

A Modular Prompting Framework for AI-Guided Unobtrusive User Engagement in Crowd-Powered Indoor Positioning with Semantic Knowledge Graphs and GPT-4-Based LLMs^{*}

Ahmed Mansour^{1,2}, Wu Chen¹, Mahmoud Adham^{1,2} and Huan Luo¹

¹Department of Land Surveying and Geo-Informatics, The Hong Kong Polytechnic University, Hong Kong

²Public Works Department, Faculty of Engineering, Cairo University, Giza, Egypt

Abstract

Crowd-powered indoor positioning systems (IPS) offer scalable and cost-effective solutions for building and updating radio maps. However, they face challenges including limited calibration resources, unreliable signal labeling, and incomplete semantic annotations. Passive data collection reduces user burden but often lacks sufficient contextual cues to ensure localization accuracy. Auxiliary aids such as GNSS, QR codes, or BLE beacons can support calibration but are confined to specific deployment zones, and GNSS is ineffective in many indoor environments. Active user engagement, by contrast, can provide scalable annotation and calibration if prompts are delivered in a timely, context-aware manner and crafted sensitively to user preferences and cognitive load, so as to avoid fatigue and disengagement. This work is motivated by unobtrusive engagement strategies employed in platforms like Google Maps, where users are prompted to report transit conditions during navigation, and by similar feedback mechanisms on YouTube and Facebook that refine recommendation algorithms. We propose the first modular, AI-guided prompting framework for unobtrusive spatial feedback collection in crowdsourced IPS, enabling users to confirm location estimates, floor levels, and Points of Interest (POI) names without disrupting primary tasks. The framework comprises five interoperable layers: contextual situation assessment, intelligent prompt selection, hybrid prompt generation, user interaction handling, and feedback integration through continuous learning. Spatial knowledge graphs (SKGs) are introduced to embed semantic context into prompt logic, while large language models (LLMs) such as GPT-4 generate linguistically optimized, user-specific queries. By selectively prompting users at opportune moments, the system transforms sporadic passive data into semantically rich, trustworthy inputs with minimal disruption.

Keywords

Crowdsourced indoor positioning, user engagement, prompt generation, radio map construction, semantic annotation, spatial knowledge graph, large language models (LLM), GPT-4, unobtrusive interaction

1. Introduction

Indoor positioning systems (IPS) have emerged as a critical technology enabling numerous applications ranging from indoor navigation and location-based services to emergency response and smart building management. Traditionally, accurate IPS solutions—particularly fingerprinting-based approaches—have relied heavily on meticulous site surveys, extensive infrastructure deployment, and manual calibration processes [1, 2]. These methods, while precise, are costly, labor-intensive, and impractical for large-scale or dynamically changing environments. Consequently, crowdsourced approaches leveraging smartphones' ubiquity and integrated sensors (Wi-Fi, magnetometers, barometers, etc.) have gained traction in recent years. By capitalizing on crowd-powered data collection, IPS can substantially reduce costs, scale more effectively, and adapt dynamically to environmental changes. However, despite these advantages, crowdsourced data introduces new challenges related to data reliability, user compliance, and annotation accuracy. One fundamental issue with crowdsourced IPS data is the lack of controlled conditions for data collection [3]. Users typically contribute passively, unaware of or uninterested in

IPIN-WCAL 2025: Workshop for Computing & Advanced Localization at the Fifteenth International Conference on Indoor Positioning and Indoor Navigation, September 15–18, 2025, Tampere, Finland

✉ ahmed.m.mostafa@connect.polyu.hk (A. Mansour)

ORCID 0000-0002-9840-7030 (A. Mansour)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

ensuring the data’s accuracy or completeness. This often leads to noisy, inconsistent, or incomplete datasets that undermine the reliability of the resulting positioning models. Moreover, the absence of semantic annotations—such as precise floor identification, building labels, or POI tags—complicates an IPS’s ability to generate rich, contextually meaningful maps. As a result, IPS built solely on passive crowdsourced data can exhibit substantial localization errors and uncertainty [4]. To mitigate these challenges, auxiliary aids such as GNSS signals, QR codes, or BLE beacons have been suggested to support calibration in passive data collection approaches. These aids, however, are confined to specific deployment zones, and GNSS remains ineffective in many indoor or underground areas. Recent research in other domains has highlighted the need for active user engagement to improve crowdsourced data quality. Traditional methods of actively soliciting user input (e.g., periodic surveys or pop-up queries) often face resistance due to inconvenience or fatigue, resulting in low response rates and limited scalability. The critical problem, therefore, is balancing the need for accurate, reliable, and contextually rich data against the requirement of minimal user disruption. Achieving this balance necessitates innovative solutions inspired by successful digital platforms that seamlessly integrate unobtrusive interactions into user workflows, fostering high engagement without negatively affecting user experience.

