This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication.
Citation information: DOI: 10.1109/TMC.2026.3674609,

LLM-Assisted Optimisation of Multi-RIS Placement and Beamforming in Smart Warehouses

Chenyang Yuan, Jinbo Hou, Gang Yu, *Graduate Student Member, IEEE,* Kehai Qiu, *Member, IEEE,* Kezhi Wang, *Senior Member, IEEE,* Haonan Hu, *Member, IEEE,* Jie Zhang, *Senior Member, IEEE*

***Abstract*—In this paper, we propose an optimisation framework for deployment of multiple reconfigurable intelligent surfaces (RISs) to meet the wireless coverage demands for smart warehouses. Specifically, we are the first to formulate a unified network optimisation task that jointly considers RIS placement and beamforming to maximize overall network coverage with a deterministic channel model to accurately describe the multipath effect for the warehouse. To address this problem, we design a hybrid optimisation framework composed of three synergistic modules. (1) A Large Language Model (LLM) acts as a semantic planner that generates physically feasible multi-RIS configurations, jointly determining the placement and beamforming directions guided by structured prompts and environment-aware embeddings. (2) A Genetic Algorithm (GA) module performs local numerical refinements to enhance the precision of LLM-generated solutions under physical constraints. (3) A Diversity Reflection and Correction (DiRect) module evaluates structural similarity among candidate configurations and triggers additional semantic regeneration to maintain exploration diversity. These three modules form an alternating iterative process in which LLM reasoning, GA-based evolution, and DiRect-driven regeneration collectively guide the optimisation toward high-coverage configurations. Extensive simulations validate the effectiveness and robustness of the proposed framework. Compared with traditional heuristics, reinforcement learning methods, and LLM-guided baselines, our hybrid framework achieves 10%–15% higher coverage within 10–20 iterations. The performance consistently scales with the number of RISs and element sizes, and remains stable under varying transmitter positions, demonstrating strong adaptability to complex smart warehouse layouts. Overall, the proposed hybrid optimisation framework provides a scalable and physically grounded solution for RIS-assisted network deployment optimisation in realistic industrial environments.**

***Index Terms*—Reconfigurable intelligent surface, smart warehouse, network optimisation, large language model.**

This research is funded by UKRI Project RISIR under grant no. EP/Y023374/1. Chenyang Yuan, Jinbo Hou, and Gang Yu (e-mail: cyuan7@sheffield.ac.uk; jhou9@sheffield.ac.uk; gangyu@ieee.org) are with the Department of Electronic and Electrical Engineering, The University of Sheffield, S10 2TN Sheffield, U.K. Jinbo Hou is also with Brunel University of London. Kehai Qiu (e-mail: kq218@cam.ac.uk) is with the Department of Computer Science and Technology, University of Cambridge, Cambridge CB3 0FD, U.K., Kehai Qiu is also with Brunel University London. Kezhi Wang (e-mail: kezhi.wang@brunel.ac.uk) is with the Department of Computer Science, Brunel University of London, UB8 3PH, U.K. Haonan Hu (e-mail: haonan.hu@sheffield.ac.uk) is with the School of Electronic and Electrical Engineering, University of Sheffield, UK. Jie Zhang (jie.zhang@ranplanwireless.com) is with the R&D Department, Cambridge AI+ Ltd., CB23 3UY Cambridge, U.K., and also with the R&D Department, Ranplan Wireless Network Design Ltd., CB23 3UY Cambridge, U.K.
(*Corresponding author: Haonan Hu*)

I. Introduction

Smart warehouses have emerged as a key application scenario for Industry 4.0 [1], where ubiquitous wireless coverage is required to support precise inventory tracking, automated storage, and intelligent material handling for the automated guided vehicle (AGV) systems. However, metallic shelves and stored cargo in warehouse infrastructures frequently block the line-of-sight (LOS) path between access points (APs) and AGVs. This may cause poor coverage in certain areas, resulting in serious consequences such as AGV crashing and cargo falling [2], [3]. Ultra-densely deployed networks (UDNs) were proposed to improve indoor wireless coverage [4], [5], which are also considered promising for smart warehouses to ensure AGV communication via LOS paths. However, they will significantly increase the capital expenditure (CAPEX) of the AGV system.

Reconfigurable Intelligent Surface (RIS), an emerging technology in beyond fifth generation (B5G) and sixth generation (6G) networks [6], has become a cost-efficient solution to address the ubiquitous wireless coverage problem in smart warehouses. RISs are planar metasurfaces comprising tens to hundreds of passive reflecting elements, with each of them being capable of adjusting the amplitude and phase of incident waves to realize programmable beamforming. Our previous work [7] demonstrated that the hybrid deployment of RISs and base stations (BSs) can reduce the network CAPEX while ensuring a coverage probability threshold as compared with the UDNs. However, [7] obtained this conclusion by randomly placing the BSs and RISs, without optimising RIS positions or phase shifts, both of which are critical for the practical deployment of AGV systems in smart warehouses.

There have already been some works investigating the indoor RIS deployment. In [8], the effect of RIS position has been investigated in an indoor residential scenario with single RIS and BS for a user equipment (UE). The results showed that the optimal RIS location is related to the distance from the RIS to the BS-UE line, which is not always near the BS and UE ends. In [9], the authors investigated the impact of RIS placement on indoor coverage using a deterministic channel model that captures realistic propagation conditions. The results showed that RIS placement is critical to coverage performance and that the optimal position is highly sensitive to the relative locations of the BS and UE. However, both of these works only consider a single RIS deployment, ignoring the potential brought by multi-RIS deployment, such as scalability and flexibility.

For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

As a result, in [10], the authors demonstrated that deploying multiple RISs can further enhance coverage compared with single RIS placement. However, this work is based on a stochastic channel model, which may not generalize well to practical indoor scenarios with diverse building layouts. In [11], the authors demonstrated that jointly optimising the placement of multiple RISs and BSs based on the deterministic channel model could achieve optimal coverage. The results indicated that placing RISs at proper locations could effectively reduce the required number of BSs in practical indoor scenarios. However, both studies rely on idealized phase alignment, ignoring the practical coupling between RIS placement and discrete phase configuration, which fundamentally limits the applicability of their methods in realistic smart warehouses. To address this gap, subsequent research has incorporated beamforming considerations into multi-RIS placement strategies to better capture the combined impact of their placement on coverage. In [12], the authors proposed a gradient-based algorithm to jointly optimise the continuous positions and orientations of multiple RISs in indoor millimeter-wave networks with blockage, aiming to maximize the overall coverage. In contrast, the authors in [13] adopted a block coordinate ascent (BCA) strategy to iteratively optimise RIS placement and beamforming parameters from a predefined discrete candidate set. Building upon this work in [14], they further integrated Q-learning to accelerate the optimisation process. However, all these works rely on manually defined candidate location sets, limiting the ability to explore the possible placement. Meanwhile, to simplify beamforming, each RIS is approximated as a wide-beam radiator with fixed coverage patterns, overlooking the directional and multipath characteristics of realistic RIS reflections. As a result of these simplified physical assumptions, these methods exhibit limited scalability and generalizability, making them not suitable for effective multi-RIS deployment planning in practical smart warehouses, where addressing poor coverage in certain areas is a critical requirement.

Recently, Large Language Models (LLMs) have attracted increasing attention as a promising paradigm for network optimisation. Unlike conventional optimisation algorithms or traditional machine learning approaches that rely on structured inputs and rigid modeling assumptions, LLMs possess an inherent capacity to process unstructured environment descriptions, reason over complex spatial configurations, and generate feasible deployment strategies expressed in natural language [15]. This capability enables LLMs to efficiently handle multi-RIS placement problems in complex indoor environments with dense obstacles. Current studies have demonstrated that LLMs can leverage in-context learning with prompts to encode task objectives and physical constraints for network optimisation [16]. However, to the best of our knowledge, only two works have investigated their application to indoor wireless placement problems. In [17], the authors first developed a framework that employed an LLM as a zero-shot combinatorial optimiser to efficiently determine the optimal number and placement of BSs in indoor scenarios with superior coverage compared with the Ant Colony Optimisation algorithm. However, due to inherent semantic biases, LLMs tend to concentrate their outputs in high-probability regions and repeatedly generate structurally similar yet suboptimal solutions, lacking mechanisms for global exploration or incremental refinement. Building upon this prior work in [18], the authors proposed an LLM-powered heuristic (LHS) framework, where the LLM generates a high-quality initial solution in each iteration that is subsequently refined by a heuristic algorithm to jointly optimise the placement and beam direction of window-mounted RISs. The results demonstrated that the LHS framework achieved superior coverage and daylight performance compared with the pure LLM-based approach. While this approach improves optimisation efficiency, the heuristic algorithm remains tightly dependent on the LLM-generated solution and incurs considerable computational cost in high-dimensional spaces such as smart warehouses. Consequently, the LHS framework still suffers from limited exploration capacity and tends to converge to local optima, without fully leveraging the semantic reasoning potential of LLMs for broader solution-space coverage. Therefore, the major motivations of this work are summarized as follows:

1) Typical approaches first select RIS locations from predefined candidate sets and then optimise phase shifts. However, in blockage-dominated warehouse environments, RIS beamforming is inherently geometry-dependent and cannot be determined independently of the RIS placement. Consequently, treating placement and beamforming as separate stages fails to capture their intrinsic coupling and often leads to structurally suboptimal coverage performance in practice. Moreover, many existing RIS deployment models are built on idealized assumptions, such as perfect channel knowledge and overly optimistic reflection behavior. These assumptions are useful for analytical studies and upper-bound analysis. However, they conceal important trade-offs in realistic indoor environments and make decoupled deployment strategies unreliable in practice.

