

Tightly Coupled SLAM with Imprecise Architectural Plans

Muhammad Shaheer¹, Jose Andres Millan-Romera¹, Hriday Bavle¹, Marco Giberna¹,
Jose Luis Sanchez-Lopez¹, Javier Civera², and Holger Voos¹

Abstract— Robots navigating indoor environments often have access to architectural plans, which can serve as prior knowledge to enhance their localization and mapping capabilities. While some SLAM algorithms leverage these plans for global localization in real-world environments, they typically overlook a critical challenge: the “as-planned” architectural designs frequently deviate from the “as-built” real-world environments. To address this gap, we present a novel algorithm that tightly couples LIDAR-based simultaneous localization and mapping with architectural plans in the presence of deviations. Our method utilizes a multi-layered semantic representation to not only localize the robot, but also to estimate global alignment and structural deviations between “as-planned” and “as-built” environments in real-time. To validate our approach, we performed experiments in simulated and real datasets demonstrating robustness to structural deviations up to 35 cm and 15°. On average, our method achieves 43% less localization error than baselines in simulated environments, while in real environments, the “as-built” 3D maps show 7% lower average alignment error.

Paper Video: <https://www.youtube.com/watch?v=9O0qwNhTuqk>

Index Terms—SLAM, Localization, Robotics and Automation in Construction

I. INTRODUCTION

PRIOR information from architectural plans can be used to enhance the localization and mapping accuracy of mobile robots. However, real-world buildings rarely match their plans perfectly due to construction tolerances and modifications [1], [2]. Incorporating such imprecise prior information into robot navigation pipelines can introduce systematic errors, potentially damaging the localization and mapping accuracy, instead of helping to improve it. To address this issue, it is necessary to tightly couple the Simultaneous Localization and Mapping (SLAM) pipeline with architectural plans. This requires a unified representation of both the “as-planned” and “as-built”

Manuscript received: January, 21, 2025; Revised: April, 14, 2025; Accepted: June, 05, 2025.

This paper was recommended for publication by Editor Hyungpil Moon upon evaluation of the Associate Editor and Reviewers’ comments.

*This research was funded in whole, or in part, by the Luxembourg National Research Fund (FNR) under the projects 17097684/RoboSAUR and C22/IS/17387634/DEUS, by a partnership between the Interdisciplinary Center for Security Reliability and Trust (SnT) of the University of Luxembourg and Stugalux Construction S.A. and by the Spanish Government under Grants PID2021-127685NB-I00 and TED2021-131150B-I00. For the purpose of Open Access, the author has applied a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission

¹Authors are with the Automation and Robotics Research Group, Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg. Holger Voos is also associated with the Faculty of Science, Technology and Medicine, University of Luxembourg, Luxembourg. {muhammad.shaheer, jose.millan, hriday.bavle, marco.giberna, joseluis.sanchezlopez, holger.voos}@uni.lu

²Author is with I3A, Universidad de Zaragoza. jcivera@unizar.es

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE




Fig. 1: *diS-Graphs* overview. Our method couples a hierarchical SLAM factor graph built online by a robot, with an architectural plan (also modeled as a hierarchical factor graph) that may contain deviations, to form a deviations-informed Situational Graph (*diS-Graph*, center). The zoomed-in view (right) illustrates that this coupling enables the estimation of the rigid transformation p_S^A between both graphs and, additionally, the wall-surface and room deviations $d_{\pi_{s_i}}^{\rho_{a_j}}$ and $d_{\rho_{s_j}}^{\rho_{a_j}}$ respectively.

environments—one that captures not only the geometry but also architectural semantics. Such a representation will first enable the matching of structural elements between the two instances, and subsequently, the estimation of their deviations. Although there are various environment representation techniques such as *occupancy grids* [3], none of them explicitly models the semantic and hierarchical information of the environment, which is needed to identify the deviated structural elements. Recent approaches such as 3D scene graphs [4], [5] or Situational Graphs (*S-Graphs*) [6], [7], represent a robot’s environment in a compact and hierarchical manner, encoding high-level semantic abstractions (for example, walls and rooms) and their relationships (e.g., a set of walls forms a room). Herein, *S-Graphs* extend 3D scene graphs by merging geometric models of the environment generated by SLAM approaches with 3D scene graphs into a multi-layered jointly optimizable factor graph. This representation, combined with the prior information extracted from architectural plans, can be used to provide fast and efficient localization.

Informed *S-Graphs* (*iS-Graphs*) [8] further extend *S-Graphs* by using architectural plans to provide accurate localization over the resulting hierarchical factor graphs. However, its success is based on the assumption that there are no deviations between the “as-built” and the “as-planned”. In reality, this is never the case, and the building elements exhibit certain deviations with respect to their planned geometries.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

The main contribution of this paper is a novel method capable of coupling architectural plans (“as-planned”) and SLAM (“as-built”) data even in the presence of deviations, as shown in Fig. 1. We call this algorithm *Deviations Informed Situational Graphs* or *diS-Graphs* in short.

In summary, *diS-Graphs* performs three interconnected tasks simultaneously:

- Tightly coupling the SLAM factor graph with architectural plans, in order to match structural elements (*walls, rooms*) between planned and built environments.
- Globally localizing the robot in imprecise architectural plans that may contain deviations from the real building.
- Detecting and estimating the deviations between the matched structural elements of “as-planned” and “as-built” environments in real-time.

II. RELATED WORKS

Most localization techniques using prior information from architectural plans assume that the environments are built precisely according to the plans. One of the most commonly used localization techniques in 2D metric prior maps is Monte Carlo Localization (MCL) [9], [10] but it is not scalable to large-scale complex environments. Boniardi et al [11] use a technique that scales to more complex environments by aligning a scan-based map with CAD-based floor plans. OGM2PGM [12] also scales to larger environments by converting the 2D floor plan to an occupancy grid map (OGM) and using a pose-graph map (PGM) to localize the robot. UKFL [13] further enhances the localization accuracy using an unscented Kalman filter to localize the robot in 3D metric meshes. Recent techniques such as [14] exploit neural networks to localize the robot using an implicit neural representation of the floor plans. These techniques rely solely on geometric information from architectural plans, ignoring available semantic data (room types, door locations, functional spaces) that could disambiguate similar areas and improve localization. Additionally, these geometric-only approaches are vulnerable to floor plan inaccuracies or outdated information.