Platforms like YouTube, Facebook, and Google Maps have demonstrated that brief, contextually relevant prompts can elicit significant user engagement with minimal disruption. YouTube’s context-sensitive video recommendations and prompts effectively guide user interaction, enhancing both satisfaction and platform metrics. Similarly, Facebook employs brief, unobtrusive questions directly integrated into users’ browsing flows, enabling effortless user feedback without interrupting primary activities. In navigation scenarios, Google Maps’ implementation of simple prompts—for instance, asking users to assess transit crowdedness during a journey—has proven effective in collecting real-time user-generated data while maintaining a smooth navigation experience. Indeed, unobtrusive user engagement has driven valuable crowd inputs in domains like traffic and transit monitoring [5].

Unlike these well-studied domains, the context of crowd-powered IPS has lacked dedicated studies on unobtrusive user engagement strategies. To fill this gap, and motivated by the effectiveness of such strategies elsewhere, this paper proposes a structured, AI-driven framework tailored for IPS data quality enhancement through user feedback that minimizes disruption. The proposed framework is designed to seamlessly engage users in contributing spatial feedback—such as confirming location estimates, floor labels, or POI names—to support accurate indoor radio map construction and semantic annotations.

Our framework comprises five interoperable layers: (1) contextual situation assessment, (2) intelligent prompt strategy selection, (3) prompt generation, (4) user interaction and feedback capture, and (5) feedback integration with continuous learning. It balances system requirements with user attention by adapting to motion state, data uncertainty, and individual user preferences. Spatial Knowledge Graphs (SKGs) are introduced at the core of the system’s intelligence to embed semantic awareness into prompt selection and generation. Additionally, large language models (LLMs) such as GPT-4 are leveraged to generate linguistically optimized and context-sensitive prompts tailored to each user. By selectively engaging users at opportune moments and in context-appropriate ways, the system transforms sporadic passive data into semantically rich, trustworthy inputs with minimal disruption.

This paper presents the design of the AI-driven prompting framework and discusses each of its components in detail. We also report on a preliminary evaluation through simulation, demonstrating the framework’s potential to improve user engagement and data quality. Future work will extend this research with full system implementation on real devices and empirical user studies. The rest of this paper is organized as follows: Section 2 describes the proposed framework architecture and methodology, detailing the functionality of each layer. Section 3 presents an evaluation of the framework’s performance in simulated scenarios, including user engagement rates and map improvement metrics. Finally, Section 4 concludes the paper and outlines directions for future work.

2. Proposed Framework and Methodology

Figure 1 provides an overview of the proposed AI-guided prompting framework for unobtrusive user engagement in crowd-powered IPS. The framework is designed as a modular pipeline of five layers that work together to decide when to prompt the user, what to ask, how to phrase the prompt, how to handle the user’s response, and how to integrate the feedback to improve the system continuously. In this section, we describe each layer in detail, along with the underlying methodology and algorithms enabling its functions.