2) Joint multi-RIS deployment constitutes a large-scale combinatorial optimisation problem with a high-dimensional search space. To manage this complexity, many existing methods adopt alternating or block-wise update schemes that optimise RIS placement and beamforming in separate stages. However, in multi-RIS scenarios, the coverage contribution of each RIS depends not only on its own configuration but also on the collective placement and beamforming decisions of other RISs. As a result, alternating optimisation breaks the global coupling among RISs, often leading to premature convergence to structurally suboptimal deployment patterns and poor scalability in large warehouse layouts.

3) Joint optimisation of multi-RIS placement and beamforming results in a rapidly expanding combinatorial search space as the environment size and the number of RISs increase. In practical warehouse planning, deployment solutions must be obtained within a reasonable time frame to enable design iteration and timely coverage improvement. However, classical heuristic algorithms scale poorly in such high-dimensional search spaces and tend to concentrate in likelihood-dominant regions, while reinforcement learning (RL) methods often suffer from slow and unstable convergence due to sparse rewards and strict geometric constraints. These challenges motivate the

development of an optimisation framework that can efficiently explore the coupled search space and significantly reduce time-to-solution under realistic deployment conditions.

Accordingly, we develop a deterministic channel model to evaluate the received signal distribution across the warehouse environment. This simulator accounts for the structural layout and physical obstructions of shelves and walls, providing spatially consistent coverage feedback throughout the deployment region. Such feedback is essential for multi-RIS deployment planning, as it enables data-driven assessment of coverage variations under different placement configurations and beam steering decisions. Building upon this foundation, we propose a hybrid optimisation framework that leverages LLMs as semantic optimisers to generate feasible and structurally meaningful multi-RIS configurations, where each configuration jointly specifies the placement and beamforming of all deployed RISs. Guided by task-specific prompts and structured environment embeddings, the LLM produces candidate solutions that jointly determine the placement and beamforming direction of each RIS. This design overcomes the limitations of conventional approaches that rely on random initialization without environmental awareness, often leading to inefficient exploration and suboptimal convergence. To further enhance solution quality, we incorporate the genetic algorithm (GA) module that performs localized adjustments to refine the LLM-generated candidates, enhancing the numerical precision of semantically generated solutions. Additionally, we propose a novel diversity reflection and correction (DiRect) module, which is designed to prevent the LLM optimiser from converging to structurally homogeneous solutions. By evaluating structural similarity between new candidates and historically best-performing configurations, DiRect ensures sufficient diversity is maintained throughout the optimisation process. Collectively, these modules are systematically integrated into an optimisation framework enabling efficient discovery of high-performance and generalizable RIS deployment strategies in practical smart warehouses. The main contributions are summarized as follows:

1) We formulate a joint optimisation framework for multi-RIS placement and beamforming to maximize coverage in blockage-dominated smart warehouse environments, explicitly accounting for their intrinsic coupling. Unlike existing studies that decouple variables under idealized assumptions, our formulation is built upon a deterministic channel model that provides spatially consistent coverage evaluation across the entire warehouse layout.

2) To address the large-scale combinatorial nature and coordination challenges of multi-RIS deployment, we propose a hybrid optimisation framework that integrates LLM-based semantic reasoning with physics-aware search mechanisms. Structured prompts and environment embeddings guide the LLM to generate globally informed and coverage-aware deployment proposals, providing effective initialization and directional guidance in the high-dimensional search space. To further improve solution quality and scalability, a GA module performs constraint-aware local refinement, while a DiRect mechanism enforces semantic diversity across candidate configurations to prevent premature convergence. Together, these components operate in a coordinated manner under a unified optimisation objective, enabling efficient exploration of the coupled multi-RIS configuration space without breaking global deployment coherence.

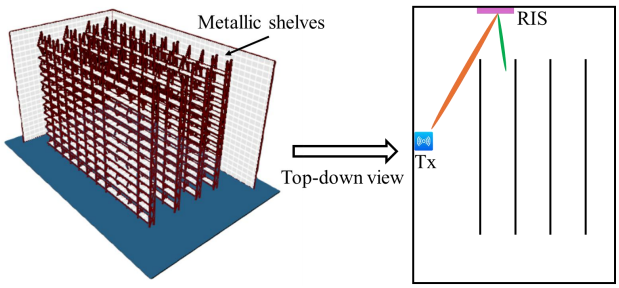

Fig. 1. Overview of the warehouse simulation scenario with metallic shelves, deterministic ray-tracing propagation, and programmable RIS deployment.

3) We conduct extensive simulations under varying RIS quantities, element sizes, and BS positions to comprehensively evaluate the proposed framework. The results demonstrate that our method consistently achieves 10–15% higher coverage, whereas classical heuristic algorithms, RL methods, and LLM-only baselines require hundreds of iterations or fail to converge to high-coverage solutions. Beyond convergence behaviour, we further evaluate the overall runtime and show that the proposed framework reaches target coverage within tens of minutes, while GA and RL based methods incur hour or day level computation under the same deployment scale, highlighting its significant advantage in time-to-solution for practical warehouse planning. Moreover, experiments with multiple representative LLMs confirm that the proposed prompting and embedding design enables stable optimisation behaviour across heterogeneous model architectures. These results further demonstrate that the framework is largely model-agnostic and robust to different LLM implementations.

The rest of this paper is organized as follows. Section II establishes the system model and formulates the joint optimisation problem. Section III provides the proposed LLM-assisted hybrid optimisation framework. Simulation results are performed in Section IV. Section V concludes the paper.

II. System Model and Problem Formulation

As illustrated in Fig. 1, we consider a smart warehouse environment modeled as a two-dimensional rectangular domain $\mathcal{D} = [0, L] \times [0, W]$, where L and W denote the warehouse length and width, respectively. This 2D layout corresponds to the top-down projection of the actual environment. Vertical metallic shelves are represented as line-segment obstructions aligned along the vertical axis. Each metallic shelf with the stored goods forms an electrically large composite surface whose effective surface height variation remains small relative to the wavelength, thereby satisfying the Rayleigh electromagnetic smoothness condition [19]. The dominant propagation contributions are governed by the LOS and the first-order specular reflection components. Therefore, the diffuse scattering effects, which primarily introduce secondary corrections to the reflected fields are neglected at the deployment scale considered in this work.

A. Signal Propagation Model

We adopt a deterministic channel model that accounts for LoS and first-order specular reflection paths, consistent with the environment characterization [20]. Given the transmitter location $\mathbf{x}_{\mathrm{Tx}} \in \mathcal{D}$, a propagation path to any receiver position $\mathbf{x}_j \in \mathcal{D}$ is regarded as valid if no segment of the path intersects any shelf boundary. Each valid path ℓ is therefore either a LoS trajectory or a first-order specular reflection from the warehouse walls. Higher-order reflections involving multiple bounces, material absorption or scattering, and small-scale fading are not considered. This simplification is adopted to maintain computational tractability and focus on the dominant contributions governing large-scale signal behavior. Denoting the total path length by d_ℓ, the corresponding complex-valued signal component is expressed as

$$s_\ell = \frac{1}{d_\ell} \exp\big(-jkd_\ell\big),$$

where λ is the carrier wavelength and $k = \frac{2\pi}{\lambda}$. Let $\mathcal{L}(\mathbf{x}_j)$ denote the set of all valid LOS and first-order reflection paths to $\mathbf{x}_j$. The aggregated signal received before any RIS contribution is given by

$$s_{\mathrm{raw}}(\mathbf{x}_j) = \sum_{\ell \in \mathcal{L}(\mathbf{x}_j)} s_\ell.$$

This model provides a spatially consistent baseline field distribution across the warehouse layout, which is subsequently combined with RIS-assisted components for coherent superposition and coverage optimisation.

B. RIS Beamforming Model

To improve wireless coverage in regions with limited direct and reflected paths, we deploy a set of RISs along the horizontal boundaries of the domain $\mathcal{D}$. Let $\mathcal{I} = \{1, 2, \ldots, I\}$ denote the set of deployed RISs. Each RIS is implemented as a uniform planar array of passive reflecting elements arranged in a rectangular grid. The total number of elements is denoted by M, defined as $M = M_x \times M_y$, where M_x and M_y are the numbers of elements along the horizontal and vertical dimensions, respectively. The element spacing is set to $\delta = \lambda/2$ in both directions to satisfy spatial sampling requirements.