To address these limitations, semantic-based localization techniques, such as Mendez et al. [15] use semantic cues from architectural plans and sensor information to improve localization accuracy. Boniardi et al. [16] exploit the semantics of the room in architectural plans to do robot localization by matching the detected rooms from sensor data. Wang et al. [17] leverage pre-labeled architectural features, such as wall intersections and corners, as landmarks in floor plans, and match them with detection from sensor data to jointly perform mapping and localization. Zimmerman et al. [18], [19] use high-level semantic information in floor plans, derived from object detection, along with geometric data from 2D LiDAR to perform long-term robot localization in floor plans. Huan et al. [20] convert architectural plans into semantically enriched point cloud maps, followed by a coarse-to-fine localization process using ICP. Gao et al. [21] used neural networks to detect vertical elements from floor plans to do LiDAR based localization. These methods are prone to inaccuracies due to misidentification and errors in the pose estimate of semantic

elements. Moreover, they treat semantic elements in isolation, failing to leverage the rich contextual information embedded in their spatial and functional relationships. Recent work, such as Shaheer et al. [8] exploits the topological relationship between semantic elements to localize the robot with respect to architectural plans. However, all the above mentioned approaches assume no deviations between the architectural plans and the actual environment.

Some recent works leverage imprecise floor plans for localization. Boniardi et al. [22] integrate localization techniques to take advantage of the information embedded in the CAD drawing, and the real-world observations acquired during navigation, which may not be reflected in the floor plan. Li et al. [23] presented a 2D LiDAR-based localization system in imprecise floor plans using stochastic gradient descent (SGD) with a scan matching algorithm. Chan et al. [24] presented a 2D LiDAR-based localization in floor plans that integrates SLAM with MCL. Blum et al. [25] use neural networks for feature segmentation and combine them with LiDAR data for localization in imprecise floor plans. Despite advances in robot localization using imprecise floor plans, no existing approach combines localization with the estimation of structural deviations between “as-planned” and “as-built” environments, limiting the ability to assess construction accuracy in real time. In this work, we address this limitation by presenting a unified framework that enables both global localization and deviation estimation.

III. SYSTEM ARCHITECTURE

We propose an algorithm that tightly couples SLAM and architectural plans and jointly optimizes them.

A. Factor Graph Structure and Notation

Our algorithm represents the SLAM backend and the architectural plans using hierarchical factor graphs (see Fig. 3). For clarity, we first introduce here the terminology and notation used throughout this paper.

Architectural Graph (A-Graph). Three-layered hierarchical factor graph containing the geometry, semantics, and topology of an environment, generated from its architectural plan [8]. It models the environment “as-planned” by the architect.

Situational Graph (S-Graph). Four-layered hierarchical optimizable factor graph, estimated online from 3D LiDAR and odometry measurements, that models the “as-built” environment [6], [7]. It also includes keyframes in addition to the geometry, semantics, and topology of the environment.

Deviations Informed Situational Graph (diS-Graph). Joint graph, coupling both the *A-Graph* and the *S-Graph* of a building. The layers of these graphs are as follows.

Origins Layer.

The *S-Graph* and *A-Graph* are defined in their respective frames \mathcal{S} and \mathcal{A} , with anchor frame \mathcal{W} set such that $\mathbf{p}_{\mathcal{W}}^{\mathcal{A}} = \mathbf{0}$.

Keyframes Layer.

Exclusive to *S-Graphs*, this layer contains robot poses in the \mathcal{S} frame: $\mathbf{x}_{r_i}^{\mathcal{S}} \in SE(3) \forall i$.

Walls Layer. In *A-Graph*: Contains wall-surfaces $\pi_{a_j}^{\mathcal{A}} \in S(2) \times \mathbb{R} \forall j$ and walls $\omega_{a_k}^{\mathcal{A}} \in SE(3) \forall k$. Wall-surfaces

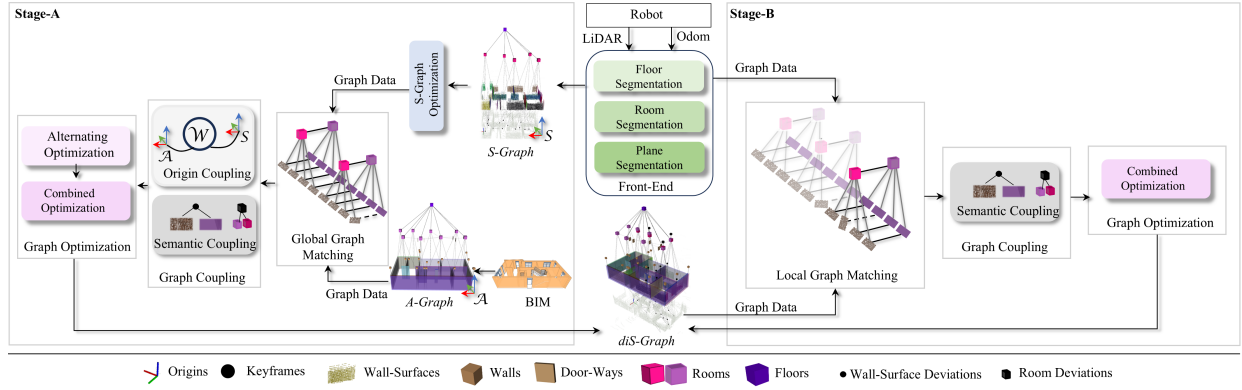


Fig. 2: **System Architecture.** The inputs to our method are an A -Graph generated from an architectural plan, and an S -Graph estimated online from the 3D LiDAR and the robot odometry. Stage-A is run first, and only once, for initial estimates of global localization and deviation. Once Stage-A is successful, Stage-B is run sequentially to match, couple, and optimize newly incorporated observations incrementally.




