Carnicer, R. Fiducial Objects: Custom Design and Evaluation. *Sensors* **2023**, *23*, 9649. <https://doi.org/10.3390/s23249649>.
55. Robotec.ai. ROS2 For Unity—Native ROS2 bindings for Unity. 2023. Available online: <https://github.com/RobotecAI/ros2-for-unity> (accessed on 26 July 2026).
56. Montgomery, D.C. *Design and Analysis of Experiments*, 9th ed.; John Wiley & Sons: Hoboken, NJ, USA, 2017.
57. Wilcoxon, F. Individual Comparisons by Ranking Methods. *Biom. Bull.* **1945**, *1*, 80–83. <https://doi.org/10.2307/3001968>.
58. McNemar, Q. Note on the Sampling Error of the Difference Between Correlated Proportions or Percentages. *Psychometrika* **1947**, *12*, 153–157. <https://doi.org/10.1007/BF02295996>.
59. Mann, H.B.; Whitney, D.R. On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *Ann. Math. Stat.* **1947**, *18*, 50–60.

60. Welch, B.L. The Generalization of 'Student's' Problem when Several Different Population Variances are Involved. *Biometrika* **1947**, *34*, 28–35.
61. Cohen, J. *Statistical Power Analysis for the Behavioral Sciences*, 2nd ed.; Lawrence Erlbaum Associates: Hillsdale, NJ, USA, 1988.
62. Brata, K.C.; Funabiki, N.; Riyantoko, P.A.; Panduman, Y.Y.F.; Mentari, M. Performance Investigations of VSLAM and Google Street View Integration in Outdoor Location-Based Augmented Reality under Various Lighting Conditions. *Electronics* **2024**, *13*, 2930. <https://doi.org/10.3390/electronics13152930>.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.