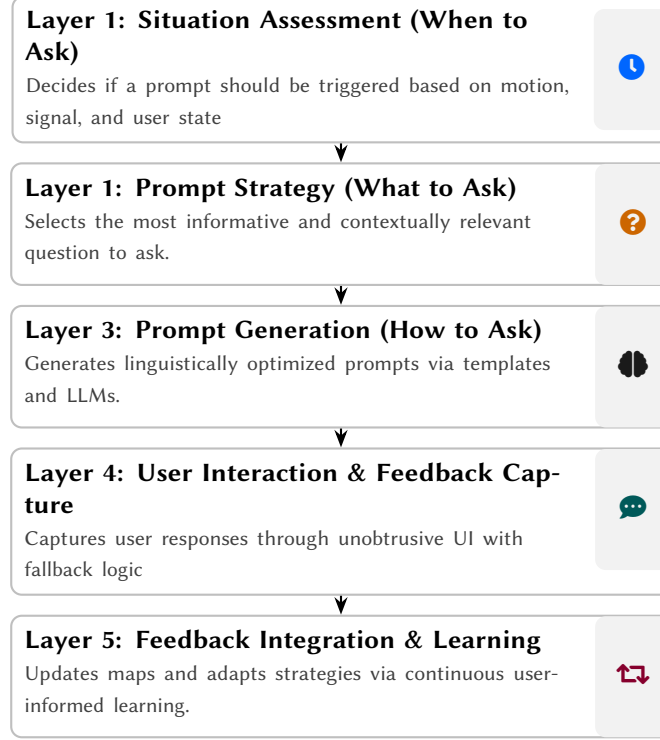


Figure 1: Modular pipeline for AI-driven prompt management in crowd-powered IPS, comprising five layers from context assessment to feedback learning.

2.1. Contextual Situation Assessment (When to Ask)

The first layer of the framework assesses whether a user should be prompted at a given moment. In crowd-powered IPS, indiscriminate prompting is counterproductive—it leads to user fatigue, reduced compliance, and potential disengagement. Thus, Layer 1 is pivotal in ensuring that prompts are issued only when necessary for improving the system and when they are unlikely to disrupt the user’s primary task. This situation assessment continuously monitors both the state of the system (e.g., localization uncertainty, map coverage) and the state of the user (motion, activity, history of prompts) to make intelligent decisions about prompting. The overall decision logic and key contextual inputs guiding this layer’s evaluations are illustrated in Figure 2.

As shown in Figure 2, the decision to trigger a prompt is informed by several key criteria. The first concerns *radio map coverage*, which refers to the availability (or lack) of existing signal fingerprints at the user’s estimated location. If the current location is under-represented in the radio map—such as when few or no prior Wi-Fi scans exist for this area—the system may prompt the user to contribute data or verify their position. The sufficiency of radio data at a given location l_i is quantified by a normalized signature count, defined as

$$R_{\text{sig}}(l_i) = \frac{|D(l_i)|}{|D_{\text{max}}|}, \quad (1)$$

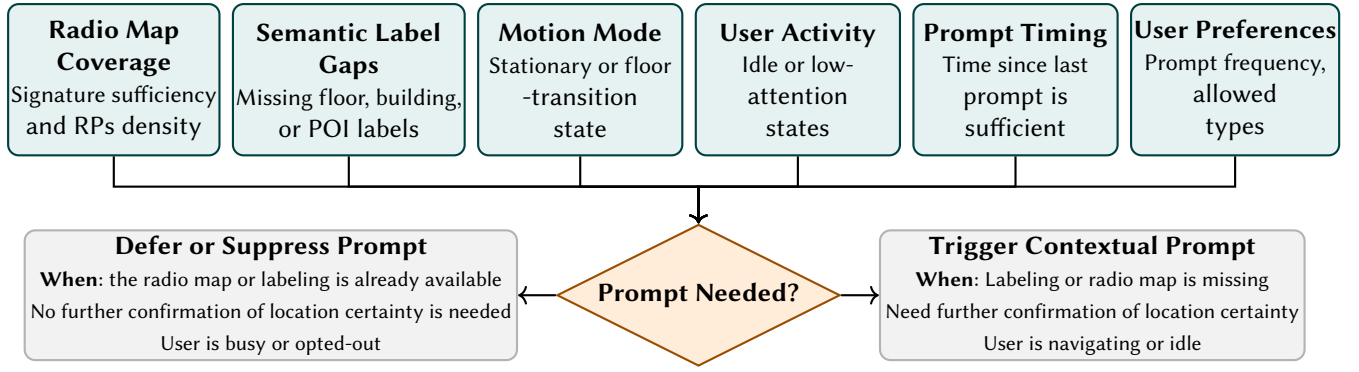


Figure 2: Decision logic for Layer 1: assessing when to prompt based on uncertainty, semantic gaps, and user context. Prompts occur only if needed; otherwise engagement is deferred.

where $|D(l_i)|$ denotes the number of sensor readings (for example, Wi-Fi RSS fingerprints) collected at l_i and $|D_{\max}|$ represents a reference value corresponding to well-surveyed locations. If $R_{\text{sig}}(l_i)$ falls below a threshold τ_R , indicating sparse data, the likelihood of triggering a prompt at this location increases. A second criterion is the presence of *semantic label gaps*, which evaluates the necessity or the likelihood of semantic annotations for the current estimated place. Missing labels, such as the floor number, building ID, or POI name, highlight opportunities where user input could meaningfully enhance the system’s understanding.