Fig. 3: **diS -Graph structure**, tightly coupling an A -Graph and a S -Graph. $d_{\rho_s}^{\rho_a}$ and $d_{\pi_s}^{\pi_a}$ refer to the deviation nodes between rooms and wall-surfaces respectively, with c_{d_ρ} and c_{d_π} as their associated cost function. \mathcal{W} is the anchor node, and, c_T is the associated cost function to estimate the transformation p_S^A between the two graph origins \mathcal{A} and \mathcal{S} .

are defined by their normals and distances to the origin [6]. A wall is constrained by two opposite wall-surfaces via the factor $\|f(\pi_{a_1}^A, \pi_{a_2}^A, s_w^A)\|_\Lambda^2$, defined in Eq. 2 of [8]. In S -Graph: Contains only wall-surfaces $\pi_{sl}^S \in S(2) \times \mathbb{R} \forall l$ extracted from 3D LiDAR scans and linked to keyframes via pose-plane constraints. We couple both graphs in our diS -Graphs via wall deviation nodes $d_{\pi_{s_i}^{\pi_a}}$ and their associated factors $c_{\pi_m} \forall m$, that we show in Fig. 3 and define in Eq. 3 in this paper.

Rooms Layer. In A -Graph: Contains rooms $\rho_{a_n}^A \in SE(3) \forall n$ composed of four wall-surfaces, and doorways $\mathcal{D}_{a_o}^A \in SE(3) \forall o$. A doorway constrains two rooms via the factor $\|f(\rho_{a_1}^A, \mathcal{D}_{a_1}^{\rho_2}) - f(\rho_{a_2}^A, \mathcal{D}_{a_2}^{\rho_1})\|_\Lambda^2$, that is defined in Eq. 3 of [8]. A room is constrained by four wall-surfaces via factor $f(\pi_{x_1}^A, \pi_{x_2}^A, \pi_{y_1}^A, \pi_{y_2}^A)\|_\Lambda^2$, defined in Eq. 7 of [7]. In S -Graph: Contains rooms of four wall-surfaces $\rho_{s_p}^S \in SE(3) \forall p$ or two wall-surfaces $\kappa_{s_q} \in SE(3) \forall q$. Both graphs are connected by room deviation factors $d_{\rho_{s_i}^{\rho_a}}$ in our diS -Graph.

Floors Layer. In both A -Graph and S -Graph this layer consists of a floor center node represented as $\xi_s^A \in SE(3)$ and $\xi_s^S \in SE(3)$ respectively, constraining all rooms present at that particular floor level with the factor $\|f(\xi_{a,s}^A, \rho_{a_n, s_p}^A)\|_\Lambda^2$,

defined in Eq. 9 of the reference [7].

More details on the type of factors between the different elements of the graphs can be found in [7] [8].

B. The Algorithm

Our algorithm has two stages, as shown in Fig. 2. In Stage-A, the A -Graph is first matched and coupled with the S -Graph, and then jointly optimized. Stage-A runs until it finds the initial match between the two graphs, and provides the initial estimates of both the transformation between graphs and potential deviations between their elements. Afterwards, in Stage-B, the joint graph is continuously refined through incremental matching and optimization as the robot navigates the environment, incorporating newly detected semantic entities. This tight coupling enables continuous refinement of both the robot's localization and the detected structural deviations during navigation. Both stages of our algorithm consist of multiple processes, introduced briefly here and fully detailed later in their respective sections.

Graph Matching (Section IV). The Global Graph Matching in Stage-A provides the first unique match, when there is sufficient information, between the S -Graph and A -Graph at room and wall-surface levels, accounting for potential deviations. In Stage-B, the Local Graph Matching extends the previously matched elements in the diS -Graph with newly detected elements following an incremental approach.

Graph Coupling (Section V). It is performed for the candidates matched by graph matching. In Stage-A, graph coupling registers the origins of two graphs along with the semantic coupling (wall-surfaces and rooms), and explicit mapping of the deviation factors. Stage-B only performs the semantic coupling between newly observed entities and their correspondences in the plan, with explicit deviation factors.

Graph Optimization (Section VI). It optimizes the coupled diS -Graph. Stage-A involves two steps: alternating optimization and joint optimization. Alternating optimization estimates two types of variables, 1) The transformation between the origins of the S -Graph and the A -Graph, yielding global localization, and 2) the potential deviations between

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

the matched graph entities. Stage-B jointly optimizes the *diS-Graph* from initial seeds for the global transformation and map deviations.

IV. GRAPH MATCHING

Our graph matching extends the method presented in [8]. In [8], a top-down potential candidate search between an *A-Graph* and a *S-Graph* is performed by leveraging their hierarchical structures. To assess the overall consistency of matching candidate, two verification steps are applied iteratively. First, the consistency of the node type and graph structure is verified. And secondly, geometric consistency is checked over a certain predefined threshold [26]. Finally, the consistencies of all candidates are compared, and the candidate with the best consistency is selected as the final unique match. Note that symmetries may affect our algorithm. Their occurrence may result in false positives when the robot has only partially explored the environment and observed one of the symmetric places. They may also cause several equally good candidates if the robot has explored a larger fraction of the building.




Fig. 4: **Illustration of the graph matching workflow** between the semantic nodes of an *A-Graph* and a *S-Graph*.

A. Deviation-based Candidates Search

The presence of deviations generates geometric inconsistencies that affect the aforementioned candidates' checks of [8]. Concretely, a deviation in the position of a wall-surface implies a slight deviation in the center of its parent room as well. To handle this, we propose a two-stage search method in this paper, searching first for non-deviated wall-surfaces and including after that those elements that are deviated.

Non-deviated wall-surfaces search To handle potential deviations, we tighten the matching criteria. Specifically, during the generation of non-deviated *room-to-room* and *room-to-wall-surface* match candidates, we look for the candidates with higher consistency at each level.

