

Hydra: A Real-time Spatial Perception System for 3D Scene Graph Construction and Optimization

Nathan Hughes, Yun Chang, Luca Carlone
Laboratory for Information & Decision Systems (LIDS)
Massachusetts Institute of Technology
Cambridge, USA
Email: {na26933, yunchang, lcarlone}@mit.edu

Abstract—3D scene graphs have recently emerged as a powerful high-level representation of 3D environments. A 3D scene graph describes the environment as a layered graph where nodes represent spatial concepts at multiple levels of abstraction (from low-level geometry to high-level semantics including objects, places, rooms, buildings, etc.) and edges represent relations between concepts. While 3D scene graphs can serve as an advanced “mental model” for robots, how to build such a rich representation in real-time is still uncharted territory.

This paper describes a *real-time* Spatial Perception System, a suite of algorithms to build a 3D scene graph from sensor data in real-time. Our first contribution is to develop real-time algorithms to incrementally construct the layers of a scene graph as the robot explores the environment; these algorithms build a local Euclidean Signed Distance Function (ESDF) around the current robot location, extract a topological map of places from the ESDF, and then segment the places into rooms using an approach inspired by community-detection techniques. Our second contribution is to investigate loop closure detection and optimization in 3D scene graphs. We show that 3D scene graphs allow defining *hierarchical descriptors* for loop closure detection; our descriptors capture statistics across layers in the scene graph, ranging from low-level visual appearance to summary statistics about objects and places. We then propose the first algorithm to optimize a 3D scene graph in response to loop closures; our approach relies on *embedded deformation graphs* to simultaneously correct all layers of the scene graph. We implement the proposed Spatial Perception System into a highly parallelized architecture, named *Hydra*¹, that combines fast early and mid-level perception processes (e.g., local mapping) with slower high-level perception (e.g., global optimization of the scene graph). We evaluate Hydra on simulated and real data and show it is able to reconstruct 3D scene graphs with an accuracy comparable with batch offline methods despite running online.

Index Terms—Robot perception, 3D scene graphs, localization and mapping, real-time scene understanding.

I. INTRODUCTION

The next generation of robots and autonomous systems will be required to build persistent high-level representations of unknown environments in real-time. *High-level* representations are required for a robot to understand and execute instructions from humans (e.g., “bring me the cup of tea I left on the dining room table”); high-level representations also enable fast planning (e.g., by allowing planning over compact abstractions rather than dense low-level geometry). Such representations

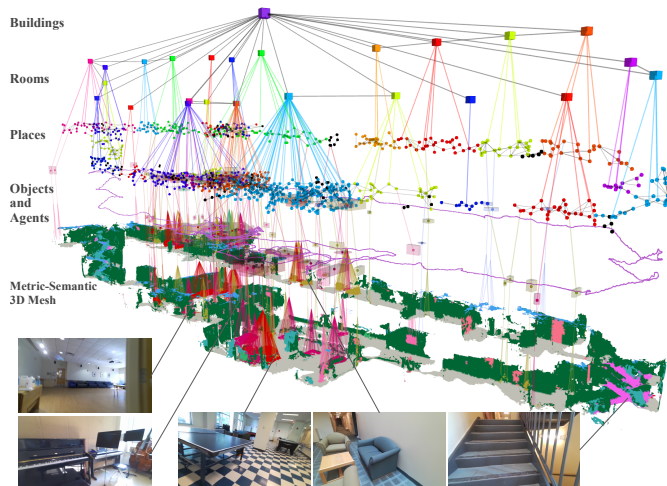


Fig. 1. We present *Hydra*, a highly parallelized architecture to build 3D scene graphs from sensor data in real-time. The figure shows sample input data and the 3D scene graph created by Hydra in a large-scale real environment.

must be built in *real-time* to support just-in-time decision-making. Moreover, these representations must be *persistent* to support long-term autonomy: (i) they need to scale to large environments, (ii) they should allow for corrections as new evidence is collected by the robot, and (iii) their size should only grow with the size of the environment they model.

3D Scene Graphs [4, 26, 49, 50, 63, 67] have recently emerged as powerful high-level representations of 3D environments. A 3D scene graph (Fig. 1 and Fig. 6) is a layered graph where nodes represent spatial concepts at multiple levels of abstraction (from low-level geometry to objects, places, rooms, buildings, etc.) and edges represent relations between concepts. Armeni *et al.* [4] pioneered the use of 3D scene graphs in computer vision and proposed the first algorithms to parse a metric-semantic 3D mesh into a 3D scene graph. Kim *et al.* [26] reconstruct a 3D scene graph of objects and their relations. Rosinol *et al.* [49, 50] propose a novel 3D scene graph model that (i) is built directly from sensor data, (ii) includes a subgraph of places (useful for robot navigation), (iii) models objects, rooms, and buildings, and (iv) captures moving entities in the environment. More recent work [22, 25, 63, 67] infers objects and relations from point clouds, RGB-D sequences, or object detections.

While 3D scene graphs can serve as an advanced “mental

¹Hydra is available at <https://github.com/MIT-SPARK/Hydra>

model” for robots, how to build such a rich representation in real-time remains uncharted territory. The works [26, 63, 67] allow real-time operation but are restricted to “flat” 3D scene graphs and are mostly concerned with objects and their relations while disregarding the top layers in Fig. 1. The works [4, 49, 50], which focus on building truly hierarchical representations, run offline and require several minutes to build a 3D scene graph ([4] even assumes the availability of a correct and complete metric-semantic mesh of the environment built beforehand). Extending our prior works [49, 50] to operate in real-time is non-trivial. These works utilize an Euclidean Signed Distance Function (ESDF) of the entire environment to build the 3D scene graph. Unfortunately, the memory required for ESDFs scale poorly in the size of the environment [43]. Moreover, the extraction of places and rooms in [49, 50] involves batch algorithms that process the entire ESDF, whose computational cost grows over time and is incompatible with real-time operation. Finally, the ESDF is reconstructed from the robot trajectory estimate which keeps changing in response to loop closures. The approaches in [49, 50] would therefore need to rebuild the scene graph from scratch after every loop closure, clashing with real-time operation.