The center position of the i-th RIS is denoted by $\mathbf{r}_i = (r_{ix}, r_{iy})$, constrained to lie along the boundary of the smart warehouse. The associated beam-steering reference point is denoted by $\mathbf{u}_i = (u_{ix}, u_{iy}) \in \mathcal{D}$. It specifies the desired reflection direction along which the RIS aligns the outgoing wavefront. In line with standard RIS beamforming [21], the resulting phase profile forms a main-lobe radiation pattern with a finite beamwidth and thus provides regional signal enhancement around $\mathbf{u}_i$. This spatially distributed effect is explicitly accounted for in our pixel-wise field evaluation across region $\mathcal{D}$, which adopts the ray-based formulation commonly used in RIS-enabled indoor propagation modeling [22]. For the m-th element of the i-th RIS, we define the distances corresponding to the two propagation segments:

$$d_{\mathrm{Tx},im} = \|\mathbf{x}_{\mathrm{Tx}} - \mathbf{x}_{im}\|_F, \quad d_{im,\mathbf{u}_i} = \|\mathbf{x}_{im} - \mathbf{u}_i\|_F, \quad (1)$$

where $\mathbf{x}_{im}$ and $\mathbf{x}_{\mathrm{Tx}}$ denote the positions of the m-th element of the i-th RIS and the transmitter, respectively. To ensure that the reflected waves from each RIS element coherently combine toward the focusing direction, we compensate for the propagation-induced phase delay accumulated along the two segments from the transmitter to the RIS element and from the RIS element toward $\mathbf{u}_i$. Specifically, the cumulative propagation phase for element (i, m) is

$$\psi_{im} = k\big(d_{\mathrm{Tx},im} + d_{im,\mathbf{u}_i}\big), \quad (2)$$

and the corresponding phase compensation applied by each RIS element is expressed as [9], [23]

$$\phi_{im} = \mathrm{mod}\big(-\psi_{im},\, 2\pi\big), \quad (3)$$

which is then quantized to the feasible hardware resolution within $[0, 2\pi)$. To evaluate the field at an arbitrary sampling location $\mathbf{x}_j \in \mathcal{D}$ under this RIS configuration, we define the second segment distance

$$d_{im,j} = \|\mathbf{x}_{im} - \mathbf{x}_j\|_F, \quad (4)$$

and the complex-valued contribution from element (i, m) to $\mathbf{x}_j$ as

$$s_{im}^{\mathrm{RIS}}(\mathbf{x}_j) = \frac{\exp\big(-jk\big(d_{\mathrm{Tx},im} + d_{im,j}\big)\big)\, \exp\big(j\phi_{im}\big)}{d_{\mathrm{Tx},im} + d_{im,j}}. \quad (5)$$

The total contribution from all RIS elements is then given by

$$s_{\mathrm{RIS}}(\mathbf{x}_j) = \sum_{i=1}^{I} \sum_{m=1}^{M} s_{im}^{\mathrm{RIS}}. \quad (6)$$

For a deployment to be feasible, the two segments of the path must be unobstructed by any metallic shelves.

C. Performance Metrics and Optimisation Problem

To characterize the coherent superposition of all propagation contributions, the domain $\mathcal{D}$ is discretized into a uniform grid with spacing δ_{grid}, yielding the candidate receiver set $\mathcal{N} = \{\mathbf{x}_n \in \mathcal{D} \mid n = 1, \ldots, N\}$. In this set, each vector $\mathbf{x}_n$ specifies the spatial location of a candidate receiver. For each $\mathbf{x}_n$, the total complex-valued received signal is expressed as

$$s(\mathbf{x}_n) = s_{\mathrm{LOS}}(\mathbf{x}_n) + \sum_{r=1}^{R_n} s_{\mathrm{ref}}^{r}(\mathbf{x}_n) + s_{\mathrm{RIS}}(\mathbf{x}_n), \quad (7)$$

where R_n denotes the number of unblocked reflection paths reaching $\mathbf{x}_n$. Under a free-space path loss model [24], the received signal strength at position $\mathbf{x}_n$ is given by

$$P_n = P_{\mathrm{Tx}} \cdot \left(\frac{\lambda}{4\pi}\right)^2 \cdot \big|s(\mathbf{x}_n)\big|^2, \quad (8)$$

where P_{Tx} denotes the transmit power. A receiver location $\mathbf{x}_n$ is considered successfully covered if the received power exceeds a minimum threshold γ. Accordingly, the coverage indicator function is thus defined as

$$\mathbb{I}_n = \begin{cases} 1, & P_n \geq \gamma, \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

The total number of covered grid points is formally given as

$$N_{\text{cov}} = \sum_{n \in \mathcal{N}} \mathbb{I}_n. \tag{10}$$

Consequently, the coverage ratio is defined as

$$\mathcal{P}_{\text{cov}} = \frac{N_{\text{cov}}}{N}. \tag{11}$$

Given a fixed number of RISs I, the optimisation aims to jointly determine the deployment positions and beamforming target positions of all RISs. The set of decision variables is defined as

$$\mathcal{S} = \left\{ \left((r_{ix}, r_{iy}),\, (u_{ix}, u_{iy}) \right) \,\middle|\, i = 1, \ldots, I \right\}. \tag{12}$$

The optimisation problem is formulated as

$$\begin{aligned}
\max_{\mathcal{S}} \quad & \mathcal{P}_{\text{cov}}(\mathcal{S}) && \text{(P1)} \\
\text{s.t.} \quad & r_{ix} \in [x_{\min}, x_{\max}], \quad \forall i, && \text{(C1)} \\
& r_{iy} \in \{y_{\min}, y_{\max}\}, \quad \forall i, && \text{(C2)} \\
& u_{ix} \in [0, L], \quad \forall i, && \text{(C3)} \\
& u_{iy} \in [0, W], \quad \forall i, && \text{(C4)} \\
& \mathbb{B}(\mathbf{r}_i, \mathbf{u}_i) = 1, \quad \forall i, && \text{(C5)}
\end{aligned}$$

where $\mathbb{B}(\cdot)$ is a boolean indicator equal to 1 if the line segment between the RIS and its beamforming target is unblocked by any metallic shelves, and 0 otherwise.

III. LLM-BASED HYBRID FRAMEWORK

This section presents the hybrid optimisation framework for multi-RIS placement and beamforming in smart warehouse environments. As illustrated in Fig. 2, the framework consists of three interoperable modules: an LLM-driven semantic planner that formulates globally meaningful RIS configurations, a GA-based refinement module that performs constraint-aware numerical adjustment, and a DiRect-based semantic diversity mechanism that maintains exploration capability by preventing convergence to structurally similar layouts. These modules operate as independent optimisation steps within a unified iterative process, enabling the framework to progressively construct high-coverage configurations while balancing global reasoning, local physical refinement, and structural diversity.

A. Initialization Module

The initialization module establishes the starting point of the proposed framework by generating an informed candidate RIS deployment plan based on a structured semantic prompt and LLM inference. This stage consists of two core components: the initialization prompt and the LLM initializer.

1) Initialization Prompt: The initialization prompt plays a central role in transforming the RIS deployment task from an unconstrained text generation problem into a structured spatial optimisation process. It consolidates four key components: a precise task definition specifying that the objective is to select RIS placements and beamforming targets to maximize coverage; a description of relevant physical constraints including LOS feasibility and boundary restrictions; structured encodings of the environment in the form of discretized matrices capturing warehouse geometry and baseline signal distributions; and an explicit JSON-style output specification enumerating each RIS deployment. In particular, as illustrated in Fig. 3 and Fig. 4(a), the prompt includes visual embeddings that encode both the floor layout information and the coverage conditions. The warehouse layout matrix assigns numeric codes to each grid cell indicating free space, obstructions, and transmitter location. Complementarily, the baseline signal power distribution and its binary coverage map enable the model to identify coverage deficits that should guide the initial RIS deployment. These embeddings compress high-dimensional spatial and propagation data into compact representations that provide the model with a clear context for reasoning.

2) LLM Initializer for Candidate Generation: Upon receiving the prompt, the LLM initializer integrates the encoded environment, task objectives, and the binary coverage representation to produce an initial RIS deployment $\mathcal{S}_0$. It specifies the positions of all RIS units along the designated warehouse boundaries and their corresponding beamforming targets. The initialization module thus establishes a baseline configuration that will be iteratively refined in subsequent optimisation cycles to improve overall coverage performance.

B. Visualization-Aware Environment Embedding

To support the LLM in understanding the spatial constraints and propagation characteristics inherent in the smart warehouse, the framework incorporates structured visual representations that serve as compact and interpretable encodings of key information. By combining continuous received power distributions with their binarized coverage maps, these embeddings provide the model with a concise yet informative description of both the baseline signal conditions and the effects of RIS deployment. This design compresses high-dimensional propagation data into structured matrices that can be consistently processed by the LLM, enabling accurate reasoning about spatial feasibility, blockage constraints, and coverage optimisation.

As illustrated in Fig. 3, the first visualization captures the warehouse layout in a discretized 2D matrix format. Each grid cell is assigned a distinct numeric code indicating its role in the environment: zero for the transmitter position, one for metallic shelves that block signal propagation, and two for free space. This encoding conveys the spatial dimensions, physical obstructions, and Tx placement in a unified representation. By providing this structured map to the initialization prompt, the model can directly infer where RIS placements are feasible, where blockage constraints apply, and how the geometry defines the propagation environment.