Deviated wall-surfaces detection To identify deviated wall-surfaces that were not matched in the first stage but are connected to already matched rooms, we now relax the matching criteria at *wall-surface-to-wall-surface* level to account for inconsistencies due to deviations. To further speed up the candidate search, we incorporate the following additional information.

Orphan wall-surfaces: We utilize wall-surfaces in the *S-Graph* without a parent room for the assessment of the

geometrical consistency of the final match candidates at the wall-surface level.

Ground orientation: We exploit the ground plane normal in the *A-Graph* and the *S-Graph*, only allowing candidates with z-axis rotations. Finally, the geometric consistency score provides a quantification of the probability of the deviation for each room and wall-surface, which is further used in the graph coupling stage (Section V).

B. Incremental Matching

To enhance the efficiency of the Graph Matching algorithm in [8], we implement an incremental approach with two stages (associated with stages A and B of the systems architecture of Fig. 2), namely **Global Graph Matching** (Fig. 4 left) and **Local Graph Matching** (Fig. 4 right), each of them based on the two deviation-based candidate search stages described in Section IV-A. Until a first unique match has been found, the Global Graph Matching is executed for every new observation in the *S-Graph*. Afterwards, the Local Graph Matching is run every time the *diS-Graph* is updated with newly observed rooms and wall-surfaces. The graph elements that have been already matched are excluded from candidate generation here, and each assessment of intra-level consistency considers the previously matched elements at the corresponding level.

V. GRAPH COUPLING

Origin Coupling. We couple the two graphs using the anchor node \mathcal{W} as a common reference frame, as shown in Fig. 3. This allows us to align both graphs in the frame \mathcal{W} by estimating the transformation \mathbf{p}_S^A between \mathcal{A} and \mathcal{S} . The associated constraint c_T is defined as:

$$c_T = \|\mathbf{p}_S^A \ominus \mathbf{p}_S^{\mathcal{W}} \oplus \mathbf{p}_A^{\mathcal{W}}\|_{\Lambda_T}^2 \quad (1)$$

$\Lambda_{\bar{T}}$ stands for the covariance of the cost, and it is assigned a high value to estimate the transformation accurately. Note that, since our approach simultaneously estimates both the global transformation and structural deviations, special care should be put in choosing the initial seed for this transformation and avoiding local minima coming from the combined effect of wall deviations and local registration.

Semantic Entities Coupling. To estimate deviations between the wall-surfaces and rooms of the *A-Graph* and the *S-Graph*, we couple them with deviation nodes $\mathbf{d}_{\pi_{s_i}^{a_i}}$ and $\mathbf{d}_{\rho_{s_i}^{a_i}}$ respectively, as shown in Fig. 3.

Room Coupling: To estimate the deviation $\mathbf{d}_{\rho_{s_i}^{a_i}}$ in the pose of a pair of matched rooms of the two graphs, we define the cost function $c_{\mathbf{d}_{\rho_i}}$ as:

$$c_{\mathbf{d}_{\rho_i}} = \|\mathbf{d}_{\rho_{s_i}^{a_i}} \ominus \rho_{s_i}^{\mathcal{W}} \oplus \rho_{a_i}^{\mathcal{W}}\|_{\Lambda_{\mathbf{d}_{\rho_i}}}^2 \quad (2)$$

Here $\rho_{a_i}^{\mathcal{W}}$ and $\rho_{s_i}^{\mathcal{W}}$ are the room poses of *A-Graph* and *S-Graph* respectively, in common frame \mathcal{W} . $\Lambda_{\bar{\rho}_i}$ is the covariance associated with the cost function depending on the probability of deviation estimated by the matching of the graph (Section IV). Rooms with a higher probability of deviation assigned by graph matching have a higher covariance assigned to their cost function than rooms with lower deviation probability. If

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

all matched rooms have the same deviation probability, they are assigned lower uniform covariances.

Wall-Surface Coupling: To estimate the deviation $\mathbf{d}_{\pi_{s_i}^{\pi_{a_i}}}$ between a pair of matched wall-surfaces of both graphs, we define the cost function of deviation factor between two wall-surfaces $c_{\mathbf{d}_{\pi}}$ as:

$$c_{\mathbf{d}_{\pi_i}} = \|\mathbf{d}_{\pi_{s_i}^{\pi_{a_i}}} \ominus \pi_{s_i}^{\mathcal{W}} \oplus \pi_{a_i}^{\mathcal{W}}\|_{\Lambda_{\mathbf{d}_{\pi_i}}}^2 \quad (3)$$

Here $\pi_{a_i}^{\mathcal{W}}$ and $\pi_{s_i}^{\mathcal{W}}$ are the poses of wall-surfaces of *A-Graph* and *S-Graph* respectively, in common frame \mathcal{W} .¹ $\Lambda_{\mathbf{d}_{\pi_i}}$ is the covariance associated with the cost function. Like rooms, wall-surfaces with a higher deviation are assigned higher covariances, and the ones with a lower deviation are assigned lower uniform covariances.

VI. GRAPH OPTIMIZATION

Before the *S-Graph* and *A-Graph* are coupled, the *S-Graph* state at time t can be defined as [7]:

$$\mathbf{s}_{1_S} = [\mathbf{x}_{r_i}^S, \pi_{s_i}^S, \rho_{s_p}^S, \kappa_{s_q}^S, \xi_s^S, \mathbf{x}_O^S]^\top \quad (4)$$

Similarly, the *A-Graph* state is defined as [8]:

$$\mathbf{s}_{1_A} = [\pi_{a_j}^A, \rho_{a_n}^A, \mathcal{D}_{a_o}^A, \xi_a^A]^\top \quad (5)$$

where $\mathcal{S}_{\mathbf{x}_O}$ models the drift between the odometry frame O and the *S-Graph* reference frame S , and the other variables are described in Section III. Note, if at time t there is no match between the *A-Graph* and *S-Graph*, we perform single *S-Graph* optimization as detailed in [7] until we detect it.

