

# Multi-Modal 3D Scene Graph Updater for Shared and Dynamic Environments

**Emilio Olivastri**

Department of Information Engineering  
University of Padova  
Italy  
olivastrie@dei.unipd.it

**Jonathan Francis**

Bosch Center for AI  
Carnegie Mellon University  
United States  
jon.francis@us.bosch.com

**Alberto Pretto**

Department of Information Engineering  
University of Padova  
Italy  
alberto.pretto@dei.unipd.it

**Niko Sünderhauf**

Centre for Robotics  
Queensland University of Technology  
Australia  
niko.suenderhauf@qut.edu.au

**Krishan Rana**

Centre for Robotics  
Queensland University of Technology  
Australia  
ranak@qut.edu.au

## Abstract:

The advent of generalist Large Language Models (LLMs) and Large Vision Models (VLMs) have streamlined the construction of semantically enriched maps that can enable robots to ground high-level reasoning and planning into their representations. One of the most widely used semantic map formats is the 3D Scene Graph, which captures both metric (low-level) and semantic (high-level) information. However, these maps often assume a static world, while real environments, like homes and offices, are dynamic. Even small changes in these spaces can significantly impact task performance. To integrate robots into dynamic environments, they must detect changes and update the scene graph in real-time. This update process is inherently multimodal, requiring input from various sources, such as human agents, the robot’s own perception system, time, and its actions. This work proposes a framework that leverages these multimodal inputs to maintain the consistency of scene graphs during real-time operation, presenting promising initial results and outlining a roadmap for future research.

**Keywords:** 3D Scene Graphs, Task Planning, Lifelong Learning

## 1 Introduction

A fundamental goal in robotics research is to enable robots to operate safely, efficiently, and autonomously in environments shared with humans. In such shared spaces—whether homes, offices, or other collaborative settings—robots must perform a wide range of tasks, from simple object manipulation to complex interactions with the environment. To achieve this, robots rely heavily on internal representations of their surroundings, typically in the form of maps. These maps allow robots to understand spatial layouts and reason about tasks at both low-level (e.g., navigation) and high-level (e.g., task planning). However, one key assumption often made when using these maps is that the environment remains static after the map is created.

In practice, this assumption breaks down in dynamic, shared environments where human and robotic activities constantly alter the state of the world. Furniture is moved, objects are relocated, and items are used or discarded. These changes are non-trivial, and any discrepancies between the robot’s internal map and the real-world environment can significantly degrade task performance. For instance, a robot instructed to retrieve a banana from the kitchen may fail if it relies on outdated map information—such as the banana being present a week ago—without accounting for the possibility that it has since been eaten or thrown away. Such inconsistencies can lead to inefficient behavior, incorrect task execution, and ultimately limit the robot’s usefulness in dynamic environments.

To prevent this degradation in performance, robots must possess the ability to detect changes in their environment and update their internal maps in real time. This capability is essential for ensuring that downstream reasoning and task planning are based on an accurate, up-to-date representation of the world. A critical insight from this work is that maintaining the consistency of a robot’s map in shared, living environments is inherently a multimodal challenge. It requires the robot to integrate information from multiple sources—such as visual perception, interactions with humans and other agents, temporal data, and its own actions—to inform map updates. The ability to leverage this diverse set of information streams is key to building robust, real-time scene representations that evolve as the environment changes.

One of the most prominent representations used in robotics for modeling complex environments is the 3D Scene Graph (3DSG)[1]. 3DSGs encode both metric and semantic information, allowing robots to reason about not only the spatial configuration of objects but also their relationships and roles in the environment. The flexibility of 3DSGs makes them a powerful tool for representing dynamic environments, as they can be extended to include different types of information. For example, hierarchical relationships between objects and their surrounding spaces (e.g., an object belonging to a room) can be incorporated[2], as well as semantic relationships between objects (e.g., objects that are typically used together) [3]. Despite the versatility of 3DSGs, existing research has primarily focused on their construction and optimization for specific tasks, with limited attention to how these representations can be dynamically updated to reflect ongoing changes in the real world.

To address this critical gap, we propose MM-3DSGU, a general framework designed to dynamically update 3D Scene Graphs in shared, dynamic environments. MM-3DSGU leverages heterogeneous sources of information, including time, human interactions, the robot’s perception system, and its own actions, to perform robust change detection and maintain an accurate representation of the environment. By integrating these multimodal streams, our framework ensures that the robot’s scene graph is consistently aligned with the real world, enabling better task execution and reasoning in dynamic settings.