As shown in Fig. 4, the framework further incorporates two pairs of complementary visual matrices to describe signal coverage conditions. Fig. 4(a) depicts the baseline received power distribution computed by the deterministic ray-tracing model in the absence of RIS deployments, together with its corresponding binarized coverage map. This binarized matrix, where grid cells exceeding the coverage threshold are set to one and all others to zero, enables the initialization

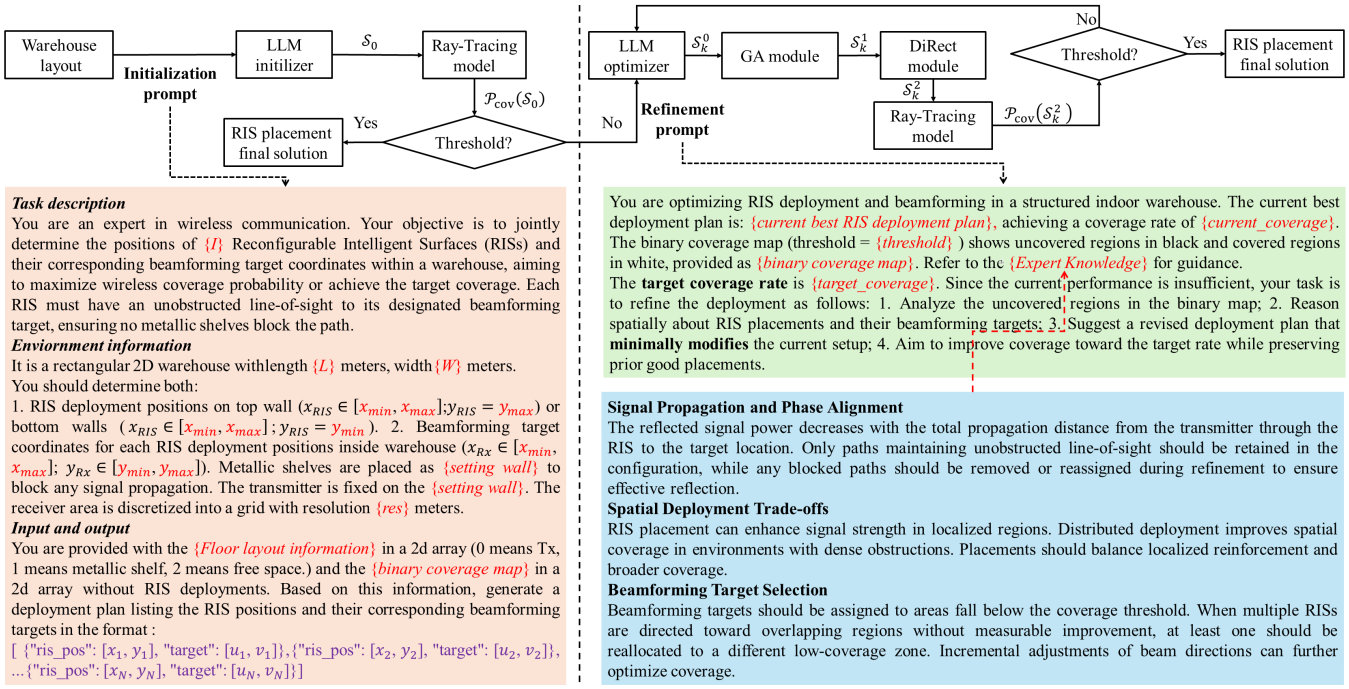

Fig. 2. The structure of the proposed optimisation framework, including three prompts.

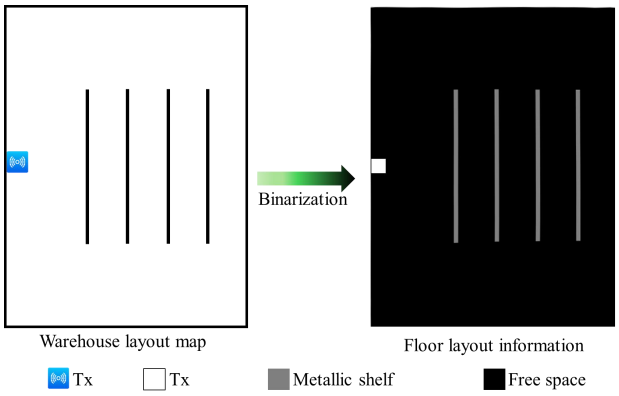

Fig. 3. Illustration of the warehouse layout encoding: (a) spatial distribution of the Tx and metallic shelves in the warehouse; (b) visualization of the floor layout information as a discretized 2D array representing the structured input provided to the initialization Prompt.

prompt to identify areas with severe coverage deficits. In the iterative optimisation stage, Fig. 4(b) shows the received power distribution achieved by a given RIS configuration and its binarized representation. This updated matrix serves as input to the refinement prompt, informing the model of the coverage improvements obtained by the current deployment and highlighting residual under-served regions.

C. Refinement Module

Once the initial candidate deployment $\mathcal{S}_0$ is established, the refinement module is activated to iteratively improve coverage performance until the target threshold is satisfied.

1) Refinement Prompt for Iterative Improvement: During each refinement iteration, the framework provides the LLM optimiser with feedback describing the current deployment status. Specifically, the prompt includes the latest RIS deployment plan with the corresponding global coverage rate achieved so far. As illustrated in Fig. 4(b), the ray-tracing model simulates the received power distribution resulting from this deployment. Then produce a binary coverage map, where each grid indicates whether the received power meets the coverage criterion. By providing this structured representation, the prompt enables the LLM to analyze the spatial distribution of uncovered regions and to reason about targeted adjustments to improve overall coverage. This structured feedback ensures that each refinement cycle begins with an up-to-date assessment of coverage performance.

2) LLM Optimiser and Candidate Update: Upon receiving the refinement prompt, the LLM optimiser generates a revised candidate deployment $\mathcal{S}_k^0$. This candidate reflects incremental adjustments to the RIS placements and beamforming targets aimed at addressing residual coverage gaps while minimally modifying effective configurations from previous iterations. The resulting candidate $\mathcal{S}_k^0$ is then forwarded to the DiRect module and the evolutionary mutation module for diversity correction and local adjustment.

D. GA-Based Local Refinement

After the LLM proposes a candidate configuration $\mathcal{S}_k^0$, the GA module performs localized refinement to improve coverage while preserving geometric feasibility. Let $\mathcal{G} = \{1, 2, \cdots, G\}$ and $\mathcal{P} = \{1, 2, \cdots, P\}$ denote the index set of GA generation and population, respectively. For the g-th generation, the p-th population candidate of i-th RIS and its corresponding beamforming focus point are initialized

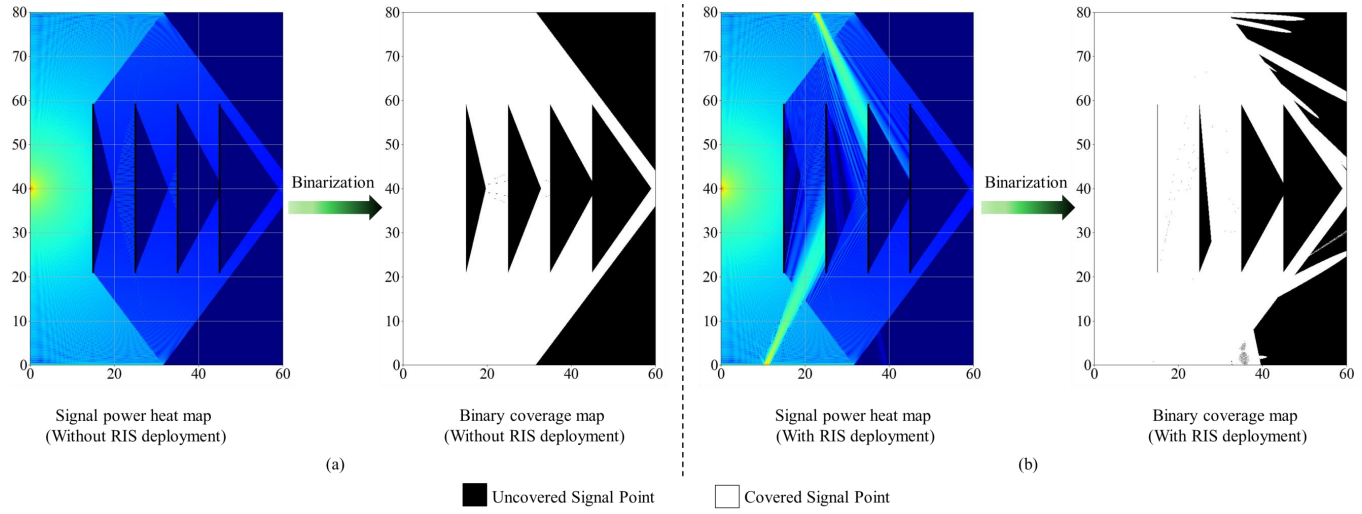

Fig. 4. Illustration of the coverage embedding process: (a) signal power distribution and corresponding binary coverage map of the smart warehouse without RIS deployment, used as input to the initialization prompt; (b) signal power distribution and binary coverage map reflecting the optimised RIS deployment, provided to the refinement prompt.

by applying zero-mean Gaussian perturbations to the LLM-generated $\mathcal{S}_k^0$ as

$$\mathbf{r}_{i,p}^{(0)} = \Pi(\mathbf{r}_i + \boldsymbol{\varepsilon}_{r,i,p}), \qquad \mathbf{u}_{i,p}^{(0)} = \Pi(\mathbf{u}_i + \boldsymbol{\varepsilon}_{u,i,p}), \tag{13}$$

where $\boldsymbol{\varepsilon}_{r,i,p}$ and $\boldsymbol{\varepsilon}_{u,i,p}$ are samples from zero-mean Gaussian distributions, and $\Pi(\cdot)$ denotes the Projection operator onto the feasible deployment set defined by boundary and LOS constraints. To gradually transition from coarse spatial adjustment to fine-grained tuning, the perturbation scale follows an annealing schedule:

$$\sigma_{g+1} = \alpha\sigma_g, \quad g \in \mathcal{G}, \tag{14}$$

where α denotes a fixed decay coefficient chosen between 0.6 and 0.9 to gradually reduce the perturbation scale across generations. Then, the g-th population member $(\mathbf{r}_{i,p}^{(g)}, \mathbf{u}_{i,p}^{(g)})$ is evaluated using the deterministic channel model to compute the coverage ratio $C_p^{(g)}$. The best-performing refined configuration is then selected as

$$\mathcal{S}_k^1 = \arg\max_{p\in\mathcal{P},\, g\in\mathcal{G}} C_p^{(g)}. \tag{15}$$

This refinement procedure does not perform global evolutionary search. Instead, it provides lightweight, constraint-preserving numerical adjustment around the LLM-proposed semantic configuration, allowing the framework to jointly exploit high-level spatial reasoning and fine-grained physical optimisation.