Alternating Optimization. Alternating optimization is further performed in two steps as follows:

Transformation Estimation: Upon receiving the match (Section IV) and performing graph coupling (Section V), we combine Eq. 4 and Eq. 5 to make a global state with additional transformation $\mathbf{s}_2 = [\mathbf{s}_{1_S}, \mathbf{s}_{1_A}, \mathbf{p}_S^A]$. It is important to note that at this stage, the deviated wall-surface and room are not included for estimating \mathbf{p}_S^A .

Deviation Estimation: After optimizing \mathbf{s}_2 we already have an initial guess of the transformation between the *A-Graph* and the *S-Graph*, and we can incorporate the deviated wall-surface and room entities into the graph with appropriate deviation factors. Our state then becomes $\mathbf{s}_3 = [\mathbf{s}_2, \mathbf{d}_{\pi_{s_i}^{\pi_{a_i}}}, \mathbf{d}_{\rho_{s_i}^{\pi_{a_i}}}]$. When optimizing \mathbf{s}_3 we keep \mathbf{s}_2 constant to obtain a good initial estimation of the deviation between the matched deviated entities.

Joint Optimization. Finally, after getting the initial estimates of the transformation between the origins and the deviations between the semantic entities, we optimize the whole state \mathbf{s}_3 to simultaneously estimate the position of each semantic entity, deviations, and the transformation between the two graphs.

¹For compactness, we slightly abused notation by overloading the \oplus and \ominus operators. In Eq. 2, they denote the pose composition and pose difference, respectively, and, in Eq. 3, they denote the plane compositions and difference, respectively.

VII. EXPERIMENTAL EVALUATION

A. Methodology

Setup. We evaluated our algorithm in both simulated and real environments. Both types of experiments were performed using a laptop computer with an Intel i9-11950H (8 cores, 2.6 GHz) with 32 GB of RAM.

Baselines. We have selected the following LiDAR-based baselines for comparisons due to their suitability, their reported results, and the availability of their code: *AMCL* [9], *UKFL* [13], *OGM2PGM* [12], *IR-MCL* [14], and *iS-Graphs* [8]. Although other works do localization in imprecise architectural plans, the absence of open-source implementations prevents direct comparison with these methods. As each selected baseline takes a different map input for localization, we used 2D occupancy grid maps for *AMCL* and *OGM2PGM*, 3D meshes for *UKFL* and *IR-MCL*, and *A-Graphs* for *iS-Graphs* and our method *diS-Graphs*, respectively.

Simulated Datasets. We validate the algorithms on five simulated datasets (*SE1* to *SE5*) using Gazebo to recreate the robot, LiDAR sensor, and 3D environments from actual architectural plans. We report absolute trajectory error (ATE) and localization convergence success rates.

Real Datasets. We collected data using a legged robot with Velodyne VLP-16 LiDAR at five construction sites (*RE1* to *RE5*) with existing architectural plans. We report Root Mean Square Error (RMSE) between estimated 3D maps and ground truth plans (indicating localization accuracy since mapping quality reflects pose estimation), plus convergence rates and computation times.

Deviations. Given the absence of actual ground truth deviations, we explicitly deviated wall-surfaces in architectural plans by 5 cm to 40 cm in translation, and 5° and 15° in rotations across five separate tests. We conducted 25 experiments per dataset to test algorithm robustness against these varying deviations.

Ablations. We conducted ablation studies on two key components of our algorithm. 1. **Covariance assignment:** As mentioned in the graph matching and graph merging (Section. IV & V), we assign covariances to the deviation factors of the semantic elements, based on their deviation likelihood. Here, we analyze the effect of assigning equal covariances, referred to as uniform covariances (UC), to both deviated and non-deviated elements. 2. **Alternating optimization:** In graph optimization (Section. VI) we discuss the use of alternating optimization for simultaneous localization and deviation estimation. Here, we study the performance when using only a single optimization cycle (SO).

B. Results and Discussion

Absolute Trajectory Error. Table I shows the average ATE for all baselines and our *diS-Graphs* in the presence of deviations between “as-planned” and “as-built” environments. Our method shows an error reduction of around 75.8% compared to *AMCL*, 64.9% compared to *OGM2PGM*, 69% compared to *IR-MCL*, 47.2% compared to *UKFL*, and a reduction of 43% over *iS-Graphs*. Fig. 6a summarizes the ATE performance of all the baseline algorithms. *AMCL* shows the highest median

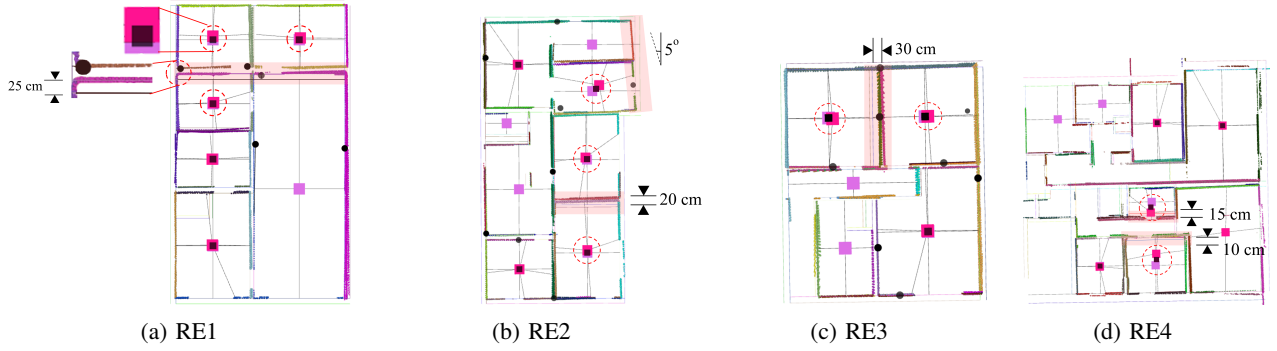


Fig. 5: *diS-Graphs* of real construction sites showing detected deviations. Red dotted circles indicate room deviations, red-shaded translucent rectangles highlight wall deviations, and black markers represent deviation factors. See in (a) a zoom-in of a room deviation (non-coincident room centers, represented by pink and purple squares) and of a wall-surface deviation.