The main motivation of this paper is to overcome these challenges and develop the first *real-time* Spatial Perception System, a suite of algorithms and implementations to build a hierarchical 3D scene graph, such as the one pictured in Fig. 1, from sensor data in real-time.

Our first contribution is to develop real-time algorithms to incrementally reconstruct the layers of a scene graph as the robot explores the environment. The proposed algorithms reconstruct a *local* ESDF of the robot’s surroundings and incrementally convert the ESDF into a metric-semantic 3D mesh as well as a *Generalized Voronoi Diagram*, from which a topological graph of places can be quickly extracted. This computation is incremental and runs in constant-time regardless of the size of the environment. Our algorithms also perform a fast and scalable room segmentation using a community-detection inspired approach that clusters the graph of places into rooms in a matter of milliseconds.

Our second contribution is to investigate loop closure detection and optimization in 3D scene graphs. We propose a novel hierarchical approach for loop closure detection: the proposed approach involves (i) a *top-down loop closure detection* that uses hierarchical descriptors —capturing statistics across layers in the scene graph— to find putative loop closures and (ii) a *bottom-up geometric verification* that attempts estimating the loop closure pose by registering putative matches. Then, we propose the first algorithm to optimize a 3D scene graph in response to loop closures; our approach relies on *embedded deformation graphs* to simultaneously correct all layers of the scene graph, *i.e.*, the 3D mesh, places, objects, and rooms.

Our final contribution is to develop a real-time architecture and implementation, and demonstrate the resulting Spatial Perception System on challenging simulated and real data. In particular, we propose a highly parallelized implementation, named *Hydra*, that combines fast early and mid-level

perception processes (*e.g.*, local mapping) with slower high-level perception (*e.g.*, global optimization of the scene graph). We evaluate Hydra in several heterogeneous environments, including an apartment complex, an office building, and a subway. Our experiments show that (i) we can reconstruct 3D scene graphs of large, real environments in real-time, (ii) our online algorithms achieve an accuracy comparable to batch offline methods, and (iii) our loop closure detection approach outperforms standard approaches based on bag-of-words and visual-feature matching in terms of quality and quantity of detected loop closures. The source code of Hydra is publicly available at <https://github.com/MIT-SPARK/Hydra>.

II. RELATED WORK

Metric-semantic and Hierarchical Mapping. The last few years have seen a surge of interest towards *metric-semantic mapping*, simultaneously triggered by the maturity of traditional 3D reconstruction and SLAM techniques, and by the novel opportunities for semantic understanding afforded by deep learning. The literature has focused on both object-based maps [10, 14, 41, 42, 53, 56] and dense maps, including volumetric models [23, 38, 40], point clouds [6, 32, 61], and 3D meshes [48, 51]. Some approaches combine objects and dense map models [31, 39, 54, 68]. These approaches are not concerned with estimating higher-level semantics (*e.g.*, rooms) and typically return dense models that might not be directly amenable for navigation [44].

A second research line focuses on building *hierarchical map* models. Hierarchical maps have been pervasive in robotics since its inception [12, 28, 29, 62]. Early work focuses on 2D maps and investigates the use of hierarchical maps to resolve the apparent divide between metric and topological representations [18, 52, 70]. More recently, 3D scene graphs have been proposed as expressive hierarchical models for 3D environments. Armeni *et al.* [4] model the environment as a graph including low-level geometry (*i.e.*, a metric-semantic mesh), objects, rooms, and camera locations. Rosinol *et al.* [49, 50] augment the model with a topological map of places, as well as a layer describing dynamic entities in the environment. The approaches in [4, 49, 50] are designed for offline use. Other papers focus on reconstructing a graph of objects and their relations [26, 63, 67]. Wu *et al.* [67] predict objects and relations in real-time using a graph neural network. Izatt and Tedrake [25] parse objects and relations into a scene grammar model using mixed-integer programming. Gothoskar *et al.* [22] use an MCMC approach.

A somewhat parallel research line investigates how to *parse the layout of a building* from 2D or 3D data. A large body of work focuses on parsing 2D maps [9], including rule-based [27] and learning-based methods [34]. Friedman *et al.* [16] compute a Voronoi graph from a 2D occupancy grid, which is then labeled using a conditional random field. Recent work focuses on 3D data. Liu *et al.* [34] and Stekovic *et al.* [57] project 3D point clouds to 2D maps, which however is not directly applicable to multi-story

buildings. Furukawa *et al.* [17] reconstruct floor plans from images using multi-view stereo combined with a Manhattan World assumption. Lukierski *et al.* [37] use dense stereo from an omni-directional camera to fit cuboids to objects and rooms. Zheng *et al.* [71] detects rooms by performing region growing on a 3D metric-semantic model.

Loop Closures Detection and Optimization. Established approaches for visual loop closure detection in robotics trace back to place recognition and image retrieval techniques in computer vision; these approaches are broadly adopted in SLAM pipelines but are known to suffer from appearance and viewpoint changes [36]. Recent approaches investigate place recognition using image sequences [20, 55] or deep learning [3]. More related to our proposal is the set of papers leveraging semantic information for loop closure detection. Gawel *et al.* [21] perform object-graph-based loop closure detection using random-walk descriptors built from 2D images. Liu *et al.* [35] use similar object-based descriptors but built from a 3D reconstruction. Lin *et al.* [33] adopt random-walk object-based descriptors and then compute loop closure poses via object registration. Qin *et al.* [45] propose an object-based approach based on subgraph similarity matching. Zheng *et al.* [71] propose a room-level loop closure detector.