E. Semantic Diversity Correction via DiRect

While the GA-based refinement improves the physical quality of the candidate configuration $\mathcal{S}_k^1$, repeated local perturbations may still drive the search toward structurally similar RIS layouts, resulting in premature convergence. To address this issue, the DiRect module is executed once after the GA refinement to introduce controlled semantic diversity, without disrupting the convergence gained from high-performing configurations.

Let $\mathcal{P}^{(g)}$ denote the set of all candidate configurations produced during the g-th GA refinement round, each evaluated using the deterministic channel model. These population members are ranked based on their achieved coverage performance $\mathrm{Cov}(\mathcal{S})$, and the top-k highest-performing candidates are retained as the elite set

$$\mathcal{E}^{(g)} = \underset{\mathcal{S}\in\mathcal{P}^{(g)}}{\text{Top-k}}\big(\mathrm{Cov}(\mathcal{S})\big). \tag{16}$$

These elite configurations serve as stable high-quality references and are kept unmodified to maintain convergence stability while guiding the diversity evaluation.
Next, for any candidate configuration $\mathcal{S}^{\mathrm{c}} \in \mathcal{P}^{(g)}$, its structural similarity to each elite configuration $\mathcal{S}^{\mathrm{e}} \in \mathcal{E}^{(g)}$ is assessed by the weighted position–target distance:

$$D\big(\mathcal{S}^{\mathrm{c}}, \mathcal{S}^{\mathrm{e}}\big) = \frac{1}{I}\sum_{i=1}^{I}\Big(w_r\big\|\mathbf{r}_i^{\mathrm{c}} - \mathbf{r}_i^{\mathrm{e}}\big\|_2 + w_u\big\|\mathbf{u}_i^{\mathrm{c}} - \mathbf{u}_i^{\mathrm{e}}\big\|_2\Big), \tag{17}$$

where w_r and w_u are weighting factors that balance the emphasis on RIS placement and beamforming target diversity, respectively. This metric captures structural similarity in both placement and beam-steering space, allowing DiRect to distinguish configurations that are geometrically different but functionally equivalent or vice versa. If the minimum distance across all elite solutions falls below a predefined threshold τ, formally

$$\min_{\mathcal{S}^{\mathrm{e}}\in\mathcal{E}^{(g)}} D\big(\mathcal{S}^{\mathrm{c}}, \mathcal{S}^{\mathrm{e}}\big) \le \tau, \tag{18}$$

then $\mathcal{S}^{\mathrm{c}}$ is considered structurally redundant. In this case, the DiRect module invokes a regeneration prompt that conditions on: (i) the current elite set $\mathcal{E}^{(g)}$, (ii) the warehouse geometry and LOS feasibility constraints, and (iii) the uncovered or weak-coverage regions in the binary coverage map. The LLM is thereby guided to synthesize a new replacement configuration $\mathcal{S}_k^2$ that remains feasible but explores spatial patterns not represented by the elite solutions. This one-time

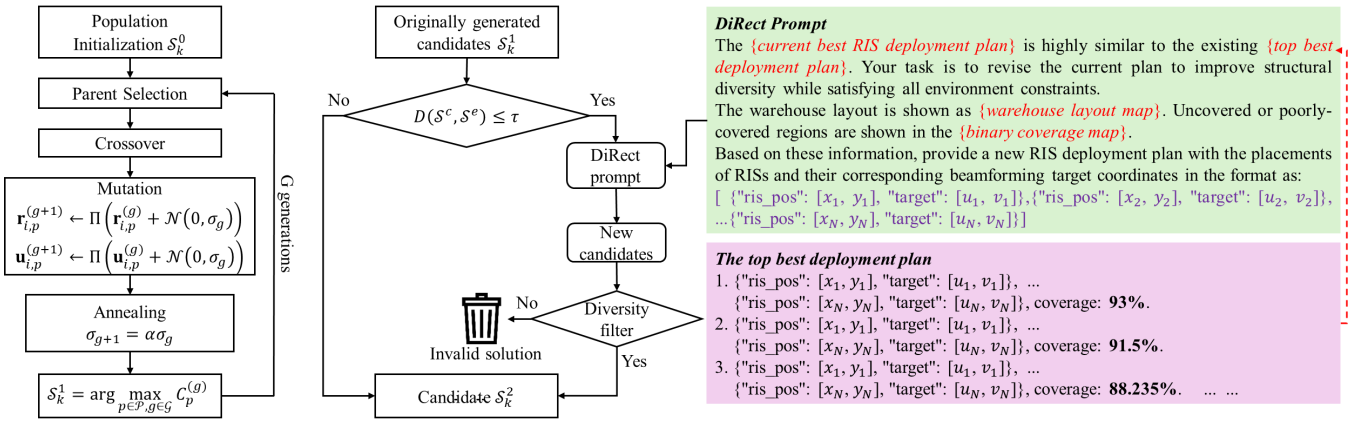

(a) GA module workflow. (b) DiRect module workflow including similarity check and prompt-guided regeneration.

Fig. 5. Framework integrating complementary refinement mechanisms: (a) GA module performing constraint-aware local optimisation of RIS placements and beamforming targets; (b) DiRect module executing semantic regeneration to ensure structural diversity of candidate configurations.

DiRect correction introduces semantically diverse yet physically meaningful, coverage-driven configurations, enabling controlled exploration beyond likelihood- or mutation-induced stagnation while preserving the high-performing geometric patterns identified so far.

IV. Simulation Results

In this section, numerical simulations will be performed to evaluate the achievable communication performance and to verify the feasibility of the proposed LLM-assisted hybrid optimisation of multiple RIS placement and beamforming under different warehouse scales. In addition to the proposed scheme, several counterparts are presented for comparison. In simulation figures, the legends are described below.

- `LLM_GA_DiRect (Proposed)`: This scheme represents our proposed hybrid framework. It integrates semantic spatial reasoning provided by the LLM, local numerical refinement via GA-based evolution, and diversity enforcement through DiRect-driven regeneration. These modules operate in an alternating and interactive manner, where the LLM generates new deployment proposals, the GA performs lightweight local refinements, and the DiRect module identifies structural similarity among candidate configurations and triggers regeneration to preserve exploration diversity.
- `LLM_GA`: This variant retains the LLM-guided semantic initialization and GA-based local search but disables the DiRect module. It is used to evaluate the role of diversity maintenance and to understand the impact of limited exploration in high-dimensional optimisation spaces.
- `LLM`: This scheme uses only the language model to infer RIS placement and beamforming targets based on environmental semantics. It serves to assess the standalone capability of LLMs in spatial reasoning and optimisation without any iterative correction or refinement.
- `GA`: This benchmark applies a traditional genetic algorithm initialized from scratch without leveraging semantic knowledge. It serves as a classical heuristic baseline to understand the effectiveness of uninformed evolutionary search.
- `TD3`: This scheme uses Twin Delayed Deep Deterministic Policy Gradient (TD3) to optimise multi-RIS deployment. The state contains the current positions of all RISs, and the action updates their positions and the beamforming target coordinates. The reward equals the coverage achieved in the warehouse environment. We train TD3 for up to 1500 episodes with early stopping once the target coverage is reached. The hyperparameters follow standard TD3 settings: two-layer 256–256 MLPs for both actor and critic network, learning rates of 3×10^{-4} for both networks, discount factor $\gamma = 0.99$, Polyak coefficient $\tau = 0.005$, batch size 256, replay buffer size 10^6, exploration noise with standard deviation 0.1, and target policy smoothing noise 0.2 clipped at 0.5.
- `Random`: This serves as the lower-bound reference, where RIS placements and beamforming targets are generated randomly without any optimisation. It reflects the achievable performance under uninformed decision making.

A. Simulation Setup and Parameters

In simulations, metallic shelves are modeled as perfect electric conductors (PECs), acting as fully opaque and reflective obstacles that block both LOS and reflected propagation paths. The shelves are placed at $x = \{20, 30, 40, 50\}$ m and span vertically from $y \in [20, 60]$ m, representing typical industrial storage aisles. Each RIS is deployed along the top or bottom warehouse boundaries to reflect practical mounting feasibility, and a minimum spacing of 0.5 m is enforced between adjacent RIS units to avoid redundant placement. The receiver region is uniformly sampled with spacing Δ_g, and a grid point is considered covered if its received power exceeds the threshold γ based on the deterministic channel model in Section II. For most important simulation parameters, their setting values are listed in Table I and some may act as coordinate variables in some simulation figures.