ATE and a relatively narrow distribution, indicating consistently high error rates. UKFL and OGM2PGM demonstrate moderate performance with similar median ATEs, although OGM2PGM shows a wider range of errors. IR-MCL exhibits the largest variability, suggesting inconsistent performance in different scenarios. AMCL, OGM2PGM, IR-MCL, UKFL and *iS-Graphs* assume no deviations between “as-planned” and “as-built” environments. These methods attempt to match laser scans with an incorrect reference map, leading to consistent misalignment. Therefore, they show higher ATEs across datasets, indicating their vulnerability to architectural deviations. Our *diS-Graphs* can detect and estimate the deviations between “as-planned” and “as-built” environments, resulting in the lowest ATE.

Point Cloud Alignment Error. Table II shows the RMSE of the point clouds with respect to the ground truth for all baseline methods and ours. In the presence of deviations between “as-planned” and “as-built”, *diS-Graphs* shows 53.5% better accuracy than AMCL, 45.8% better than OGM2PGM, 51.8% better than IR-MCL, and 7% better than *iS-Graphs*. Although UKFL’s average error is equal to *diS-Graphs*’, it has a very low convergence rate (see Table III) rendering the comparison unfair. Moreover, it cannot estimate the deviations between “as-planned” and “as-built” environments. Fig. 6b summarizes the performance of all algorithms in real environments. AMCL exhibits the highest median and widest interquartile range,

TABLE I: Average ATE [cm] for simulated experiments. Each entry represents the mean of 5 tests with different amounts of deviation. **Bold** values are the best and the second best are underlined. ‘-’ refers to an unsuccessful run.

Method	Dataset					
	ATE [cm]					
	SE1	SE2	SE3	SE4	SE5	Avg.
AMCL [9]	—	17.2	—	20.1	22.4	19.9
UKFL [13]	12.6	15.3	—	8.7	11.1	9.1
OGM2PGM [12]	15.2	18.1	10.7	10.3	14.3	13.7
IR-MCL [14]	14.7	<u>6.4</u>	<u>9.6</u>	28.4	18.8	15.5
<i>iS-Graphs</i> [8]	<u>5.4</u>	6.7	16.6	<u>4.6</u>	<u>9.5</u>	<u>8.5</u>
<i>diS-Graphs</i> (UC)	6.2	7.3	13.7	8.4	8.6	8.8
<i>diS-Graphs</i> (SO)	4.4	6.1	15.2	7.7	6.2	7.9
<i>diS-Graphs</i> (Ours)	3.3	4.1	6.4	4.4	5.7	4.8

indicating greater variability in performance. IR-MCL presents a large spread of results, while OGM2PGM shows moderate performance with a smaller range of variability compared to AMCL and IR-MCL. Although UKFL and *diS-Graphs* show the lowest median RMSE, suggesting superior accuracy, the comparison with UKFL is not fair because of its low convergence rate. Because of our simultaneous estimation of deviations and initial transformation, we can not only simultaneously globally localize the robot but also estimate the deviations between semantic elements of “as-planned” and “as-built” environments, moreover improving the overall accuracy compared to other algorithms.

Convergence Rate. Our approach demonstrates superior convergence rates across both simulated and real environments (Table III). In simulations, our method achieves a convergence rate 60% better than AMCL, 44% better than UKFL, and 12% better than OGM2PGM. While our convergence rate matches that of IR-MCL in simulations, IR-MCL exhibits inconsistent performance in terms of ATE and requires retraining for each new dataset. Similarly, in real environments (*RE1-RE4*), our method maintains the highest convergence rate among all baselines, leveraging its ability to detect deviated elements in the environment. The only exception is in *RE5*, where our algorithm fails due to insufficient room detection for effective




Fig. 6: **a)** Distribution of Average Trajectory Error (ATE) in simulated datasets. **b)** Distribution of point cloud alignment error (RMSE) in real datasets. UKFL’s only 2 convergent cases prevent fair comparison.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 7: Comparison of Average Trajectory Error (ATE) against the amount of deviation in simulated datasets.

TABLE II: Point cloud RMSE [cm] for real experiments. **Bold** values are the best and the second best are underlined. ‘-’ refers to an unsuccessful run.

Method	Dataset						
	Point Cloud RMSE [cm]						
	RE1	RE2	RE3	RE4	RE5	Avg	
AMCL [9]	0.48	0.88	0.31	—	—	0.56	
UKFL [13]	0.31	—	0.20	—	—	0.26*	
OGM2PGM [12]	0.46	0.57	0.33	0.49	0.55	0.48	
IR-MCL [14]	0.51	0.76	0.36	—	—	0.54	
<i>iS-Graphs</i> [8]	<u>0.27</u>	<u>0.29</u>	<u>0.23</u>	<u>0.33</u>	—	<u>0.28</u>	
<i>diS-Graphs</i> (UC)	0.27	0.32	0.23	0.34	—	0.29	
<i>diS-Graphs</i> (SO)	0.27	0.30	0.23	0.33	—	0.28	
<i>diS-Graphs</i> (ours)	0.25	0.28	0.20	0.31	—	0.26	

* Omitted due to low convergence.

graph matching, because of noisy sensor data. Moreover, AMCL and UKFL’s heavy reliance on accurate initial position estimates leads to degraded convergence rates in environments with complex geometries. OGM2PGM’s accuracy suffers particularly in symmetric environments. Our method overcomes these limitations through its robust deviation detection capabilities, though it requires a sufficient number of distinguishable rooms to function effectively.

TABLE III: Convergence rate [%] of simulated and real experiments.