After a loop closure is detected, the map needs to be corrected accordingly. While this process is easy in sparse (*e.g.*, landmark-based) representations [11], it is non-trivial to perform in real-time when using dense representations. Stückler and Behnke [58] and Whelan *et al.* [66] optimize a map of *surfels*, to circumvent the need to correct structured representations (*e.g.*, meshes or voxels). Dai *et al.* [13] propose reintegrating a volumetric map after each loop closure. Reijgwart *et al.* [47] correct drift in volumetric representations by breaking the map into submaps that can be rigidly re-aligned after loop closures. Whelan *et al.* [65] propose a 2-step optimization that first corrects the robot trajectory and then deforms the map (represented as a point cloud or a mesh) using a deformation graph approach [59]. Rosinol *et al.* [50] unify the two steps into a single pose graph and mesh optimization. None of these works is concerned with simultaneously correcting multiple hierarchical representations.

III. REAL-TIME INCREMENTAL 3D SCENE GRAPH LAYERS CONSTRUCTION

This section describes how to construct the layers of a 3D scene graph given an odometric estimate of the robot trajectory (*e.g.*, from visual-inertial odometry). Then, Section IV discusses how to correct the graph in response to loop closures.

We focus on indoor environments and adopt the 3D scene graph model introduced in [49] and visualized in Fig. 1 and Fig. 6. In this model, *Layer 1* is a metric-semantic 3D mesh. *Layer 2* is a subgraph of objects and agents; each object has a semantic label, a centroid, and a bounding box, while each agent is modeled by a pose graph describing its trajectory (in our case the robot itself is the only agent). *Layer 3* is a subgraph of *places* (essentially, a topological map) where each

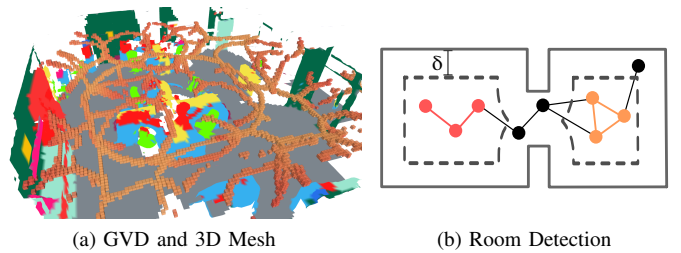


Fig. 2. (a) GVD (orange blocks) and mesh inside the active window. (b) Room detection: connected components (in orange and red) in the subgraph of places induced by a dilation distance δ (walls are shown in gray, dilated walls are dashed, places that disappear after dilation are in black).

place is an obstacle-free location and an edge between places denotes straight-line traversability. *Layer 4* is a subgraph of rooms where each room has a centroid, and edges connect adjacent rooms. *Layer 5* is a building node connected to all rooms (we assume the robot maps a single building). Edges connect nodes within each layer (*e.g.*, to model traversability between places or rooms) or across layers (*e.g.*, to model that mesh vertices belong to an object, that an object is in a certain room, or that a room belongs to a building).

Next, we present an approach to construct Layers 1-3 (Section III-A) and to segment places into rooms (Section III-B).

A. Layers 1-3: Mesh, Objects, and Places

Mesh and Objects. The real-time construction of the metric-semantic 3D mesh (*Layer 1* of the 3D scene graph) is an extension of Kimera [50], with small but important modifications. Kimera [50] uses Voxblox [43] to integrate semantically-labeled point clouds into a Truncated Signed Distance Field (TSDF) and an ESDF of the environment, while also performing Bayesian inference over the semantic label of each voxel.

Contrarily to Kimera [50] and building on the implementation in [43], we spatially window the TSDF and ESDF and only form a volumetric model of the robot’s surroundings within a user-specified radius (8m in our implementation); the radius is chosen to bound the amount of memory used by the ESDF. Within this “active window” we extract the 3D metric-semantic mesh using Voxblox’ marching cubes implementation and the places (as described in the next paragraph); as the mesh and places move outside the active window, they are passed to the Scene Graph Frontend (Section IV). We also modify the marching cubes algorithm to label the TSDF voxels that correspond to zero-crossings (*i.e.*, the voxels containing a surface); we call these voxels “parents” and keep track of the corresponding mesh vertices. Then, for each ESDF voxel—which already stores a distance to the closest obstacle—we additionally keep track of which parent is closest to the voxel. When extracting the places from the ESDF we use the parents to associate each place to the closest vertex in the 3D mesh.

After extracting the 3D mesh within the active window, we segment objects (*Layer 2* of the 3D scene graph) by performing Euclidean clustering of the 3D metric-semantic mesh vertices; in particular, we independently cluster vertices belonging to each semantic class. As in Kimera [50], the result

of the Euclidean clustering is then used to estimate a centroid and bounding box for each putative object. During incremental operation, if a putative object overlaps² with an existing object node of the same semantic class in the scene graph, we merge them together by adding new mesh vertices to the previous object node; if the new object does not correspond to an existing object node it is added as a new node.

Places. Kimera [50] builds a monolithic ESDF of the environment and then uses [44] to extract the places subgraph (*Layer 3* in the 3D scene graph). We instead implement an approach that incrementally extracts the subgraph of places using a Generalized Voronoi Diagram (GVD, shown in Figure 2a) built on the fly during the ESDF integration. The GVD is the set of voxels that are equidistant to at least 2 obstacles (“basis points” or “parents”), and intuitively forms a skeleton of the environment [44]. We obtain the GVD as a byproduct of the ESDF integration inside the active window, following the approach in [30]. In particular, the voxels belonging to the GVD can be easily detected from the wavefronts of the brushfire algorithm used to update the ESDF.