TABLE I
SIMULATION PARAMETERS

Notation	Description	Value
L	Warehouse length	60 m
W	Warehouse width	80 m
$(x_{\mathrm{Tx}}, y_{\mathrm{Tx}})$	BS location	$(0, 40)$
f_c	Carrier frequency	3 GHz
P_{Tx}	Transmit power	23 dBm
I	Number of RISs	4
M	Number of RIS elements	256
δ	RIS element spacing	$\lambda/2$
Δ_g	Receiver grid resolution	0.1 m
γ	Coverage threshold	-50 dBm
P	GA population size	5
G	GA generations	4
k	DiRect elite set size	3
σ_0	Initial perturbation scale	1.2 m
α	Perturbation decay coefficient	0.75
k	DiRect elite set size	3
τ	Structural-similarity threshold	0.5

TABLE II
OVERALL RUNTIME AND RESOURCE UTILIZATION UNDER IDENTICAL HARDWARE CONDITIONS.

Method	Runtime/Iter. (min)	Total Runtime (min)	GPU Util. (%)	Memory Util. (%)
Proposed (LLM-IT)	2.69±0.51 (0.42±0.17)	**20.32±9.55 (7.72±0.76)**	3.2±1.4 (3.2±1.4)	64.2±7.8 (64.2±7.8)
LLM_GA (LLM-IT)	1.55±0.19 (0.32±0.47)	62.06±7.91 (10.72±0.93)	3.7±0.5 (3.7±0.5)	61.2±4.8 (61.2±4.8)
LLM	**0.35±0.09**	162.20±25.20	4.8±1.3	63.5±2.1
GA	1.13±0.37	1562.67±483.17	0.9±1.1	76.3±8.5
TD3	1.64±0.39	1394.44±186.40	56.3±21.3	76.8±12.1

B. Overall Runtime and Computational Efficiency

We adopt a unified metric based on actual runtime across all methods to evaluate computational cost. All experiments deploy 6 RISs and target a coverage level of 90%. We conduct twenty independent trials for each method. We set a maximum of 1500 iterations and apply an early stopping rule that terminates the run once the target coverage is achieved. As shown in Table II, the `Proposed` framework achieves the shortest total runtime. Its cumulative LLM inference time (LLM-IT) is approximately 7.7 minutes, accounting for only a limited fraction of the overall runtime. It improves coverage more rapidly than the other methods because the LLM provides strong global guidance, GA performs effective local refinement, and DiRect maintains solution diversity. These components allow the method to reach the target with far fewer effective optimisation rounds. A key advantage of LLM-based methods is their ability to leverage semantic priors and spatial reasoning. The LLM identifies promising deployment regions from the warehouse layout and provides strong global initialization. GA then refines these solutions efficiently. In contrast, `GA` relies on random mutations and `TD3` relies on gradient-based updates. Both mechanisms lack semantic guidance and therefore explore the search space slowly. As a result, LLM-based optimisation converges much faster and achieves significantly lower total runtime than `GA` and `TD3`. `LLM_GA` converges faster than `GA` because the LLM generates better starting points for the search. `LLM` runs extremely quickly per iteration, but it needs many more iterations to correct suboptimal placements. This slow improvement leads to a significantly longer total runtime despite its low per-iteration cost.

TABLE III
PERFORMANCE COMPARISON OF DIFFERENT LLMS AS SEMANTIC PLANNERS UNDER THE SAME MULTI-RIS OPTIMISATION TASK, PROMPTING STRUCTURE, AND DECODING CONFIGURATION.

Model	Params (B)	Final Cov. (%)	Avg. Iter. Time (s)
GPT-4o	≈ 200	85.6	13.1
GPT-3.5-Turbo	N/A	74.3	4.9
DeepSeek-3.2	N/A	83.0	7.0
Gemini-2.0	N/A	85.0	9.0
Grok-4	N/A	81.0	7.0
Qwen-3	N/A	78.0	4.0
Kimi-K2	≈ 32	80.0	3.5
Claude-4	N/A	82.0	3.5

C. Evaluation Across Different LLM Models

To evaluate the robustness and generalization of the proposed hybrid framework, we conduct a comparative analysis using a set of representative LLMs acting as semantic planners for the identical multi-RIS placement and beamforming task. While GPT-4o serves as the primary planner in this work, we extend our evaluation to include other state-of-the-art models, specifically GPT-3.5-Turbo, DeepSeek-3.2, Gemini-2.0, Grok-4, Qwen-3, Kimi-K2, and Claude-4. To ensure a fair comparison, all models are accessed with their official API interfaces and subjected to identical structured prompts and visual embeddings, comprising floor layout maps and binary coverage maps. Furthermore, the decoding configuration remains consistent across all models, with the temperature and top-p values fixed at 0.5 and 0.9, respectively. The simulation results in Table III demonstrate a clear correlation between reasoning capacity and the resulting coverage probability of all models. Specifically, frontier-level models with larger parameter scales, such as GPT-4o, Gemini-2.0, and DeepSeek-R1, consistently achieve higher final coverage in the multi-RIS deployment task. This superior performance is attributed to their enhanced spatial and semantic reasoning capabilities, which allow them to effectively process the complex coupling between RIS variables and their resulting geometric blockage relationships. While these high-capacity models exhibit longer average inference latency in each iteration because of their deeper computation and the generation of more comprehensive internal reasoning chains, they produce more globally consistent deployment proposals. In contrast, lightweight models, such as Qwen-3, Kimi-K2, and GPT-3.5-Turbo, offer significant advantages in terms of inference efficiency but deliver moderately reduced coverage. Their solutions tend to prioritize local geometric constraints rather than optimising the global coverage probability across the warehouse layout. These results highlight a fundamental performance-efficiency trade-off within the proposed framework, where larger LLMs

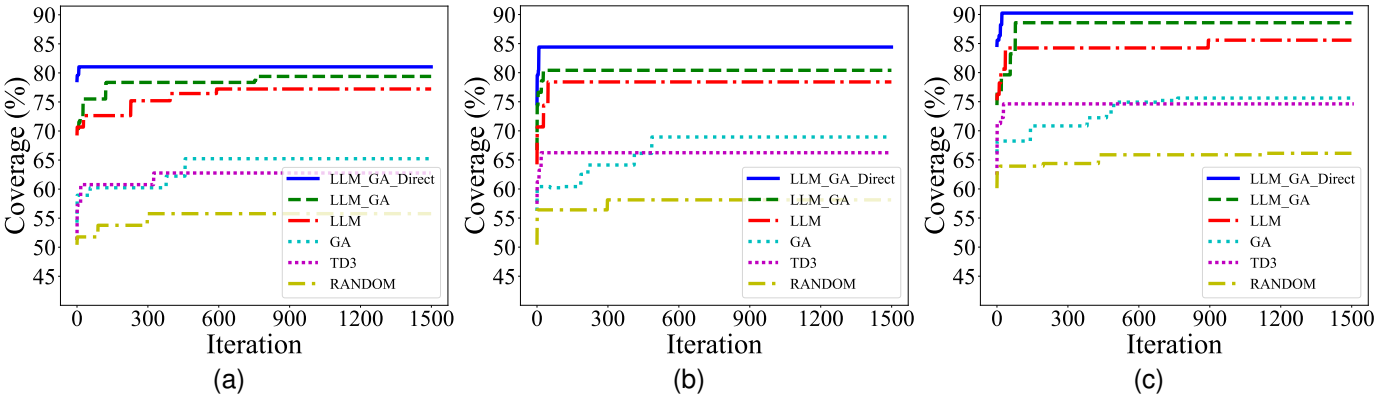

Fig. 6. Coverage behavior of all schemes in the joint optimisation of multi-RIS deployment and beamforming in a smart warehouse, including the proposed scheme and several other counterparts. Subfigures (a)–(c) correspond to scenarios with 2, 4, and 6 deployed RISs, respectively.

provide more robust spatial planning for complex multi-RIS scenarios and more compact models can be selected for time-sensitive applications that require a balance between reasoning depth and computational latency.

D. Convergence Behaviour under Large-Scale Deployment

Fig. 6 presents the coverage achieved at each iteration within a large warehouse environment of 60×80 m^2 when the number of RISs is set to two, four, and six, respectively. It should be noted that in all convergence plots, the iteration index represents the total number of alternating optimisation steps executed by the LLM, GA, and DiRect modules. The proposed `LLM_GA_DiRect` framework consistently delivers the best performance, demonstrating accelerated convergence behavior and superior steady-state coverage compared with all baselines. Notably, in the most challenging six-RIS case, it approaches near-optimal coverage within the first twenty-two iterations, highlighting its ability to efficiently explore and refine deployment configurations even in high-dimensional settings. This strong performance is attributed to the synergy between semantic initialization by the LLM, spatial refinement via genetic mutation, and diversity enforcement through the DiRect module. Together, these components enable the framework to generate high-quality initial proposals, adjust them based on environmental constraints, and avoid redundant solutions during the search process. When the DiRect module is removed, the performance of the `LLM_GA` variant begins to deteriorate, especially in higher-RIS configurations. Although the LLM still provides effective initialization, insufficient semantic diversity in the form of structurally similar RIS configurations leads to premature saturation and restricts the exploration of distinct and physically feasible solutions. This observation suggests that maintaining sufficient diversity becomes increasingly important as the number of RISs grows, due to the larger and more complex solution space. The `LLM` method also shows relatively fast early-stage convergence, but its inability to revise geometrically infeasible or suboptimal placements results in lower final performance, reinforcing the necessity of feedback-driven refinement in realistic deployment scenarios.