In this work, we present the design and initial implementation of MM-3DSGU, along with promising results from preliminary experiments. We also outline a detailed roadmap for future research, aimed at further refining our approach and addressing challenges related to scalability, efficiency, and real-time performance in more complex environments.

## 2 Related Work

**3D Scene Graphs:** 3D Scene Graphs (3DSGs) [1] are a powerful representation for modeling large-scale environments by abstracting them into graphs, where nodes represent objects and spatial concepts, and edges capture relationships between them. This structure enables more efficient high-level reasoning for robotic tasks such as navigation and manipulation [4, 5]. 3DSGs are particularly well-suited for capturing both geometric and semantic attributes, as nodes and edges can carry various properties, such as object affordances, spatial hierarchies, and semantic relationships [3, 6]. These representations have been shown to improve planning and reasoning capabilities, as demonstrated in systems like SayPlan [7] and language-guided robots [8]. While 3DSGs offer a rich framework for representing static environments, their ability to remain up-to-date in dynamic, real-world settings is limited. Most existing works focus on constructing and optimizing 3DSGs from static data, which can quickly become outdated in shared environments where objects and spatial layouts are

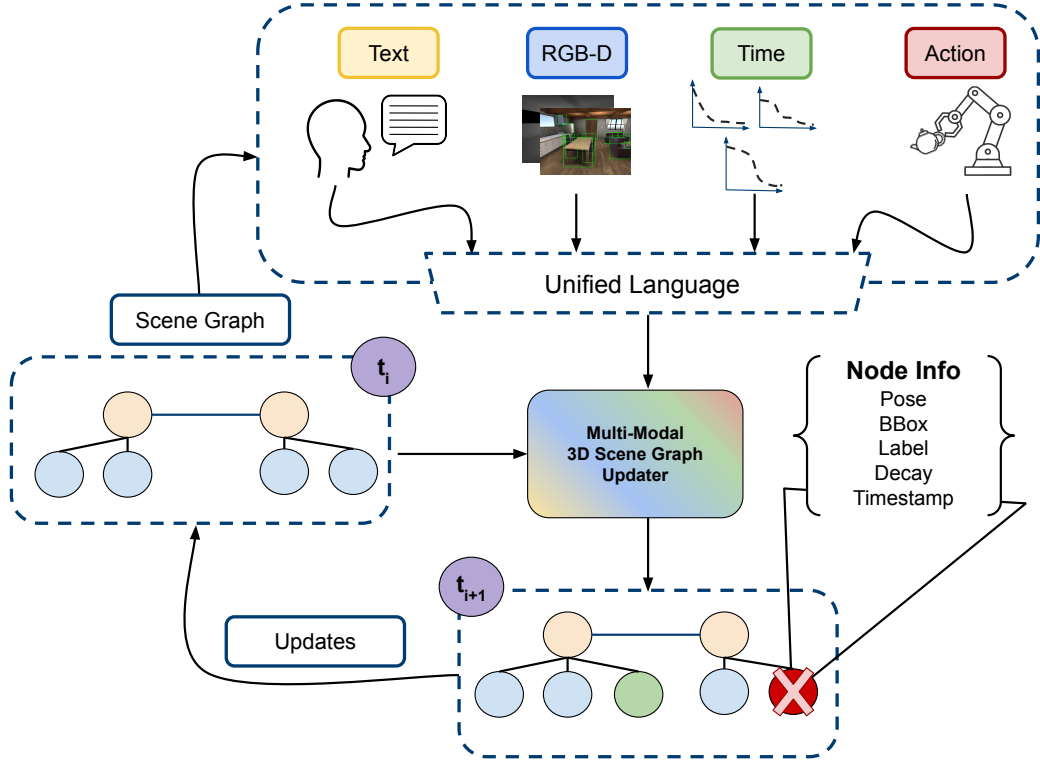


Figure 1: **High-level overview of the proposed framework.** The various modalities proposed for change detection are displayed (top). These modules leverage the latest version of the 3D scene graph to assess whether change has occurred. Update operations on the 3D scene graph (bottom) are applied agnostically, due to the use of a unified language.

constantly changing. This limitation is particularly problematic in non-stationary environments, as robots relying on outdated 3DSG representations are prone to errors in task execution. Our work addresses this gap by proposing MM-3DSGU, a framework that enables 3DSG updates using multi-modal information sources, allowing robots to maintain a consistent and relevant understanding of their environment over time.