After the GVD of the active window is computed, we incrementally sparsify the GVD into a subgraph of places by modifying the batch approach of [44]. Intuitively, we select a subset of GVD voxels to become places nodes and connect them with edges to form the graph of places. After each update to the GVD, we iterate through each new voxel member of the GVD with enough basis points (3 in our implementation) to create a node or edge. Voxels are considered nodes if the new GVD voxel either has enough basis points (*i.e.*, ≥ 4) or if the neighborhood of the voxel matches a template proposed by [44] to identify corner voxels. To identify edges between nodes, we alternate between two phases. First, we label GVD voxels with the nearest node ID via flood-fill, starting from the labels produced from the previous ESDF integration to generate a putative set of edges from all neighboring node IDs. As a second phase, we split putative edges where the straight-line edge deviates too far from the GVD voxels connecting the two nodes by inserting a new node at the GVD voxel with the maximum deviation. During the first phase (flood-fill), we also merge nearby nodes. After a fixed number of iterations of the two phases, we then add the identified edges to the sparse graph, and remove any disconnected nodes. Finally, we add extra edges between disconnected components inside the active window to make the subgraph of places connected.

B. Layer 4: Room Detection

The room detection approach in Kimera [50] requires a volumetric representation of the entire environment and makes assumptions on the room geometry (*e.g.*, ceiling height) that do not easily extend to arbitrary (and possibly multi-story) buildings. To resolve these issues, we present a novel approach for constructing *Layer 4* of the 3D scene graph that segments rooms directly from the sparse subgraph of places.

²In practice, we consider two objects overlapping if the centroid of one object is contained in the other object’s bounding box, a proxy for spatial overlap measures such as Intersection over Union (IoU).

The subgraph of places, that we denote as \mathcal{G}_p , can be the one produced by the approach in Section III-A, or the optimized one computed after loop closures, as described in Section IV.

Our approach is based on two key insights. The first is that dilation operations on the voxel-based map help expose rooms in the environment: if we inflate obstacles, small apertures in the environment (*i.e.*, doors) will gradually close, naturally partitioning the voxel-based map into disconnected components (*i.e.*, rooms). The second insight is that each node in our place subgraph \mathcal{G}_p stores a distance to its closest obstacle (Section III-A); therefore, dilation operations in the voxel-based map can be directly mapped into topological changes in \mathcal{G}_p . More precisely, if we dilate the map by a distance δ , every place with obstacle distance smaller than δ will disappear from the graph (since it will no longer be in the free space). A visualization of this idea is given in Fig. 2b.

These insights motivate our approach for room detection. We dilate the map by increasing distances δ (*e.g.*, 10 distances uniformly spaced in [0.45, 1.2]m). For each dilation distance, we prune the subgraph of places by discarding nodes with distance smaller than δ (and their edges); we call the pruned subgraph $\mathcal{G}_{p,\delta}$. We count the number of connected components in $\mathcal{G}_{p,\delta}$ (intuitively, for a suitable choice of δ , the connected components will correspond to the rooms in the environment).³ Then we compute the median number of connected components n_r (to gain robustness to the choice of δ) and select the largest \mathcal{G}_{p,δ^*} that has n_r connected components. Finally, since \mathcal{G}_{p,δ^*} might miss some of the nodes in the original graph \mathcal{G}_p , we assign these unlabeled nodes via a partially seeded clustering technique. In particular, we use a greedy modularity-based community detection approach from [8] that involves iteratively attempting to assign each node in the graph to a community (*i.e.*, a room) that would result in the largest increase in modularity. We seed the initial communities to be the detected connected components in \mathcal{G}_{p,δ^*} and only iterate through unlabeled nodes. This both produces qualitatively consistent results and scales significantly better than related techniques (*e.g.*, spectral clustering is a popular method for clustering graphs but requires a more expensive eigen-decomposition of the Laplacian of \mathcal{G}_p).

Our novel room detection method provides two advantages. First, reasoning over 3D free-space instead of a 2D occupancy grid (*e.g.*, [27]) reduces the impact of clutter in the environment. Second, the parameters for the method can be set to work for a variety of environments (*i.e.*, the two main parameters are the minimum and maximum opening size between rooms). At the same time, our room detection only reasons over the topology of the environment; as such it will fail to segment semantically distinct rooms in an open floor-plan.

³In practice, we combine the dilation and the connected component computation, *i.e.*, we never explicitly compute the dilated subgraph and instead just prevent nodes and edges from being visited during the breadth-first search for connected components if they fall below the distance threshold δ .

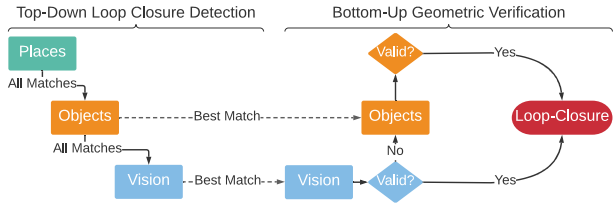


Fig. 3. Loop closure detection (left) and geometric verification (right). To find a match, we “descend” the 3D scene graph layers, comparing descriptors. We then “ascend” the 3D scene graph layers, attempting registration.

IV. PERSISTENT REPRESENTATIONS: LOOP CLOSURE DETECTION AND 3D SCENE GRAPH OPTIMIZATION

While the previous section describes how to incrementally build the layers of an “odometric” 3D scene graph (the layers are built from current odometry estimates), this section describes how to detect loop closures (Section IV-A) and how to correct the scene graph after loop closures (Section IV-B).

A. Loop Closure Detection and Geometric Verification

We augment visual loop closure detection and geometric verification by using multiple layers in the scene graph.

Top-down Loop Closure Detection. The agent layer in the 3D scene graph stores a pose graph describing the robot trajectory; we refer to these poses as the *agent nodes*. In our implementation, we store a keyframe for each agent node, from which appearance information can be extracted. Loop closure detection then aims at finding a past agent node that matches (*i.e.*, observed the same portion of the scene seen by) the latest agent node (corresponding to the current robot pose).