The system also accounts for the *user’s motion mode*, classified using device sensors into states such as stationary, walking, running, or undergoing an elevator or stair transition. Prompts are ideally delivered when the user is stationary, moving slowly, or immediately after completing a significant transition like a floor change. Active movements, especially rapid walking or navigation that demands the user’s attention, lead to prompt suppression. A binary condition formally captures this:

$$\text{Prompt}_{\text{motion}} = \begin{cases} \text{True}, & \text{if motion state} \in \{\text{stationary, just stopped, floor change}\}, \\ \text{False}, & \text{if motion state} \in \{\text{walking, running, driving}\}. \end{cases} \quad (2)$$

Only when $\text{Prompt}_{\text{motion}}$ evaluates to true does the system proceed to consider prompting the user. Beyond physical movement, the framework incorporates the *user activity state* by monitoring current device engagement through activity recognition APIs or analysis of the foreground application context. If the user is actively engaged in tasks that demand attention—such as being on a phone call, typing a message, or gaming—the prompt is deferred to avoid disruption. This decision uses predefined sets: $\mathcal{A}_{\text{promptable}}$, which includes activities conducive to prompting (such as using a maps app or being idle on the home screen), and $\mathcal{A}_{\text{busy}}$, which contains activities like video watching or extensive typing. A binary flag is then defined as

$$U_{\text{act}} = \begin{cases} 1, & \text{if } a_t \in \mathcal{A}_{\text{promptable}}, \\ 0, & \text{if } a_t \in \mathcal{A}_{\text{busy}}, \end{cases} \quad (3)$$

where the system will only consider triggering a prompt if $U_{\text{act}} = 1$ at that particular moment. Moreover, *prompt timing and historical context* are integral to prevent overwhelming the user. Letting t_{last} denote the timestamp of the most recent prompt shown or answered by the user, a new prompt is only generated if the current time t_{now} satisfies $t_{\text{now}} - t_{\text{last}} > \tau_T$, where τ_T is a configurable timeout parameter, typically set to several minutes or hours depending on the application’s requirements. This spacing ensures that prompts are not clustered too closely together. Additionally, users may have personalized settings that limit the number of prompts per day. If a user specifies—or the system infers—a preference for at most N prompts daily, the framework respects this by enforcing an additional gating mechanism that suppresses further prompts once this limit is reached.

The system also adapts based on the *user’s engagement history*. For users who consistently ignore or dismiss prompts of a certain type, the model deprioritizes those prompts or raises the required

significance threshold (such as uncertainty or semantic gaps) before issuing them again. Conversely, if a user reliably responds to specific kinds of queries, the system may preferentially select these when appropriate. This layered decision-making process ensures that prompt delivery aligns with moments that are both critical for system information and appropriate for the user, laying a solid foundation for unobtrusive engagement within Layer 1 of the framework.

2.2. Prompt Strategy Engine (What to Ask)

Once the framework determines it is an opportune moment to prompt, Layer 2 selects the most relevant question to ask. In a crowd-powered IPS, some prompts (like confirming a floor or naming an unknown POI) can substantially improve the system, while redundant or irrelevant queries risk user annoyance. Thus, the engine combines rule-based logic, uncertainty estimation, and learning-based policies to prioritize prompts.

It first evaluates the *current uncertainty vector* across dimensions such as position, floor, or building, selecting the highest uncertainty component to target. It then checks for *missing semantic information*, for example prompting to label an “unknown” room. The Spatial Knowledge Graph (SKG) ensures contextual relevance by anchoring prompts to nearby known entities, preventing questions about unrelated floors or wings. The engine also integrates *historical trends and user profiles*, prioritizing queries in areas flagged by prior user corrections or tailoring prompt complexity to user expertise. Each candidate prompt a is scored by

$$U(a) = r(a) - \lambda c(a), \quad (4)$$

where $r(a)$ measures expected uncertainty reduction and $c(a)$ the user effort, balanced by λ . Prompts are ranked by $U(a)$ to maximize utility. Over time, a multi-armed bandit and Q-learning approach refine these choices based on observed user interactions, balancing exploration and exploitation. Redundancy filters further prevent asking for data that has already confirmed. This ensures every prompt contributes meaningfully to improving IPS coverage and accuracy.