**Spatio-temporal Mapping:** To handle dynamic environments, spatio-temporal mapping methods have been developed to model object motion and environmental changes over time. Bore et al.[9] introduced a probabilistic approach to account for both short-term and long-term object motions within the environment. Building on this, Soares et al.[10] enhanced the robustness of ORB-SLAM3 [11] by filtering out dynamic points and applying a more reliable data association scheme to handle long-term dynamics. Similarly, Qian et al.[12, 13] proposed an object-aware factor-graph formulation for SLAM that explicitly incorporates long-term dynamics of objects, allowing the map to evolve with changes in object locations. Khronos[14] further improves the handling of short-term dynamics by embedding objects’ histories, ensuring that maps capture the temporal evolution of the environment. While these approaches provide valuable insights into dynamic mapping, they primarily focus on updating lower-level geometric maps. Our work extends beyond this by incorporating higher-level, semantic updates into 3DSGs, maintaining the consistency of the entire scene graph—both at the object and semantic relationship level—over time. This enables robots to not only account for spatial and geometric changes but also for semantic shifts, such as changes in object functions or relationships, which are critical for high-level reasoning.

**Change Detection:** Change detection is a key problem in dynamic environments, involving the identification of differences between two observations of the same scene [15, 16, 17, 18]. This

problem is often addressed by comparing sensor data before and after changes occur to detect differences that warrant updates to the robot’s internal map. Schmid et al.[19] proposed using semantic consistency to detect changes and update the map accordingly, while Fu et al.[20, 21] leveraged the prior that objects tend to exist on planes, reducing the complexity of change detection in indoor environments. However, these approaches often focus on detecting changes in the geometric structure of the environment or object-level properties. Our framework, MM-3DSGU, expands upon these methods by incorporating multimodal inputs (such as temporal data, human interactions, and the robot’s actions) to identify both geometric and semantic changes. This multimodal approach allows our system to not only detect when changes have occurred but also understand their context and relevance to the robot’s task, ensuring that the updated scene graph is more aligned with the robot’s current operational needs.

### 2.1 Robot Learning in Non-Stationary Environments

Robots operating in non-stationary environments must continually learn and adapt to changes in their surroundings. This has been an ongoing area of research in robotics, dating back to early work in Reinforcement Learning (RL) and Bayesian belief modeling, where robots attempt to learn policies that can generalize across varying environmental conditions [22, 23, 24]. However, there is no standardized approach for continual learning in non-stationary environments [25, 26], with existing methods often being task-specific or reliant on simplistic assumptions about the dynamics of the environment. Recent advances in lifelong learning for robots have focused on enabling adaptation in changing environments through techniques such as transfer learning, online learning, and meta-learning [27, 28, 29]. Yet, these approaches rarely address how scene representations themselves should be updated in real-time as the environment evolves. Our work fills this gap by proposing a structured update mechanism for 3D Scene Graphs, where the robot can dynamically modify its scene representation to reflect changes in the environment, ensuring continual alignment with real-world conditions. Moreover, unlike previous work that often assumes an explicit model of the environment’s dynamics, our framework leverages a flexible, multimodal approach that can generalize to a wide range of tasks and environments. By integrating diverse streams of information (e.g., visual perception, human-robot interactions, temporal context), we move toward a more generalisable solution for continual learning in dynamic settings, addressing a key limitation in existing literature.

## 3 Problem Definition

The problem we address can be framed as a **Partially Observable Markov Decision Process (POMDP)** [30, 31], where a robot operates in a dynamic environment represented by a **3D Scene Graph (3DSG)**. The 3DSG encodes both geometric and semantic information about objects and spaces, which the robot uses for reasoning and task execution. The environment is dynamic, with changes introduced by human activities and the robot’s own actions. At each time step, the robot receives partial observations from multiple modalities—such as human input, robot perception, and time—while it performs actions like exploring the environment or executing tasks. The robot’s challenge is to maintain an accurate scene graph by updating it as changes occur, ensuring its internal representation aligns with the real-world environment.

The core of the problem lies in defining an efficient **update function**  $\mathcal{U} : O \times G_t \rightarrow G_{t+1}$ , which takes the robot’s observations and the current scene graph and outputs an updated graph that reflects the latest changes. This multimodal approach integrates data from robot actions, human interactions, temporal semantics, and sensor inputs to detect and incorporate changes in both object locations and semantic relationships. The goal is to minimize discrepancies between the scene graph and the actual environment, ensuring that the 3DSG remains consistent and up-to-date, which in turn enables more effective task performance and decision-making in dynamic, non-stationary settings.