For each agent node, we construct a hierarchy of descriptors describing statistics of the node’s surroundings, from low-level appearance to objects semantics and places geometry. The descriptors are computed only once, when a new agent node is instantiated. At the lowest level, our hierarchical descriptors include standard DBoW2 appearance descriptors [19]. We augment the appearance descriptor with an object-based descriptor and a place-based descriptor (computed from the objects and places within a radius from the agent node). The former is computed as the histogram of the object labels in the node’s surroundings, intuitively describing the set of nearby objects. The latter is computed as the histogram of the distances associated to each place in the node surroundings, intuitively describing the geometry of the map near the node. While computing the descriptors, we also keep track of the IDs of the objects and places in the agent node’s surroundings, which are used for geometric verification.

For loop closure detection we compare the hierarchical descriptor of the current (query) node with all the past agent node descriptors, searching for a match. When performing loop closure detection, we walk down the hierarchy of descriptors (from places, to objects, to appearance descriptors). In particular, when comparing the descriptors of two nodes, we compare the places descriptor and –if the descriptor distance is below a threshold– we move on to comparing object descriptors and then appearance descriptors. If any of the

descriptor comparisons return a putative match, we perform geometric verification; see Fig. 3 for a visual summary.

Bottom-up Geometric Verification. After we have a putative loop closure between our query and match agent nodes (say i and j), we attempt to compute a relative pose between the two by performing a bottom-up geometric verification. In particular, whenever we have a match at a given layer (*e.g.*, between appearance descriptors at the agent layer, or between object descriptors at the object layer), we attempt to register frames i and j . For registering visual features we use standard RANSAC-based geometric verification as in [48]. If that fails, we attempt registering objects using TEASER++ [69], discarding loop closures that also fail object registration. This bottom-up approach has the advantage that putative matches that fail appearance-based geometric verification (*e.g.*, due viewpoint or illumination changes) can successfully lead to valid loop closures during the object-based geometric verification. Section VI indeed shows that the proposed hierarchical descriptors improve the quality and quantity of detected loop closures.

B. 3D Scene Graph Optimization

In order to correct the 3D scene graph in response to a loop closure, the *Scene Graph Frontend* “assembles” the outputs of the modules described in Section III into a single 3D scene graph, and then the *Scene Graph Backend* (i) optimizes the graph using a deformation graph approach and (ii) post-processes the results to remove redundant subgraphs corresponding to the robot visiting the same location multiple times.

Scene Graph Frontend. The frontend builds an initial estimate of the 3D scene graph that is uncorrected for drift. More precisely, the frontend takes as input the result of the modules described in Section III): the latest mesh, places subgraph, objects, and pose graph of the agent (all windowed to a radius around the current robot pose). The corresponding nodes and edges are incrementally added to the 3D scene graph data structure (which stores the entire scene graph up to the current time). Then, the frontend populates inter-layer edges from each object or agent node to the nearest place node in the active window using nanoflann [7]. Finally, the frontend computes a subsampled version of the mesh that will be optimized in the deformation graph approach described below. The subsampled mesh is computed via a octree-based vertex clustering mesh simplification approach, resulting in a smaller subset of nodes (which we refer to as the *mesh control points*) and edges representing connectivity between nodes.

Scene Graph Backend. When a loop closure is detected, the backend optimizes an *embedded deformation graph* [59] built from the frontend scene graph and then reconstructs the other nodes in the scene graph via interpolation as in [59] (Fig. 4). More precisely, we form the deformation graph as the subgraph of the 3D scene graph that includes (i) the agent layer, consisting of a pose graph that includes both odometry and loop closures edges, (ii) the mesh control points and the corresponding edges, and (iii) the minimum spanning tree of the places layer. By construction, these layers form

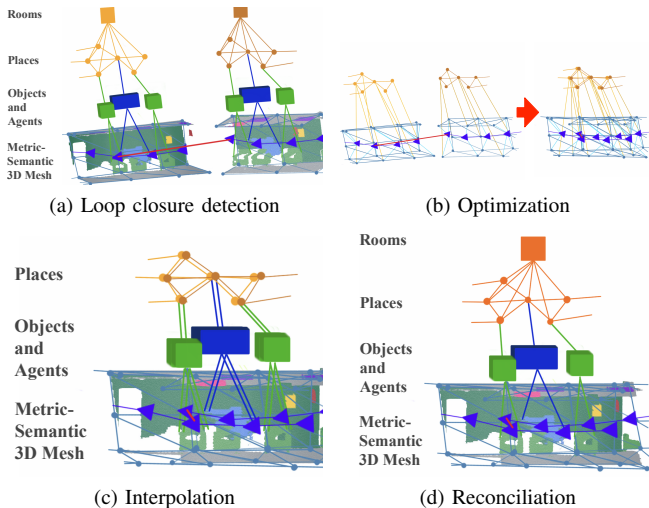


Fig. 4. Loop closure detection and optimization: (a) after a loop closure is detected, (b) we extract and optimize a subgraph of the 3D scene graph—the *deformation graph*—that includes the agent poses, the places, and a subset of the mesh vertices. (c) We then we reconstruct the rest of the graph via interpolation as in [59], and (d) reconcile overlapping nodes.

a connected subgraph via the inter-layer edges added by the frontend. The choice of using the minimum spanning tree of places is mostly motivated by computational reasons: the use of the spanning tree preserves the sparsity of the graph.

The embedded deformation graph approach associates a local frame (*i.e.*, a pose) to each node in the deformation graph and then solves an optimization problem to adjust the local frames in a way that minimizes deformations associated to each edge (including loop closures). In hindsight, this step transforms a subset of the 3D scene graph into a *factor graph* [11], where edge potentials need to be minimized. We refer the reader to [59] for the details about the optimization and note that we use the reformulation of the deformation graph with rigid transforms from [50] (instead of affine as in [59]) to obtain a standard pose graph optimization problem that is amenable to off-the-shelf solvers. In particular, we use the Graduated Non-Convexity (GNC) solver in GTSAM [2], which is also able to reject incorrect loop closures as outliers.