2.3. Prompt Generation (How to Ask)

Layer 3 formulates how to phrase the selected question, aiming for clarity, brevity, and contextual appropriateness. Since prompt wording strongly affects user engagement, our framework combines predefined templates with adaptive large language models (LLMs) to generate natural, tailored queries. We employ three strategies (Figure 3):

- *Template-Based*: Direct patterns with placeholders, e.g., “Are you on Floor [X]?” ensure consistency and low computational cost, ideal for straightforward yes/no or multiple-choice queries.
- *LLM-Based*: In complex or nuanced contexts, the system feeds situational data into an LLM to generate polite, context-aware questions like “This area isn’t labeled—do you know what it’s used for?” enhancing engagement for open-ended prompts.
- *Hybrid*: Starts from a template but lets the LLM refine phrasing based on context, balancing reliability with conversational tone (e.g., after elevator use: “Just to confirm, are you now on Floor 5?”).

This selection follows:

$$\text{Prompt}_{\text{final}} = \begin{cases} \text{Template}(c), & \text{if simple context} \\ \text{LLM}(c), & \text{if complex or novel context} \\ \text{Hybrid}(\text{Template}, c), & \text{if needing nuance} \end{cases} \quad (5)$$

Localization ensures prompts use the user’s language and local terms (e.g., “ground floor” vs. “first floor”), while personalization tailors tone by learning user preferences—some get concise checks, others a friendly nudge. All LLM outputs undergo validation to avoid off-topic or verbose phrasing, falling back to templates if necessary. In sum, Layer 3 blends deterministic and AI-generated approaches to maximize user response quality, essential for high-fidelity data collection in the IPS.

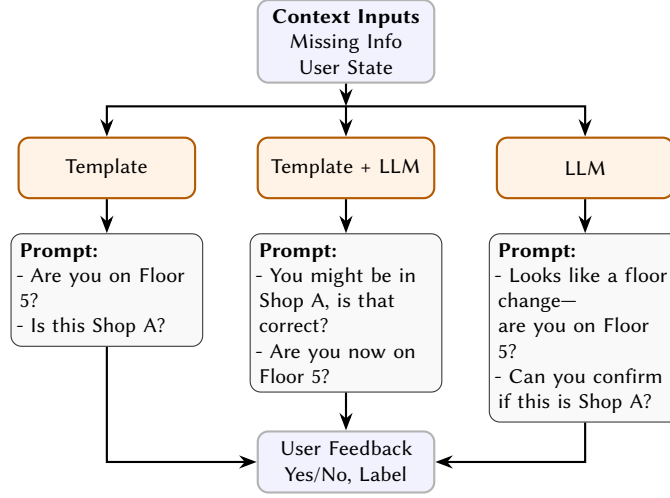


Figure 3: Three prompt generation strategies—template, hybrid, and LLM-based—are triggered from context and feed into user feedback collection.

2.4. Spatial Knowledge Graph Integration

A core innovation of our framework is the Spatial Knowledge Graph (SKG), which provides semantic context and consistency checks for prompting. The SKG is modeled as a directed multigraph

$$G = (V, E) \quad (6)$$

where each node $v \in V$ represents a spatial entity (room, floor, building), and each edge $(v_i \xrightarrow{r} v_j) \in E$ captures topological, containment, proximity, or functional relationships. This structure, initially derived from floor plans or BIM data, is dynamically refined as users contribute new labels and relations (see Layer 5).

The SKG ensures prompts are contextually valid. For example, if a user is in Room B on Floor 1 of Building X, the system knows nearby entities like Room C or the Library and might ask about these if unlabeled, while avoiding irrelevant queries about distant locations such as Office G on Floor 2 unless recent sensor data suggests a transition. This preserves user trust by preventing nonsensical prompts.

Formally, we compute a contextual relevance score

$$\phi(v_{\text{candidate}}, v_{\text{current}}) \quad (7)$$

that is high for short graph paths indicating close spatial or semantic proximity. The prompt strategy engine only considers candidates satisfying $\phi(v, v_{\text{current}}) > \tau_\phi$. This enables the system to prioritize questions about nearby or related entities. Moreover, the SKG aids in validating user feedback. If a user labels a location inconsistently with known floor numbering or spatial layout, the system detects such anomalies and may trigger follow-ups or delay integration. As users provide new place names or relations, the SKG expands, enhancing prompt selection and interpretation. Thus, the SKG acts as the semantic backbone of the framework, connecting individual user inputs into a coherent spatial model that guides intelligent prompting.