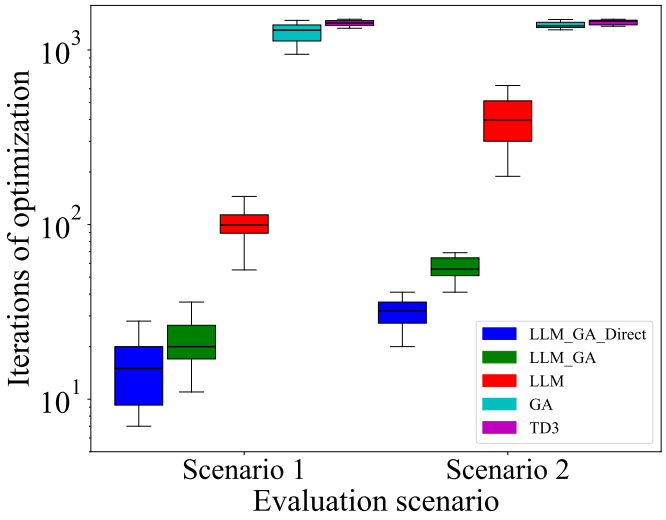

Fig. 7. Comparison of iteration counts required by all methods to reach the coverage threshold under two warehouse layouts: Scenario 1 (30×40 m^2) and Scenario 2 (60×80 m^2).

In contrast, traditional optimisation methods like `GA` and `TD3` exhibit significantly slower and more unstable convergence dynamics. The genetic algorithm suffers from inefficient search due to random initialization and limited guidance from environmental semantics, which becomes more pronounced in larger deployment spaces with complex blockage constraints. Meanwhile, the TD3-based deep reinforcement learning agent struggles with policy fluctuation across episodes and poor final results. This can be explained by the difficulty of applying model-free learning in combinatorial settings where actions must satisfy strict geometric feasibility, and where meaningful rewards are sparse and delayed. The `RANDOM` baseline consistently yields the lowest coverage and shows no improvement over iterations, serving as a lower bound for comparison.

E. Iteration Efficiency Analysis

Iteration counts required to reach the coverage threshold for different optimisation methods. Fig. 7 presents the statistical distribution of iteration counts required by each optimisation

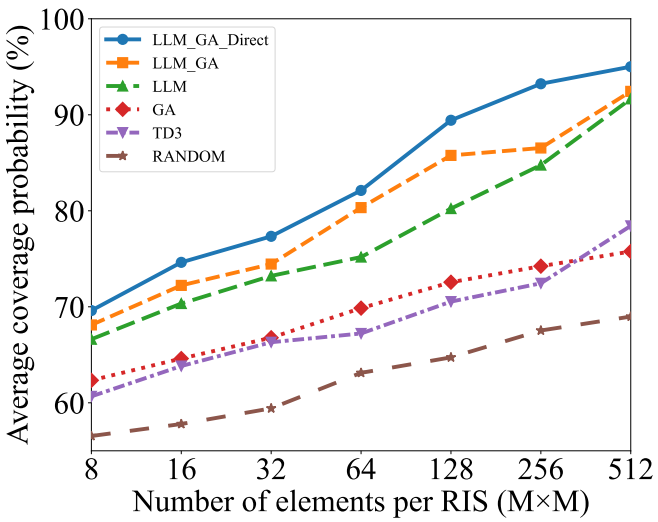

Fig. 8. Impact of the number of elements per RIS on average signal coverage under fixed RIS quantity, illustrating each method's ability to translate beamforming resolution into effective coverage.

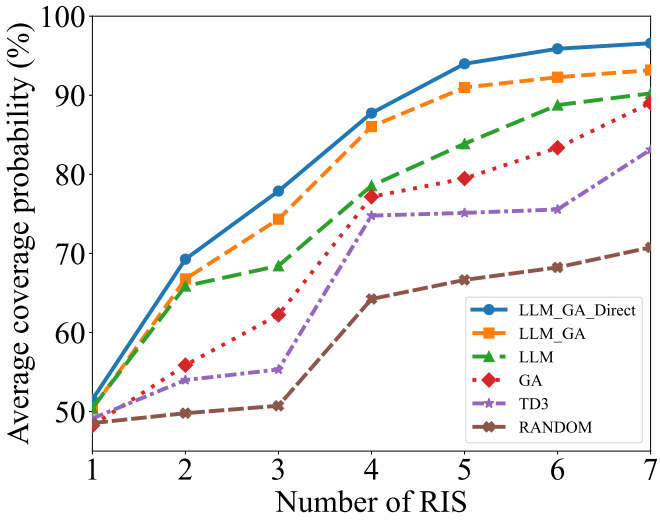

Fig. 9. Average coverage probability under different numbers of deployed RISs in a 60×80 m^2 warehouse, illustrating the scalability and effectiveness of various optimisation strategies as deployment density increases.

method to reach a predefined signal coverage threshold, evaluated across two distinct warehouse environments. Compared with average convergence curves, the boxplot provides a more complete statistical view, showing the median, spread, and outliers. In both environments, the proposed `LLM_GA_DiRect` not only achieves the lowest median iteration but also maintains the narrowest spread across runs, highlighting its strong convergence stability. Specifically, in the large-scale layout, it converges in 10–20 iterations for nearly all runs, with minimal variance. This indicates the method's ability to generate consistently high-quality proposals and refine them efficiently, regardless of environment complexity. By contrast, `LLM_GA` becomes more sensitive to the quality of initial solutions. In particular, the spread of results becomes noticeably wider in the smaller layout, indicating that some runs are slowed down by repeated or similar configurations. This suggests that insufficient diversity can negatively impact the effectiveness of the search process. The `LLM` exhibits even more pronounced variance. Although some runs converge quickly due to lucky initial placements, many require over 500 iterations, revealing the limitations of relying on static one-shot planning in complex spatial environments with geometric constraints.

Traditional methods, `GA` and `TD3` perform significantly worse in both environments. Their distributions are not only heavily skewed but also contain numerous extreme outliers. Most runs require over 1300 iterations, and the variation is large. `TD3` algorithm also shows high fluctuation, and often hits the maximum allowed number of iterations. This is likely because reinforcement learning methods have trouble dealing with environments where actions must follow strict geometric rules, and rewards are sparse or delayed.

F. Impact of RIS Elements on Coverage Performance

Fig. 8 shows the coverage performance as the number of elements per RIS increases, with the total number of RISs fixed. In all cases, coverage improves as the element size grows, reflecting the general benefit of higher beamforming resolution and reflective gain. The proposed `LLM_GA_DiRect` consistently outperforms all baselines throughout the range. It shows a particularly sharp improvement beyond the 64×64 configuration, demonstrating its ability to fully leverage large RIS arrays through coordinated spatial deployment and effective phase alignment. The integration of semantic planning, local refinement, and diversity correction ensures that high-capacity RISs are accurately directed toward weak-signal regions, avoiding misalignment or redundancy. In comparison, `LLM_GA` and `LLM` yield moderate gains. The performance gap between them widens as the element count increases, indicating that iterative refinement plays a key role in translating RIS capacity into actual coverage. Without diversity correction, `LLM_GA` tends to revisit similar configurations, limiting its ability to adjust beam directions in more complex settings.

Conventional approaches such as `GA`, `TD3`, and `RANDOM` perform noticeably worse. Their improvements are marginal and tend to saturate early, suggesting a lack of spatial adaptiveness and phase control. Despite having access to stronger hardware, these methods are unable to generate meaningful beamforming strategies, highlighting the importance of coupling algorithmic design with physical-layer capability.

G. Impact of RIS Quantity on Average Coverage

Fig. 9 shows the average coverage as the number of deployed RISs increases from one to seven. Coverage generally improves with more RISs, but the growth rate varies at different RIS counts. These variations arise because adding or removing a single RIS changes the set of available LoS and reflected paths for all other RISs. Such geometric interactions produce discrete shifts in achievable coverage under different deployment scales.

The proposed `LLM_GA_DiRect` method achieves the highest coverage across all RIS quantities. Its improvement is steady and pronounced because each RIS is placed in a

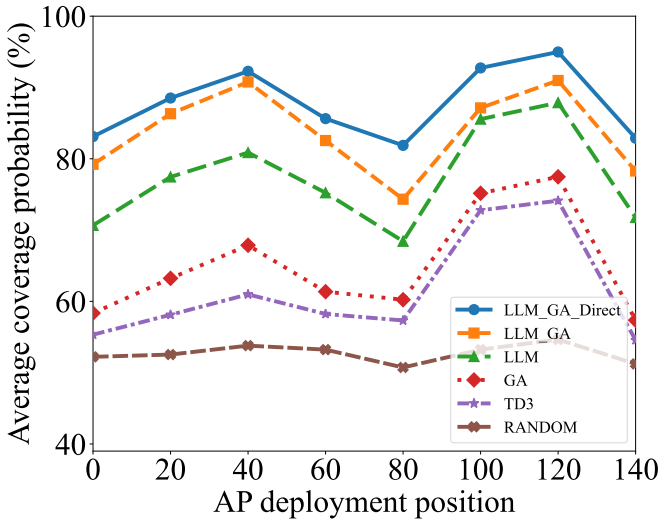

Fig. 10. Impact of BS deployment position on the average coverage across multiple RIS deployment in a 60×80 m^2 warehouse, illustrating how initial signal conditions affect multi-RIS optimisation effectiveness.

complementary region rather than a redundant one. Semantic initialization provides meaningful starting positions, GA performs local refinement, and the DiRect module maintains structural diversity so that each additional RIS contributes new useful paths rather than overlapping with existing ones. The `LLM_GA` and `LLM` variants also benefit from more RISs but exhibit smaller gains. `LLM` often places RISs in reasonable regions but cannot correct misalignments or redundant placements when the number of units increases. `LLM_GA` improves these placements but tends to converge to similar configurations without explicit diversity control, which reduces the benefit of adding more RISs.