## 4 Methodology

The proposed framework is agnostic to the change detection algorithms employed, provided that the algorithm adheres to the language and syntax needed to operate changes on the 3DSG through the usage of general primitives. In this work, we assume that the robot’s pose is known, as in [2][3], and that an initial 3DSG has been built using any preferred state-of-the-art algorithm [4][32].

### 4.1 Minimal Required 3D Scene Graph

There have been numerous variations of 3DSGs [2][3][4], with differing hierarchical levels used to organize information and types of data embedded within the nodes. In this work, the minimal granularity required consists of two hierarchical levels: the *Object* and *Room* layers. Depending on the extension of the workspace, it can be easily extended to include additional layers, such as *Floors* and *Buildings*, by applying the same logic used for the *Room* layer. Here we focus on discussing the necessary information for the first two layers. In the *Object* layer, as the name suggests, only object nodes are present. Each node contains the following information:

- **Pose:** The 6DoF pose of the object  $o$  with respect to the world  $w$ , represented using the matrix  ${}^w\mathbf{X}_o \in SE(3)$ .
- **3D Bounding Box:** Minimal bounding box required to encapsulate the object  $o$ , represented using the vector  $\mathbf{b}_o \in R_+^3$ , where  $\mathbf{b}_o = (w, h, d)$ , that are respectively the width, height and depth of the bounding box.
- **Label:** String that represents the semantic class of the object, represented with the vector  $\mathbf{l}_o \in L$ , where  $L$  is the set representing human language.
- **Decay Rate:** Parameter  $\lambda_o \in R_+$  that controls the evolution of the likelihood that the object  $o$  remains stationary after a time interval  $\Delta t$ . If the object is immovable, then  $\lambda_o = 0$ .
- **Timestamp:** Value  $t^o \in R_+$  representing when the object  $o$  was last observed.

The *Object* layer doesn’t require edges between objects as opposed to [3]. Only inter-layer edges, expressing the *belonging* relation (i.e., an object  $o$  belonging to a room  $r$ ), are used. In the *Room* layer, as the name suggests, only room nodes are present. Each node contains a subset of the object node’s elements: the pose  ${}^w\mathbf{X}_r \in SE(3)$ , 3D Bounding Box  $\mathbf{r}_o$ , and its label  $\mathbf{l}_r$ . Additionally, room nodes are connected using intra-layer edges that express *access to* relation (from one room to another).

### 4.2 Multi-Modal Change Detection

During the execution of a mission, the robot perceives, interacts with, and acts on the environment. While the research community has primarily focused on passive perception for map updates, we propose incorporating new potential sources of meaningful information, such as external agents and time. The Multi-Modal Change Detection module, one of our key contributions, mimics how a human would notice, be informed about, and predict changes in the environment in which they live or operate. Depending on the type of information received, the module will select the most appropriate change detection algorithm.

**Human Input:** In a scenario where both robots and humans have access to the same spaces, the robot is no longer the only entity causing changes and moving objects. Humans are highly dynamic entities, and as such they will be responsible for most of the changes. Humans can acquire new objects, dispose of old ones, or move them from one location to another. Thus, it would be extremely beneficial to leverage interactions with humans, for the robot to be informed of the changes that have occurred. With the advent of LLMs [33][34] and their high-level reasoning, processed text has become a reliable source of information. We leverage the capabilities of LLMs to identify the type of operation, the object of interest  $o$ , a possible supporting object  $so$  (which could be used to disambiguate semantically similar objects), the source room  $sr$  representing the location from which

the object is moved or removed, and the target room  $tr$  representing the location to which the object is moved or added. Humans communicate in a topological manner: “I moved the cup from the kitchen to the table in the living room”. In this example, which represents a common scenario, no accurate geometric indication object’s placement is provided. Thus, when adding a new object or moving an existing one, the update operation will also be topological, involving a switch in the room to which the object belongs. The perception module will be then responsible for updating the node with accurate geometric information.

**Robot Action:** If the robot lives in a shared environment with humans, we want it to assist us in our daily chores and tasks to make our lives easier. Although there is a wide variety of missions a human could assign to a robot, most share the same key components of interest: namely *pick* and *place*. These action primitives are the ones that cause changes in the environment. All the information needed to propagate these changes is found in the task description and in the task status. From the task description, we can infer the object  $o$  of interest, along with the source room  $sr$  and target room  $tr$ . From the task status, we can determine when the object  $o$  is being *picked* (temporarily removed from the 3DSG) and *placed* (re-attached to the 3DSG).