2.5. User Interaction, Feedback Capture, and Continuous Learning

Layer 4 ensures prompts are delivered unobtrusively through adaptive UI elements like overlays or notifications that adjust to user context—larger text or speech output if walking, detailed widgets when stationary—and favor simple one-tap inputs that achieved over 86% compliance in our evaluations. Ignored or dismissed prompts are treated as signals, with deferred retries governed by back-off to minimize fatigue. Multi-modal inputs are normalized and validated, with even partial or uncertain responses used as weak evidence. Feedback immediately updates operational models, such as radio

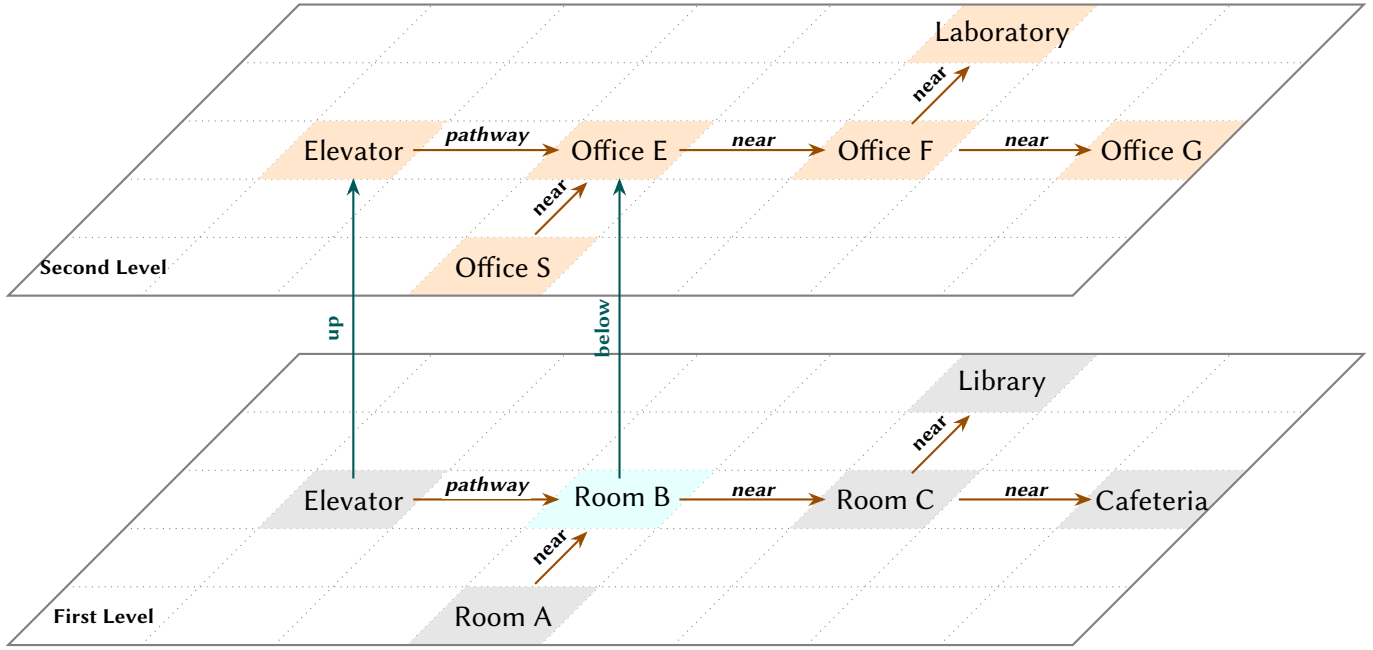


Figure 4: Multi-floor Spatial Knowledge Graph (SKG) for prompt filtering and expansion. When the user is located at Room B, prompts associated with adjacent Room C or nearby POIs, such as the Library, are semantically relevant. SKGs help suppress prompts about distant or disconnected locations.

maps and SKGs, while also populating engagement logs that inform Layer 5’s adaptive strategy. Here, reinforcement learning and bandit algorithms dynamically adjust prompt types, timing, and phrasing, balancing information gain against user burden. Validated inputs calibrate barometric floor detection, refine semantic labels, and strengthen or adjust fingerprints, all weighted to mitigate erroneous or malicious data. This integrated approach establishes a sustainable loop where user interactions progressively enhance IPS accuracy and personalization, improving both functional reliability and user experience over time.