Method	Dataset									
	Convergence Rate [%]									
	SE1	SE2	SE3	SE4	SE5	RE1	RE2	RE3	RE4	RE5
AMCL [9]	0	100	0	80	20	80	100	100	0	0
UKFL [13]	80	80	0	40	80	60	0	60	0	0
OGM2PGM [12]	80	100	100	80	80	60	80	100	60	80
IR-MCL [14]	100	100	100	20	100	60	100	20	0	0
<i>iS-Graphs</i> [8]	60	20	100	80	100	60	20	40	60	0
<i>diS-Graphs</i>	100	100	100	100	100	100	100	100	100	0

Convergence and Computation Time. Table IV shows the convergence time for each algorithm. The convergence time is the time it takes the algorithm to globally localize. IR-MCL and OGM2PGM have the best convergence time. *iS-Graphs* and *diS-Graphs* have considerably longer convergence times because they need to detect a certain number of rooms in the environment for graph matching to find a unique match. The graph matching of [8] struggles to resolve symmetries during the matching process, resulting in longer convergence times. However, the modifications in the graph matching algorithm (Section IV) proposed in this work improve the symmetry resolution ability and reduce the convergence time by almost 20%. Table IV shows the computation time for each algorithm.

TABLE IV: Convergence time [s] and computation time [ms] on real experiments. **Bold** values are the best and the second best are underlined. ‘-’ refers to an unsuccessful run.

Method	Dataset									
	Convergence Time [s]					Computation Time [ms]				
	RE1	RE2	RE3	RE4	RE5	RE1	RE2	RE3	RE4	RE5
AMCL [9]	24	89	29	—	—	2	2	2	—	—
UKFL [13]	126	—	8	—	—	104	—	119	—	—
OGM2PGM [12]	16	38	27	35	79	2	2	2	2	2
IR-MCL [14]	16	26	13	—	—	92	90	89	—	—
<i>iS-Graphs</i> [8]	155	101	78	139	—	57	78	78	64	—
<i>diS-Graphs</i>	81	43	46	<u>139</u>	—	<u>56</u>	<u>77</u>	<u>76</u>	70	—
Seq. Len. [s]	657	170	488	657	559	657	170	488	657	559

We specifically report the time spent for each pose update. On average, the computation time of our algorithm is 70 milliseconds per pose update, showing real-time performance. Our computation time is the best compared to other 3D algorithms. OGM2PGM and AMCL. The smaller computation time of the others is due to their simpler optimization in 2D.

Deviation Estimation. Fig. 8 shows the amount of deviation our *diS-Graphs* can correctly estimate in real environments. The minimum detectable deviation is bounded by the LiDAR resolution. With our LiDAR errors having a standard deviation of 3 cm [27], and considering that 99.7% of measurements fall within three standard deviations, our *diS-Graph* can reliably detect deviations only when they exceed 9 cm. The maximum translational and rotational wall-surface deviation that our *diS-Graphs* can detect is 35 cm and 15° respectively. Fig. 5 shows several qualitative results from our real experiments. The robot can successfully localize, map, and estimate deviations in these environments. Note that we only show the rooms and walls for better understanding, and all the other semantic elements and the robot are not shown. We do not show RE5 in Fig. 5, as the robot could not localize itself in this sequence.




Fig. 8: Average deviation detection rate for real datasets.

Robustness against Deviations. Figure 7 shows the effect of deviations on the localization performance of *iS-Graphs* and *diS-Graphs*. The results consistently show that *diS-Graphs*,

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

which incorporates deviation detection and modeling, maintains a relatively stable ATE even as the deviation increases from 10 cm to 35 cm. In contrast, *iS-Graphs*, which lacks explicit deviation modeling capabilities, exhibits a clear upward trend in ATE as deviations increase.

Ablation Study. Table I shows that associating ‘uniform covariance’ in simulated datasets (*diS-Graphs* (UC)) the algorithm cannot differentiate between deviated and non-deviated elements which results in poor pose estimation. Fig. 7 demonstrates that *iS-Graphs* with uniform covariance shows higher error variability compared to our method which uses deviation probability-based covariance assignment. In some cases, the use of uniform covariances results in even worse performance than *iS-Graphs*. Similarly, when using single optimization (*diS-Graphs* (SO)) instead of alternating optimization, the algorithm cannot differentiate between the transformation between two graphs and the deviations between their elements, resulting in higher ATE as shown in Table I. Table II shows the ablation of uniform covariances and single optimization in real-world datasets. Using uniform covariances (*diS-Graphs* (UC)) and single optimization (*diS-Graphs* (SO)) results in lower mapping accuracy due to the algorithm’s inability to accurately map the deviated elements and differentiate between initial transformation and deviations simultaneously.

VIII. CONCLUSIONS AND FUTURE WORK

This paper presents a novel approach to tightly couple SLAM with imprecise architectural plans. Our algorithm establishes direct correspondences between structural elements in “as-built” environments and their “as-planned” counterparts, enabling simultaneous robot localization and deviation estimation between the two representations. By detecting and estimating the deviations between “as-built” and “as-planned” environments through tight coupling, our method outperforms the current best approach with 43% better localization in simulations, 7% improved mapping accuracy in real environments, and enhanced robustness to architectural deviations. Our algorithm provides an estimate of existing deviations up to 35 cm in translation and 15° in rotation. Our algorithm is limited by the need for sufficient distinctive semantic elements (i.e., wall-surfaces and rooms) to provide unique matches. As future work, we plan to detect and match more semantic elements and support rooms with non-rectangular geometries, including curved walls and partial floorplans. Additionally, we aim to integrate multi-modal sensing (RGB-D, vision-based semantic segmentation) to enhance robustness in visually degraded environments. These improvements will enhance the convergence rate and enable deviation detection for additional detected entities beyond just rooms and walls.