**Robot Perception:** Most robots are equipped with cameras to perceive the external world similarly to humans. In this work, we assume the use of an RGB-D camera, allowing the robot to perceive simultaneously both visual and geometric features. To detect changes, the robot needs to compare the objects that it is expected to see with those currently perceived. Thus, the first step consists of querying the latest available version of the 3DSG for the list of objects that should be detected from its current pose  ${}^w\mathbf{X}_{rb}^t$ . We define the set of expected visible objects as  $VO = \{o \in 3DSG \mid o \in DO \wedge o \text{ visible from } {}^w\mathbf{X}_{rb}^t\}$ , where  $DO$  is the set of all dynamic objects. Next, we define the set of currently observed objects  $BO$ , where  $bo \in BO$  is represented as  $bo = ({}^w\mathbf{X}_{bo}, \mathbf{b}, \mathbf{l})$ . The semantic classes of the objects can be inferred using a closed-world object detector [35], while their poses and bounding boxes can be estimated using the point cloud obtained by projecting the depth values corresponding to the masks obtained from an object segmentation system [36]. The next step involves semantically and geometrically associating  $o \in VO$  and  $bo \in BO$ . If both the semantic and geometric tests are passed, it signifies that  $o$  has remained static. If only the semantic test is passed, then it is likely that  $o$  has been moved to the location where  $bo$  was observed. The non-associated objects in  $VO$  are those we expected to observe but did not detect, making them candidates for removal from the 3DSG. Instead, the non-associated observed objects in  $BO$  are those that were detected when they were not expected, making them candidates for being added to the 3DSG as new objects. To reduce data association complexity, we filter out the objects that are either immovable or unlikely to be moved, such as refrigerators and pantries.

**Time:** Time is an additional source of information that has never been explicitly utilized to inform the robot of potential changes in the environment. As humans, we calibrate our expectations of how an environment evolves when we are not observing it for a period of time  $\Delta t$ , based on our prior semantic knowledge about persistency of objects in the environment and how others would interact with the various objects within that same environment. In this work, we simulate human’s prior knowledge of how objects are interacted by utilizing the following function:

$$p(d(o_t, o_{t_0}) < \epsilon \mid t, \mathbf{l}) = \frac{2}{1 + e^{\lambda_1(t-t_0)}}, \quad (1)$$

$$d(o_t, o_{t_0}) = \|{}^w\mathbf{X}_o^t \ominus {}^w\mathbf{X}_o^{t_0}\| \quad (2)$$

Where  $d(o_t, o_{t_0}) = \|{}^w\mathbf{X}_o^t \ominus {}^w\mathbf{X}_o^{t_0}\|$  represents the distance between the poses of the same object when observed at times  $t$  and  $t_0$ .  $\epsilon$  represents the maximum displacement magnitude allowed for an object  $o$  to still be considered static. The probability  $p(d(o_t, o_{t_0}) < \epsilon \mid t, \mathbf{l})$  represents the likelihood that the object  $o$  remains static after  $t - t_0$ , where  $t_0$  is when it was last observed. The parameter  $\lambda_1$  represents the decay rate of the function, controlling how rapidly the probability  $p(d(o_t, o_{t_0}) < \epsilon \mid t, \mathbf{l})$  decreases. In other words, it indicates how quickly we expect an object to be interacted with by another human. This allows for a ‘living’ 3DSG which can evolve nodes on the graph as time passes.

Humans expect that the dynamics of different objects evolve differently based on their semantic class. To emulate this behaviour, we utilise LLMs to estimate the decay rate  $\lambda_{\mathbf{I}}$  of an object  $o$  given the semantic class  $\mathbf{I}$ . This choice stems from the general-purpose knowledge and semantic understanding of LLMs, which are derived from their training on internet-scale information.  $\lambda_{\mathbf{I}} = 0.0$  indicates that the object is hardly moved or not moved at all. Each time an object  $o$  is added, moved, or observed, the value  $t_0$  is updated. Objects  $o \in 3DSG$  with  $p(d(o_t, o_{t_0}) < \epsilon \mid t, \mathbf{I})$  below a certain threshold represent targets for the robot to actively perceive, as it becomes likely that these objects have been interacted with.

### 4.3 Scene Graph Updater

All these different modules provide a heterogeneous source of information. To establish a general framework that performs updates to the 3DSG, which is agnostic to the change detection algorithm used, it is necessary to define a *Unified Language* to be used as the output of the change detection algorithm. It can be thought of as a dictionary with the following fields:

- Source Room: Room’s label  $\mathbf{I}_r$  in which the object  $o$  belonged initially.
- Target Room: New room’s label  $\mathbf{I}_r$  in which the object belongs to now.
- Target Object: Object’s label  $\mathbf{I}_o$ .
- Pose: Object’s estimated pose  ${}^w\mathbf{X}_{bo}$ .
- 3D Bounding Box: Object’s estimated bounding box  $\mathbf{b}_o$ .
- Action: Three possible values that represent the outcome of the change detection algorithm: added/moved/removed.
- Support Object: Semantic class of an object  $so$  that has some topological relation with the object  $o$ .