2.6. Experimental Setup

Because deploying a live system to a large user base was beyond the scope of this initial study, we built a simulation environment that models user movement, sensor readings, and user behavior in response to prompts. The simulated environment comprised a two-story building with 20 distinct locations of interest (rooms, corridors, POIs), some fully labeled and some initially unknown. A set of 50 virtual users was generated, each with a profile of responsiveness and movement patterns. *Environment Model:* Each floor of the building was represented by a grid with certain nodes designated as named locations (e.g., Lab, Office, Elevator area). A Wi-Fi radio map was synthetically generated: certain grid points had associated Wi-Fi fingerprints, and signal strengths were perturbed with noise. Initially, about 60% of the grid had sufficient fingerprint data; the rest were “sparse” areas (to test how the system handles low coverage). Semantic labels for about half of the points of interest were provided, while others started as unknown to simulate missing information. We also defined a spatial knowledge graph for the building (similar to Figure 4), encoding which rooms were connected or adjacent. *Simulated User Trajectories:* Each user was assigned random start points and movement patterns. Some followed regular routes (e.g., repeatedly going from the entrance to a particular office), while others roamed more randomly. We introduced events like elevator usage (to test floor transitions) and pauses (to simulate stopping in a hallway or room for a while). The simulation ticked in discrete time-steps, and at each step, each user could either move to a neighboring cell, stay still, or change activity (some users “checked their phone” at certain intervals, etc.). *User Prompt Response Model:* We modeled user behavior in a probabilistic manner. Each user had a base willingness to respond to prompts (ranging from 20% to 90% chance to

respond when prompted, reflecting different engagement levels). This probability was dynamically adjusted by factors such as current activity (if the user was “busy” in the simulation, response probability dropped near 0), prompt frequency (if they had been prompted recently, probability dropped due to fatigue), and prompt type (we assumed users are more likely to respond to easier prompts: yes/no > multiple choice > text input). If a user decided to respond, the response content was generated based on ground truth with some chance of error (e.g., 5% chance they hit the wrong button or gave a wrong label by mistake, adding noise).

2.7. Prompt Triggering Performance

We evaluated the framework’s ability to trigger prompts precisely when needed, guided by the Layer 1 situation assessment that ensures prompts are issued at the right locations and times—specifically where data gaps exist and the user context is appropriate. In our simulation, the adaptive system generated only 420 prompts, compared to 1200 by the naive baseline, yet it effectively targeted approximately 75% of critical missing information opportunities. By contrast, the baseline covered only about 50%, often wasting prompts at irrelevant or poorly timed moments. These findings underscore the importance of Layer 1’s context-sensitive logic in delivering prompts when they are most likely to yield informative, non-intrusive responses. Future real-world trials will be essential to measure actual response rates and verify that this context-driven approach sustains high engagement and data quality beyond simulation.

2.8. Effect of the Spatial Knowledge Graph on Prompt Generation

The integration of the Spatial Knowledge Graph (SKG) proved critical in enhancing both the relevance and clarity of generated prompts. By encoding topological, containment, and proximity relationships among spatial entities—such as rooms, corridors, floors, and POIs—the SKG ensures that prompts are tightly coupled to the user’s immediate context. As shown in Table 1, without SKG support, the system might issue generic or even spatially inconsistent queries (for example, asking “Is this Room 305?” when the user is actually on Floor 2 where only 200-series rooms exist). In contrast, SKG-informed prompts leverage the graph to filter candidate questions to those that make semantic sense, such as distinguishing between Office 201 or 202 when the user is known to be on Floor 2. Moreover, the SKG enables more natural and persuasive phrasing by incorporating nearby landmarks into the question. For instance, instead of a blunt “What is the name of this room?”, the system can ask “Near the Library—could you tell us what this adjacent room is called?”, which not only improves clarity but also signals to the user that the system understands the environment. This contextual anchoring was also instrumental in validating user feedback; when a response conflicted with known SKG structures (such as labeling a space “Room 101” on a floor without 100-series rooms), the system could gracefully initiate follow-up checks. Overall, the SKG improved prompt targeting by reducing irrelevant or confusing questions, increased linguistic adaptability by embedding local context into phrasing, and thereby directly contributed to higher response quality and user trust.