Once the optimization is finished, the place nodes are updated with their new positions and the full mesh is interpolated based on the deformation graph approach in [59]. We then recompute the object centroids and bounding boxes from the position of the corresponding vertices in the newly deformed mesh. During this update, overlapping nodes are also merged: for places nodes, we merge nodes within a distance threshold (0.4m in our implementation); for object nodes we merge nodes if the corresponding objects have the same semantic label and if one of nodes is contained inside the bounding box of the other node. We maintain a version of the scene graph where the nodes are not merged; this enables undoing wrong loop closures if an accepted loop closure is deemed to be an outlier by GNC later on. Finally, we re-detect rooms from the merged places using the approach described in Section III.

V. THINKING FAST AND SLOW: THE HYDRA ARCHITECTURE

We implement our Spatial Perception System into a highly parallelized architecture, named *Hydra*. Hydra involves a combination of processes that run at sensor rate (*e.g.*, feature tracking for visual-inertial odometry), at sub-second rate (*e.g.*, mesh and place reconstruction), and at slower rates (*e.g.*, the scene graph optimization, whose complexity depends on the map size). Therefore these processes have to be organized such that slow-but-infrequent computation (*e.g.*, scene graph optimization) does not get in the way of faster processes.

We visualize Hydra in Fig 5. Each block in the figure denotes an algorithmic module matching the discussion in the previous sections. Hydra starts with fast *early* perception processes (Fig 5, left), which perform low-level perception tasks such as feature detection and tracking (at frame-rate), 2D semantic segmentation, and stereo-depth reconstruction (at keyframe rate). The result of early perception processes are passed to mid-level perception processes (Fig 5, center). These include algorithms that incrementally construct (an odometric version of) the agent layer (*e.g.*, the visual-inertial odometry backend), the mesh and places layers, and the object layer. Mid-level perception also includes the scene graph frontend, which collects the result of the other modules into an “unoptimized” scene graph. Finally, the high-level perception processes perform loop closure detection, execute scene graph backend optimization, and perform room detection.⁴ This results in a globally consistent, persistent 3D scene graph.

Hydra runs in real-time on a multi-core CPU; the only module that relies on GPU computing is the 2D semantic segmentation, which uses a standard off-the-shelf deep network. Running on CPU has the advantage of (i) leaving the GPU to learning-oriented components, and (ii) being compatible with the power limitations imposed by current mobile robots.

VI. EXPERIMENTS

This section shows that Hydra builds 3D scene graphs in real-time with an accuracy comparable to batch offline methods.

A. Experimental Setup

Datasets. We utilize two datasets for our experiments: uHumans2 (uH2) [50] and SidPac. The uH2 dataset is a Unity-based simulated dataset [50] that includes three scenes: a small apartment, an office, and a subway station. The dataset provides visual-inertial data as well as ground-truth depth and 2D semantic segmentation. The dataset also provides ground truth robot trajectories that we use for benchmarking purposes.

The SidPac dataset is a real dataset collected in a graduate student housing building using a visual-inertial hand-held device. We used a Kinect Azure camera as the primary collection device with an Intel RealSense T265 rigidly attached to the Kinect to provide external odometry input. The dataset consists of two separate recordings. The first recording covers two

⁴While room detection can be performed quickly, it still operates on the entire graph, hence it is more suitable as a slow high-level perception process.

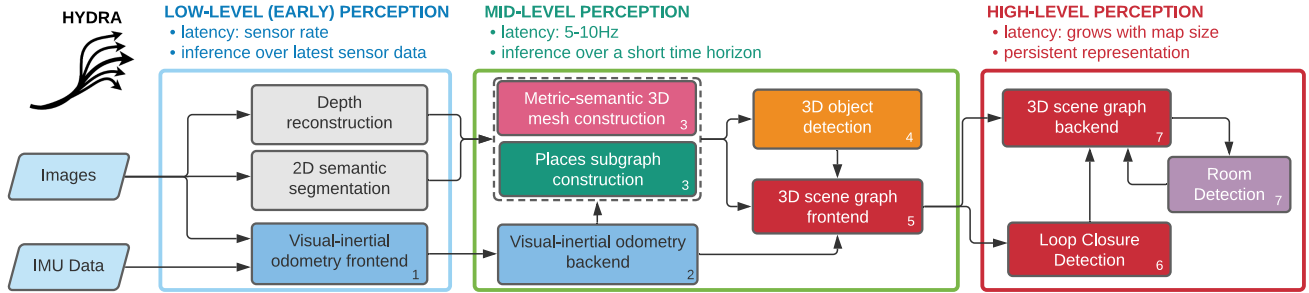


Fig. 5. Hydra’s functional blocks. We conceptualize three different functional block groupings: low-level perception, mid-level perception, and high-level perception in order of increasing latency. Each functional block is labeled with a number that identifies the “logical” thread that the module belongs to.

floors of the building (Floors 1 & 3), where we walked through a common room, a music room, and a recreation room on the first floor of the graduate residence, went up a stairwell, through a long corridor as well as a student apartment on the third floor, then finally down another stairwell to revisit the music room and the common room, ending where we started. The second recording also covers two floors (Floors 3 & 4), where we map student apartments as well as lounge and kitchen areas which are duplicated across both floors. These scenes are particularly challenging given the scale of the scenes (average traversal of around 400 meters), the prevalence of glass and strong sunlight in regions of the scenes (causing partial depth estimates from the Kinect), and feature-poor regions in hallways. We obtain a proxy for the ground-truth trajectory for both SidPac datasets via a hand-tuned pose graph optimization with additional height priors, to reduce drift and qualitatively match the building floor plans.