Using this information, the primitives for updating the 3DSG can be simplified to basic primitives that operate on any generic graph. The specific primitives and their sequence of use depend on the *Action* field. The list of implemented primitives is the following:

- Find( $\mathbf{I}$ ): Primitive used to find the node in the graph that corresponds to the input’s label.
- Add( $tr, \mathbf{I}_o, {}^w\mathbf{X}_o, \mathbf{b}_o$ ): Primitive used to create a new object node and assign it to the target room  $tr$ .
- Remove( $sr, \mathbf{I}_o$ ): Primitive used to remove an existing object node from its source room  $sr$ .
- Move( $sr, tr, \mathbf{I}_o, {}^w\mathbf{X}_o$ ): Primitive used to update the object’s node information, which can be described as sequentially executing Remove and Add operations.

## 5 Experimental Setup

**Simulated Environment:** To test the pipeline, a Unity environment has been created to enable control over all possible changes while also providing ground truth information about the objects within the environment. This environment consists of a single-story house with four rooms: a Bathroom, a Bedroom, a Living Room, and a Kitchen. Each room is populated with various sets of objects typically found in those spaces, for which we possess both CAD models and ground truth poses.

**Mission Specifications:** The robot, equipped with the MM-3DSGU module, is initially provided with a 3DSG built using the ground truth data available from the simulator before any changes are applied to the environment. A set of *virtual actions*, simulating interactions a real human might have had with the environment, is then introduced. The goal, following the robot’s mission, is for it to detect all the changes made and accordingly update the 3DSG it was initially provided with. The simulated mission is conducted as follows: the robot is informed by a human about a change they

caused, such as, “I ate the banana that was in the kitchen”. It is then assigned a mission that begins with a *pick* action and end with a *place* action, for example, “Pick the mug that is in the kitchen and take it to the bedroom.” Throughout the mission, the robot passively monitors the environment for any changes that may have occurred.

The specifications of the simulated mission are as follows:

- Human Action/Query: “I removed the towel from the bathroom because it was too old.”
- Mission: “Pick the mug in the kitchen and take it to the bedroom.”
- List of virtual actions: Removed “banana”, moved “tv remote” from “table” to “sofa’s arm”, added “book” on top of “bed”.

**Evaluation:** During the experimentation, the update operations performed on the 3DSG using MM-3DSGU are logged and annotated. For a preliminary and exploratory evaluation, these logged operations are compared with the real modifications, which were manually annotated. To ensure more fair and robust testing, the mission described was executed 10 times, and the reported results are averaged across those runs.

## 6 Results

Tab. 6 summarizes the results obtained from the preliminary experiments. For each type of update operation, we measure the success rate and highlight the modules responsible for any occurred failures. The *Time* module was not evaluated due to its passive nature and the need for more rigorous formalization and metrics. The *Action* module could never fail because grasping failures from the robot were not simulated. The *RGB-D* module successfully recognized all the new objects introduced in the environment, but it incorrectly removed the “tv remote” from the 3DSG, failing to recognize that it just had been moved to another location. YOLOv8 [35] struggles to detect smaller objects, such as the “tv remote”, resulting in the incorrect removal of these items. GPT-4 [34] was employed for semantic association. The errors in failing to associate objects of the same semantic class and incorrectly associating objects of different semantic classes were mitigated by verifying that the associations were consistent during sequential observations. The *Text* module success rate can be attributed to the topological nature of the updates, and the continuous improvement of GPT-4 [34]. The geometric attributes of added or moved objects will be updated by the *RGB-D* module when perceived. The removal operation doesn’t require any refinement.

Table 1: Results of the framework obtained on the simulated scenario described above. The table showcases the different types of updates and the modules that were the cause of failure.