Traditional methods such as `GA`, `TD3`, and `RANDOM` show limited scalability. Their performance rises slowly and often plateaus once four or five RISs are deployed. These methods struggle to explore the exponentially growing combinatorial search space, so additional RISs do not translate into meaningful new coverage paths. The `RANDOM` baseline remains consistently low, indicating that unguided placement cannot leverage the additional spatial degrees of freedom.

H. Impact of AP Deployment Position on Average Coverage

Fig. 10 presents the average coverage at eight BS locations placed along the warehouse perimeter. Coverage varies across these positions because each location produces a distinct LoS and blockage topology. Moving the BS from one discrete position to another can change which shelves become dominant blockers and which reflection paths become available, resulting in noticeable changes in achievable coverage. For each BS location, coverage is averaged over deployments with 2, 4, and 6 RISs to assess robustness under different deployment scales. This design highlights how each strategy adapts to different initial signal conditions rather than being tied to a single RIS density.

The proposed `LLM_GA_DiRect` method delivers the highest coverage across all BS positions. It adapts well even when the BS is placed near central shelving areas such as $(0, 40)$, $(40, 0)$, or $(60, 0)$ where strong blockage creates challenging NLoS conditions. Semantic initialization identifies feasible regions, GA performs structured local refinement, and the DiRect module maintains diversity to avoid collapsing to repetitive layouts. This combination enables the method to adjust RIS placements to each BS position and exploit new reflection opportunities. `LLM` and `LLM_GA` show more fluctuations across BS locations. `LLM` can perform well under favourable BS positions but struggles in heavily obstructed regions because it lacks iterative correction. `LLM_GA` improves robustness through local refinement, but without diversity preservation it may converge to similar solutions when the initial semantic proposal is weak.

Traditional methods such as `GA`, `TD3`, and `RANDOM` remain consistently inferior and show limited sensitivity to BS placement. Without semantic modeling or diversity-aware search, these baselines fail to exploit favourable AP positions and cannot overcome severe blockage when the BS is placed in challenging regions.

V. Conclusion

In this paper, we have proposed a novel hybrid optimisation framework that jointly optimises the placement and beamforming of multiple RISs to maximize coverage in smart warehouse environments. The framework integrates an LLM-driven semantic planner, an evolutionary mutation module, and a DiRect mechanism. Through extensive simulations, we evaluated the convergence, iteration efficiency, number of RISs and RIS elements, as well as the adaptability to different BS locations for all methods. The proposed method consistently outperformed traditional heuristics, reinforcement learning approaches, and LLM-guided baselines. Notably, it converged rapidly to high coverage even in high-dimensional search spaces, and maintained stable improvements as RIS quantity or element size increased. Moreover, the DiRect module played a critical role in enabling exploration across distinct deployment structures, highlighting its contribution to robust performance across varying network and layout configurations. These results have collectively demonstrated that coupling semantic reasoning with physically grounded refinement and diversity preservation offers a powerful and scalable solution for RIS-assisted network deployment in large-scale industrial scenarios. Furthermore, the observed sensitivity of coverage performance to the BS placement underscores the potential benefit of jointly optimising BS and RIS configurations. Thus, an important research direction lies in extending the proposed framework to coordinated BS–RIS co-deployment with incremental or predictive reconfiguration mechanisms, enabling the system to adapt to moderately dynamic warehouse environments while preserving the semantic planning advantages demonstrated in this study.

References

[1] L. Bassi, "Industry 4.0: Hope, hype or revolution?," in *Proc. IEEE 3rd Int. Forum Res. Technol. Soc. Ind.*, 2017, pp. 1–6.

[2] J. Cañete-Martin, *et al.*, "Integration of 5G and industrial digital models: A case study with AGVs," *IEEE 29th Int. Conf. Emerg. Technol. Fact. Autom. ETFA*, Padova, Italy, 2024, pp. 1-4.

[3] Z. Zhang, *et al.*, "Collision-free route planning for multiple AGVs in an automated warehouse based on collision classification," in *IEEE Access*, vol. 6, 2018, pp. 26022-26035.

[4] M. Pan, *et al.*, "Discretized optimal AP deployment algorithm for indoor environment," in *2023 IEEE Int. Conf. iThings and IEEE GreenCom and IEEE CPSCom and IEEE SmartData and IEEE Cybermatics*, Danzhou, China, 2023, pp. 137-142.

[5] V. Gadiraju, *et al.*, "Novel indoor sensor/access-point placement for full line-of-sight (LoS) coverage," in *IEEE Sens. J.*, vol. 25, no. 6, pp. 10218-10232, Mar. 2025.

[6] M. Di Renzo, *et al.*, "Smart radio environments empowered by reconfigurable intelligent surfaces: How it works, state of research, and the road ahead," in *IEEE J. Sel. Areas Commun.*, vol. 38, no. 11, pp. 2450-2525, Nov. 2020.

[7] Z. Li, *et al.*, "Enhancing indoor mmwave wireless coverage: Small-cell densification or reconfigurable intelligent surfaces deployment?," in *IEEE Wirel. Commun. Lett.*, vol. 10, no. 11, pp. 2547-2551, Nov. 2021.

[8] Y. Ren, *et al.*, "On deployment position of RIS in wireless communication systems: Analysis and experimental results," in *IEEE Wirel. Commun. Lett.*, vol. 12, no. 10, pp. 1756-1760, Oct. 2023.

[9] J. Huang, *et al.*, "A novel ray tracing based 6G RIS wireless channel model and RIS deployment studies in indoor scenarios," in *Proc. IEEE 33rd Annu. Int. Symp. PIMRC.*, Kyoto, Japan, 2022, pp. 884-889.

[10] M. Issa, *et al.*, "Using reflective intelligent surfaces for indoor scenarios: Channel modeling and RIS placement," in *Proc. 17th Int. Conf. Wireless Mobile Comput., Netw. Commun.*, Bologna, Italy, 2021, pp. 277-282.

[11] D. Kong, *et al.*, "Base-station and RIS deployment optimisation for indoor coverage enhancement," in *IEEE CAMA*, Genoa, Italy, Dec. 2023, pp. 246-249.

[12] J. Zhang, *et al.*, "Optimising coverage with intelligent surfaces for indoor mmwave networks," in *IEEE INFOCOM 2022*, London, United Kingdom, Jun. 2022, pp. 830-839.

[13] A. Albanese, *et al.*, "RIS-aware indoor network planning: The Rennes railway station case," in *ICC 2022 - IEEE International Conference on Communications*, Seoul, Korea, Republic of, Aug. 2022, pp. 2028-2034.

[14] G. Encinas-Lago, *et al.*, "A cost-effective RISs deployment to abate the coverage problem in B5G networks," in *IEEE Trans. Wirel. Commun.*, vol. 23, no. 10, pp. 15276-15290, Oct. 2024.

[15] G. O. Boateng et al., "A survey on large language models for communication, network, and service management: Application insights, challenges, and future directions," in *IEEE Communications Surveys & Tutorials*, Apr. 2025.

[16] H. Zhang, *et al.*, "Large language models in wireless application design: In-context learning-enhanced automatic network intrusion detection," in *GLOBECOM 2024 - 2024 IEEE Global Communications Conference*, Cape Town, South Africa, Mar. 2024, pp. 2479-2484.

[17] K. Qiu, *et al.*, "Large language model-based wireless network design," in *IEEE Wirel. Commun. Lett.*, vol. 13, no. 12, pp. 3340-3344, Dec. 2024.

[18] J. Hou, *et al.*, "Wireless-friendly window position optimisation for RIS-aided outdoor-to-indoor networks based on multi-modal large language model," *arXiv preprint arXiv:2410.20691*, 2024.

[19] N. Pinel, *et al.*, "Degree of roughness of rough layers: Extensions of the Rayleigh roughness criterion," *Progress In Electromagnetics Research B*, vol. 26, pp. 213–250, 2010.

[20] Z. Yun, *et al.*, "Ray tracing for radio propagation modeling: Principles and applications," in *IEEE Access*, vol. 3, pp. 1089-1100, Jul. 2015.

[21] C. Cai, *et al.*, "Hierarchical passive beamforming for reconfigurable intelligent surface aided communications," in *IEEE Wirel. Commun. Lett.*, vol. 10, no. 9, pp. 1909-1913, Sept. 2021.

[22] S. Sandh, *et al.*, "Ray tracing algorithm for reconfigurable intelligent surfaces," in *arXiv preprint arXiv:2402.13034*, 2024.

[23] Y. Sun, *et al.*, "A 3D non-stationary channel model for 6G wireless systems employing intelligent reflecting surfaces with practical phase shifts," in *IEEE Trans. Cogn. Commun. Netw.*, vol. 7, no. 2, pp. 496-510, Jun. 2021.

[24] S. Kurt, *et al.*, "Path-loss modeling for wireless sensor networks: A review of models and comparative evaluations," in *IEEE Antennas Propag. Mag.*, vol. 59, no. 1, pp. 18-37, Feb. 2017.

[25] Z. Zhou, *et al.*, "A survey on efficient inference for large language models," *arXiv preprint arXiv:2404.14294*, 2024.