Hydra. For the real datasets, we use the depth reconstruction from the Kinect (*cf.* Fig. 5) and we use HRNet [64] for 2D semantic segmentation, using the pre-trained model from the MIT Scene Parsing challenge [72]. While newer and more performant networks exist (*e.g.*, [5, 15, 24]), few had pre-trained semantic segmentation models for ADE20k [72] and were compatible with our inference toolchain (ONNX and TensorRT). For the simulated datasets, we use the provided depth and segmentations. For both simulated and real datasets we use Kimera-VIO [48] for visual-inertial odometry, and in the real scenes we fuse the Kimera-VIO estimates with the output of the RealSense T265 to improve the quality of the odometric trajectory. All the remaining blocks in Fig. 5 are implemented in C++, following the approach described in this paper. In the experiments we use a workstation with an AMD Ryzen9 3960X with 24 cores and two Nvidia GTX3080s, though we also report timing results on an embedded computer (Nvidia Xavier NX) at the end of this section.

B. Results and Ablation Study

We present an extensive evaluation of the accuracy and runtime of our real-time approach against the batch offline scene graph construction approach from [50].

Accuracy Evaluation: Objects and Places. Fig. 7 evaluates the object and place layers by comparing three different configurations of Hydra. The first configuration (“*GT-*

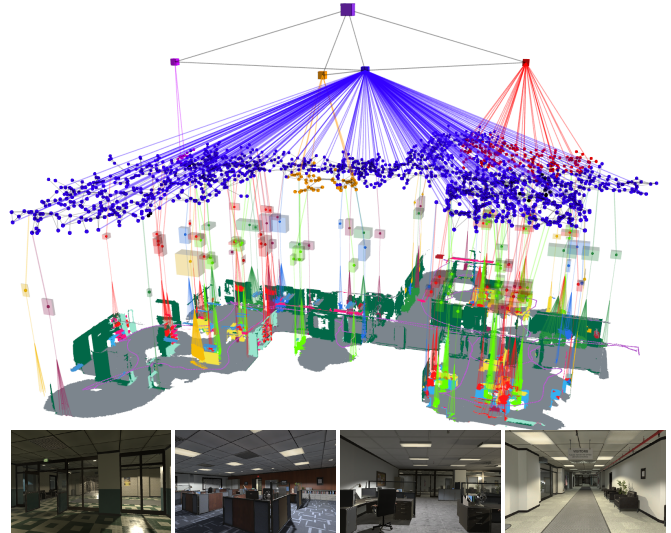


Fig. 6. 3D scene graph created by Hydra in the uH2 Office dataset.

Trajectory”) uses ground-truth poses to incrementally construct the scene graph. The second and third configurations (“*VIO+V-LC*” and “*VIO+SG-LC*” respectively) use visual-inertial odometry (VIO) for odometry estimation and then use vision-based loop closures (*VIO+V-LC*) or the proposed scene graph loop closures (*VIO+SG-LC*). For the objects and places evaluation, we consider the batch scene graph constructed from ground-truth poses using [50] as the ground-truth scene graph.

For the object layer, we report two metrics: the percentage of objects in the ground-truth scene graph that have an estimated object with the correct semantic label within a specified radius (“% *Found*”) and the percentage of objects in the estimated scene graph that have a ground-truth object with the correct semantic label within a specified radius (“% *Correct*”).⁵ For the places layer, we measure the mean distance of an estimated place node to the nearest voxel in the ground-truth GVD (“*Position Error*”).

We note some important trends in Fig. 7. First, when using *GT-Trajectory*, Hydra’s performance is close to the ground-truth scene graph (80-100% found and correct objects,

⁵As distance thresholds, we use 0.2m for the apartment scene, 0.3m for the office scene, 4m for the subway scene, 0.9m for floors 1&3 of SidPac and 1.4m for floors 3&4 of SidPac. These thresholds were chosen to roughly correspond to the mean Absolute Trajectory Error (ATE) for each scene, in order to normalize the metrics according to the environment size.

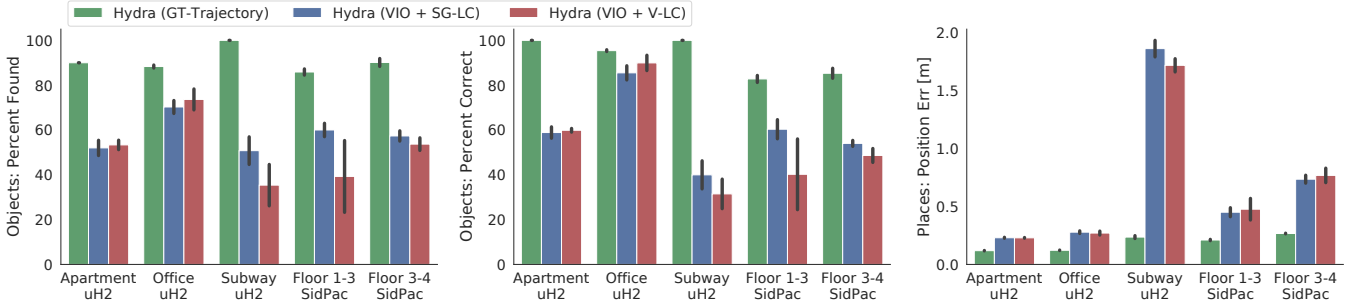


Fig. 7. Accuracy of the objects and places estimated by Hydra. Each plot reports the mean across 5 trials along with the standard deviation as an error bar.