Update Type	Success Rate		Failure Rate		
		Text	RGB-D	Action	Time
Add	100.0%	0.0%	0.0%	0.0%	-
Remove	66.67%	0.0%	33.33%	0.0%	-
Move	66.67%	0.0%	33.33%	0.0%	-

**Future Work:** MM-3DSGU incorporates the concept of time into our object-centric representation, although this has yet to be explicitly exploited or evaluated. This research opens up the possibility of incorporating time into the robot’s path planning to periodically check whether objects that are likely to have been moved remain static during the execution of other tasks. New metrics to assess the impact of this integration on overall mission performance still need to be formalized. Currently, the decay rate  $\lambda$  is inferred using GPT-4 [34] from a set of discretized values, with each value associated with a *representative* object serving as a basis for inferring  $\lambda$ . These values were fine-tuned based on the authors’ expectations regarding the dynamics of the selected elements. While this decay-tuning



serves as a proof-of-concept, it would be beneficial to develop a more general and informed method for setting the parameter  $\lambda$  or to incorporate learning for adapting these parameters to the specific environments where the robot operates. Further experimentation in both simulation and real-world settings is needed to effectively validate the system performance. This work highlights the necessity for new types of annotations in datasets, such as temporal data, especially if the robot is expected to share workspaces with humans, along with the need for new metrics to evaluate these performances.

## **7 Conclusion**

In this work, we presented preliminary insights about MM-3DSG, the first general framework for updating a 3DSG that combines human interactions, robot actions, and perception while accounting for the evolution of the environment over time. To address heterogeneous sources of information, we introduced a common language and primitives for unifying and generalizing update operations on a 3DSG. The initial simulated experimentation indicates that this approach shows promising results, opening up various new research directions.

## Acknowledgments

I would like to thank my friend and colleague, Nicolas Chapman, for always being open to discussion and for helping me navigate various challenges.

## References

- [1] I. Armeni, Z.-Y. He, J. Gwak, A. R. Zamir, M. Fischer, J. Malik, and S. Savarese. 3D scene graph: A structure for unified semantics, 3D space, and camera. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5664–5673, 2019.
- [2] A. Werby, C. Huang, M. Büchner, A. Valada, and W. Burgard. Hierarchical open-vocabulary 3D scene graphs for language-grounded robot navigation. In *First Workshop on Vision-Language Models for Navigation and Manipulation at ICRA 2024*, 2024.
- [3] Q. Gu, A. Kuwajerwala, S. Morin, K. M. Jatavallabhula, B. Sen, A. Agarwal, C. Rivera, W. Paul, K. Ellis, R. Chellappa, et al. Conceptgraphs: Open-vocabulary 3D scene graphs for perception and planning. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5021–5028. IEEE, 2024.
- [4] N. Hughes, Y. Chang, and L. Carlone. Hydra: A real-time spatial perception system for 3D scene graph construction and optimization. *Robotics: Science and Systems (RSS)*, 2022.
- [5] E. Greve, M. Büchner, N. Vödisch, W. Burgard, and A. Valada. Collaborative dynamic 3D scene graphs for automated driving. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11118–11124. IEEE, 2024.
- [6] S.-C. Wu, J. Wald, K. Tateno, N. Navab, and F. Tombari. Scenegrappfusion: Incremental 3D scene graph prediction from RGB-D sequences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7515–7525, 2021.
- [7] K. Rana, J. Haviland, S. Garg, J. Abou-Chakra, I. Reid, and N. Suenderhauf. Sayplan: Grounding large language models using 3D scene graphs for scalable robot task planning. In *7th Annual Conference on Robot Learning*, 2023.
- [8] D. Honerkamp, M. Büchner, F. Despinoy, T. Welschehold, and A. Valada. Language-grounded dynamic scene graphs for interactive object search with mobile manipulation. *IEEE Robotics and Automation Letters*, 2024.
- [9] N. Bore, J. Ekekrantz, P. Jensfelt, and J. Folkesson. Detection and tracking of general movable objects in large three-dimensional maps. *IEEE Transactions on Robotics*, 35(1):231–247, 2018.
- [10] J. C. Virgolino Soares, V. S. Medeiros, G. F. Abati, M. Becker, G. Caurin, M. Gattass, and M. A. Meggiolaro. Visual localization and mapping in dynamic and changing environments. *Journal of Intelligent & Robotic Systems*, 109(4):95, 2023.
- [11] C. Campos, R. Elvira, J. J. Gómez, J. M. M. Montiel, and J. D. Tardós. ORB-SLAM3: An accurate open-source library for visual, visual-inertial and multi-map SLAM. *IEEE Transactions on Robotics*, 37(6):1874–1890, 2021.
- [12] J. Qian, V. Chatrath, J. Yang, J. Servos, A. P. Schoellig, and S. L. Waslander. POCD: Probabilistic object-level change detection and volumetric mapping in semi-static scenes. *arXiv preprint arXiv:2205.01202*, 2022.
- [13] J. Qian, V. Chatrath, J. Servos, A. Mavrinac, W. Burgard, S. L. Waslander, and A. P. Schoellig. POV-SLAM: Probabilistic object-aware variational slam in semi-static environments. *Robotics: Science and Systems (RSS)*, 2023.