REFERENCES

- [1] M. Jenewein and P. Ferschin, “Validation of construction deviations using parametric bim and augmented reality,” in *Proceedings of the 19th International Conference on Construction Applications of Virtual Reality (CONVR2019)*, 2019, p. 10.
- [2] S. Jia, C. Liu, X. Guan, H. Wu, D. Zeng, and J. Guo, “Bidirectional interaction between bim and construction processes using a multisource geospatial data enabled point cloud model,” *Automation in Construction*, vol. 134, p. 104096, 2022.
- [3] H. Moravec and A. Elfes, “High resolution maps from wide angle sonar,” in *IEEE International Conference on Robotics and Automation*, 1985.
- [4] I. Armeni, Z.-Y. He, A. Zamir, J. Gwak, J. Malik, M. Fischer, and S. Savarese, “3d scene graph: A structure for unified semantics, 3d space, and camera,” in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*.
- [5] N. Hughes, Y. Chang, and L. Carlone, “Hydra: A real-time spatial perception system for 3d scene graph construction and optimization,” in *Robotics: Science and Systems XVIII*.
- [6] H. Bavle, J. L. Sanchez-Lopez, M. Shaheer, J. Civera, and H. Voos, “Situational graphs for robot navigation in structured indoor environments,” vol. 7, no. 4, pp. 9107–9114, *IEEE Robotics and Automation Letters*.
- [7] —, “S-graphs+: Real-time localization and mapping leveraging hierarchical representations,” *IEEE Robotics and Automation Letters*, vol. 8, no. 8, pp. 4927–4934, 2023.
- [8] M. Shaheer, J. A. Millan-Romera, H. Bavle, J. L. Sanchez-Lopez, J. Civera, and H. Voos, “Graph-based global robot localization informing situational graphs with architectural graphs,” in *2023 International Conference on Intelligent Robots and Systems*, pp. 9155–9162.
- [9] D. Fox, W. Burgard, F. Dellaert, and S. Thrun, “Monte carlo localization: Efficient position estimation for mobile robots,” *Aaai/iaai*, vol. 1999, no. 343-349, pp. 2–2, 1999.
- [10] M. Shaheer, H. Bavle, J. L. Sanchez-Lopez, and H. Voos, “Robot localization using situational graphs (s-graphs) and building architectural plans,” *Robotics*, vol. 12, no. 3, p. 65, 2023.
- [11] F. Boniardi, T. Caselitz, R. Kümmerle, and W. Burgard, “A pose graph-based localization system for long-term navigation in cad floor plans,” *Robotics and Autonomous Systems*, vol. 112, pp. 84–97, 2019.
- [12] M. V. Torres, A. Braun, and A. Borrmann, “Ogm2pgbm: Robust bim-based 2d-lidar localization for lifelong indoor navigation,” in *ECPPM 2022-eWork and eBusiness in Architecture, Engineering and Construction 2022*, 2023, pp. 567–574.
- [13] K. Koide, J. Miura, and E. Menegatti, “A portable three-dimensional LIDAR-based system for long-term and wide-area people behavior measurement,” vol. 16, no. 2, p. 1729881419841532.
- [14] H. Kuang, X. Chen, T. Guadagnino, N. Zimmerman, J. Behley, and C. Stachniss, “Ir-mcl: Implicit representation-based online global localization,” *IEEE Robotics and Automation Letters*, vol. 8, no. 3, pp. 1627–1634, 2023.
- [15] O. Mendez, S. Hadfield, N. Pugeault, and R. Bowden, “SeDAR - semantic detection and ranging: Humans can localise without LiDAR, can robots?” in *2018 International Conference on Robotics and Automation*, pp. 6053–6060.
- [16] F. Boniardi, A. Valada, R. Mohan, T. Caselitz, and W. Burgard, “Robot localization in floor plans using a room layout edge extraction network,” in *2019 International Conference on Intelligent Robots and Systems*.
- [17] X. Wang, R. J. Marcotte, and E. Olson, “GLFP: Global localization from a floor plan,” in *2019 International Conference on Intelligent Robots and Systems*, pp. 1627–1632.
- [18] N. Zimmerman, T. Guadagnino, X. Chen, J. Behley, and C. Stachniss, “Long-term localization using semantic cues in floor plan maps,” *IEEE Robotics and Automation Letters*, vol. 8, no. 1, pp. 176–183, 2022.
- [19] N. Zimmerman, M. Sodano, E. Marks, J. Behley, and C. Stachniss, “Constructing Metric-Semantic Maps Using Floor Plan Priors for Long-Term Indoor Localization,” in *International Conference on Intelligent Robots and Systems*, Oct. 2023, pp. 1366–1372.
- [20] Y. Huan, L. Zhiyi, and Y. K.W.Justin, “Semantic localization on BIM-generated maps using a 3D LiDAR sensor,” *Automation in Construction*, vol. 146, p. 104641, 2023.
- [21] L. Gao and L. Kneip, “FP-loc: Lightweight and drift-free floor plan-assisted LiDAR localization,” in *2022 International Conference on Robotics and Automation*, pp. 4142–4148.
- [22] F. Boniardi, T. Caselitz, R. Kummerle, and W. Burgard, “Robust LiDAR-based localization in architectural floor plans,” in *2017 International Conference on Intelligent Robots and Systems*, pp. 3318–3324.
- [23] Z. Li, M. H. Ang, and D. Rus, “Online localization with imprecise floor space maps using stochastic gradient descent,” in *2020 International Conference on Intelligent Robots and Systems*, pp. 8571–8578.
- [24] C. L. Chan, J. Li, J. L. Chan, Z. Li, and K. W. Wan, “Partial-map-based monte carlo localization in architectural floor plans,” in *Social Robotics*. Springer International Publishing, vol. 13086, pp. 541–552.
- [25] H. Blum, J. Stiefel, C. Cadena, R. Siegwart, and A. Gavel, “Precise robot localization in architectural 3d plans,” *arXiv preprint arXiv:2006.05137*, 2020.
- [26] P. C. Lusk, K. Fathian, and J. P. How, “Clipper: A graph-theoretic framework for robust data association,” in *International Conference on Robotics and Automation*, 2021, pp. 13 828–13 834.
- [27] C. L. Glennie, A. Kusari, and A. Facchin, “Calibration and Stability Analysis of the VLP-16 Laser Scanner,” vol. XL-3/W4, pp. 55–60.