3. Conclusion

In this paper, we proposed the first framework for unobtrusive active user intervention in crowd-powered indoor positioning systems through intelligent prompt generation. Our innovative approach combines large language models (LLMs) and a Spatial Knowledge Graph (SKG) to dynamically determine not only what to ask users, but also how and when to do so—ensuring prompts are contextually relevant, timely, minimally disruptive, and ultimately user-friendly. The framework explicitly incorporates multiple essential aspects, including spatial semantics for consistency, adaptive language for natural engagement, fatigue-aware pacing, and personalized prompting strategies, all designed to maximize the quality and quantity of user contributions without imposing undue burden. To preliminarily evaluate our approach, we conducted a set of simulation experiments that demonstrated how the framework effectively triggers prompts at appropriate moments and locations, while highlighting the added value

Table 1

Examples of prompts generated with and without SKG integration.

Situation	Without SKG	With SKG
User near Library entrance, unlabeled adjacent room	“What is the name of this room?”	“Near the Library—could you tell us what this adjacent room is called?”
User on Floor 2 in Office cluster, but ambiguous about exact office	“Is this Room 305?” (possible mismatch; Room 305 may be on another floor)	“In this office area on Floor 2, is this Office 201 or 202?” (filtered by SKG floor associations)
User standing in a corridor next to a Cafeteria	“What type of place is this?”	“Next to the Cafeteria—would you classify this space as a corridor or seating area?”
User feedback says “Room 101” on a floor without 100-series rooms	Accepts the label without question, risking inconsistency	Follows up: “We usually see Room 101 on Floor 1—could you check the floor number?”
User stopped at an unlabeled POI on Floor 3	“Please provide a label for this place.”	“On Floor 3 near the elevators, do you know what this place is called or used for?”

of the SKG in guiding both prompt selection and phrasing. These early results validate the potential of our design to achieve higher engagement rates and more targeted data collection compared to naive prompting methods. Looking ahead, this work lays the foundation for an ongoing and more extensive investigation. We plan to deploy the framework in real-world environments and perform rigorous assessments of its impact on user response rates, positioning accuracy, and overall system improvement, complemented by user surveys to gauge subjective experience. Through this line of research, we aim to establish a robust pathway for leveraging intelligent, human-in-the-loop interactions to continuously refine and enhance indoor positioning systems at scale.

Declaration on Generative AI

During the preparation of this work, the author(s) employed ChatGPT-4o and Grammarly to improve writing quality, grammar, and language, as well as for general proofreading. All ideas, methodologies, and architectural frameworks have been independently developed and proposed by the author(s). After utilizing these tools, the author(s) carefully reviewed and refined the content and assume full responsibility for the integrity and accuracy of the publication.

References

- [1] A. Pérez-Navarro, J. Torres-Sospedra, R. Montoliu, J. Conesa, R. Berkvens, G. Caso, C. Costa, N. Dorigatti, N. Hernández, S. Knauth, et al., Challenges of fingerprinting in indoor positioning and navigation, in: *Geographical and Fingerprinting Data to Create Systems for Indoor Positioning and Indoor/Outdoor Navigation*, Elsevier, 2019, pp. 1–20.
- [2] A. Mansour, W. Chen, Suns: A user-friendly scheme for seamless and ubiquitous navigation based on an enhanced indoor-outdoor environmental awareness approach, *Remote Sensing* 14 (2022) 5263. URL: <https://doi.org/10.3390/rs14205263>.
- [3] E. S. Lohan, J. Torres-Sospedra, H. Leppäkoski, P. Richter, Z. Peng, J. Huerta, Wi-fi crowdsourced fingerprinting dataset for indoor positioning, *Data* 2 (2017) 32.
- [4] A. Mansour, J. Ye, Y. Li, H. Luo, J. Wang, D. Weng, W. Chen, Everywhere: A framework for ubiquitous indoor localization, *IEEE Internet of Things Journal* 10 (2023) 5095–5113. URL: <https://doi.org/10.1109/JIOT.2022.3222003>.
- [5] A. Ali, N. Ayub, M. Shiraz, N. Ullah, A. Gani, M. A. Qureshi, Traffic efficiency models for urban traffic management using mobile crowd sensing: A survey, *Sustainability* 13 (2021) 13068.