- [14] L. Schmid, M. Abate, Y. Chang, and L. Carlone. Khronos: A unified approach for spatio-temporal metric-semantic slam in dynamic environments. In *Proc. of Robotics: Science and Systems (RSS)*, Delft, Netherlands, July 2024.
- [15] E. Langer, T. Patten, and M. Vincze. Robust and efficient object change detection by combining global semantic information and local geometric verification. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8453–8460. IEEE, 2020.
- [16] J.-M. Park, J.-H. Jang, S.-M. Yoo, S.-K. Lee, U.-H. Kim, and J.-H. Kim. Changesim: Towards end-to-end online scene change detection in industrial indoor environments. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8578–8585. IEEE, 2021.
- [17] L. Zhu, S. Huang, K. Schindler, and I. Armeni. Living scenes: Multi-object relocalization and reconstruction in changing 3D environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28014–28024, 2024.
- [18] T. Sun, Y. Hao, S. Huang, S. Savarese, K. Schindler, M. Pollefeys, and I. Armeni. Nothing stands still: A spatiotemporal benchmark on 3D point cloud registration under large geometric and temporal change. *arXiv preprint arXiv:2311.09346*, 2023.
- [19] L. Schmid, J. Delmerico, J. L. Schönberger, J. Nieto, M. Pollefeys, R. Siegwart, and C. Cadena. Panoptic multi-TSDFs: a flexible representation for online multi-resolution volumetric mapping and long-term dynamic scene consistency. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 8018–8024. IEEE, 2022.
- [20] J. Fu, C. Lin, Y. Taguchi, A. Cohen, Y. Zhang, S. Mylabathula, and J. J. Leonard. PlaneSDF-based change detection for long-term dense mapping. *IEEE Robotics and Automation Letters*, 7(4):9667–9674, 2022.
- [21] J. Fu, Y. Du, K. Singh, J. B. Tenenbaum, and J. J. Leonard. Robust change detection based on neural descriptor fields. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2817–2824. IEEE, 2022.
- [22] I. J. Cox and J. J. Leonard. Modeling a dynamic environment using a bayesian multiple hypothesis approach. *Artificial intelligence*, 66(2):311–344, 1994.
- [23] S. Thrun. Bayesian landmark learning for mobile robot localization. *Machine learning*, 33:41–76, 1998.
- [24] S. Padakandla, P. KJ, and S. Bhatnagar. Reinforcement learning algorithm for non-stationary environments. *Applied Intelligence*, 50(11):3590–3606, 2020.
- [25] G. Graffieti, G. Borghi, and D. Maltoni. Continual learning in real-life applications. *IEEE Robotics and Automation Letters*, 7(3):6195–6202, 2022.
- [26] S. Padakandla. A survey of reinforcement learning algorithms for dynamically varying environments. *ACM Computing Surveys (CSUR)*, 54(6):1–25, 2021.
- [27] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu. A survey on deep transfer learning. In *Artificial Neural Networks and Machine Learning–ICANN 2018: 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4-7, 2018, Proceedings, Part III 27*, pages 270–279. Springer, 2018.
- [28] S. C. Hoi, D. Sahoo, J. Lu, and P. Zhao. Online learning: A comprehensive survey. *Neurocomputing*, 459:249–289, 2021.
- [29] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey. Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):5149–5169, 2021.

- [30] M. Lauri, D. Hsu, and J. Pajarinen. Partially observable markov decision processes in robotics: A survey. *IEEE Transactions on Robotics*, 39(1):21–40, 2022.
- [31] M. T. Spaan. Partially observable markov decision processes. In *Reinforcement learning: State-of-the-art*, pages 387–414. Springer, 2012.
- [32] D. Maggio, Y. Chang, N. Hughes, M. Trang, D. Griffith, C. Dougherty, E. Cristofalo, L. Schmid, and L. Carlone. Clio: Real-time task-driven open-set 3D scene graphs, 2024.
- [33] A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [34] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [35] R. Varghese and M. Sambath. YOLOv8: A novel object detection algorithm with enhanced performance and robustness. In *2024 International Conference on Advances in Data Engineering and Intelligent Computing Systems (ADICS)*, pages 1–6. IEEE, 2024.
- [36] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023.