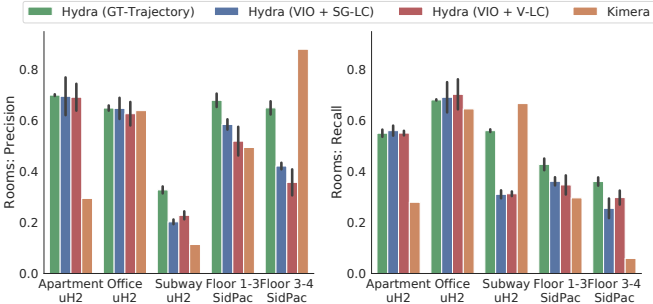


Fig. 8. Room detection accuracy of Hydra versus Kimera [50]. Each plot reports the mean across 5 trials along with the standard deviation as an error bar except Kimera, for which one trial is shown (due to its slower runtime).

sub-25cm places position error). This demonstrates that – given the trajectory– the real-time scene graph from Hydra is comparable to the batch and offline approaches at the state of the art. Second, *VIO+V-LC* and *VIO+SG-LC* maintain reasonable levels of accuracy for the objects and places and attain comparable performance in small to medium-sized scenes (e.g., Apartment, Office). In these scenes, the drift is small and the loop closure strategy does not radically impact performance (differences are within standard deviation, shown as black confidence bars). However, in larger scenes (e.g., SidPac) loop closures are more important and *VIO+SG-LC* substantially outperforms *VIO+V-LC* in terms of object accuracy. Importantly, using *VIO+SG-LC* typically leads to a reduction in the standard deviation of the results, confirming that the proposed approach for scene graph loop closure detection leads to more reliable loop closure results (more details and ablations below). The place positions errors remain similar for both *VIO+SG-LC* and *VIO+V-LC* and are larger in the subway dataset, which includes larger open-spaces with more distant nodes in the subgraph of places.

Accuracy Evaluation: Rooms. Fig. 8 evaluates the room detection performance, using the precision and recall metrics defined in [9] (here we compute precision and recall over 3D voxels instead of 2D pixels). More formally, these metrics are:

$$\begin{aligned}
 Precision &= \frac{1}{|R_e|} \sum_{r_e \in R_e} \max_{r_g \in R_g} \frac{|r_g \cap r_e|}{|r_e|} \\
 Recall &= \frac{1}{|R_g|} \sum_{r_g \in R_g} \max_{r_e \in R_e} \frac{|r_e \cap r_g|}{|r_g|}
 \end{aligned} \tag{1}$$

where R_e is the set of estimated rooms, R_g is the set of ground-truth rooms, and $|\cdot|$ returns the cardinality of a set; here, each room r_e (or r_g) is defined as a set of free-space voxels. We hand-label the ground-truth rooms R_g from the ground-truth reconstruction of the environment. In particular, we manually define sets of bounding boxes for each room and assign a unique (ground-truth) label to the free-space voxels falling within each room. For the estimated rooms R_e , we derive the free-space voxels from the places comprising each estimated room. In eq. (1), *Precision* then measures the maximum overlap in voxels with a ground-truth room for every estimated room, and *Recall* measures the maximum overlap in voxels with an estimated room for every ground-truth room. Intuitively, low precision corresponds to under-segmentation, *i.e.*, fewer and larger room estimates, and low recall corresponds to over-segmentation, *i.e.*, more and smaller room estimates. For benchmarking purposes, we also include the approach in [50] (*Kimera*) as a baseline for evaluation.

Fig. 8 shows that while *Kimera* [50] estimates a room segmentation with similar precision and recall to Hydra (*GT-Trajectory*) for the Office scene (the only single-floor scene), Hydra excels in multi-floor environments. For the split-level Apartment scene, we achieve significantly higher precision and recall as compared to *Kimera*. For SidPac Floor 3-4, the difference is particularly dramatic, where the *Kimera* achieves 0.88 precision but only 0.06 recall, as it only is able to segment 2 of 10 ground-truth rooms. In general, our approach estimates a room segmentation with consistent precision and recall (if slightly over-segmented), while *Kimera* oscillates between either low-precision and high-recall estimates (*i.e.*, extreme under-segmentation), or high-precision and low-recall estimates (*i.e.*, where it fails to segment most of the rooms). These differences stem from the difficulty of setting an appropriate height to attempt to segment rooms at for *Kimera*. Finally, it is worth noting that our room segmentation approach is able to mostly maintain the same levels of precision and recall for *VIO+SG-LC* and *VIO+V-LC* despite drift.

Runtime Evaluation. Fig. 9 reports the runtime of Hydra versus the batch approach in [50]. This plot shows that the runtime of the batch approach increases over time and takes more than 40 seconds to generate the entire scene graph for moderate scene sizes; as we mentioned, most processes in the batch approach [50] entail processing the entire ESDF

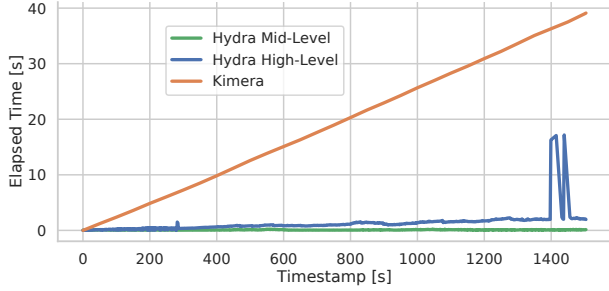


Fig. 9. Runtime required for scene graph construction vs. timestamp for the SidPac Floor 1-3 dataset for a batch approach (Kimera) and for the proposed incremental approach (Hydra). For timing of the low-level processes in Hydra, we refer the reader to the analysis in [50], as we also rely on Kimera-VIO.

		Layer		
		Objects [ms]	Places [ms]	Rooms [ms]
uH2	Apartment	32.4 ± 12.9	5.3 ± 1.4	4.4 ± 2.1
	Office	24.1 ± 12.8	8.1 ± 1.3	19.0 ± 12.3
	Subway	9.8 ± 9.3	5.9 ± 0.7	16.5 ± 10.6
SP	Floor 1-3	50.4 ± 30.3	3.4 ± 1.0	11.4 ± 14.4
	Floor 3-4	75.3 ± 37.0	4.2 ± 2.1	15.0 ± 14.6

TABLE I
HYDRA: TIMING BREAKDOWN

(e.g., place extraction and room detection), inducing a linear increase in the runtime as the ESDF grows. On the other hand, our scene graph frontend (*Hydra Mid-Level* in Fig. 9) has a fixed computation cost. In Fig. 9, a slight upward trend is observable for *Hydra High-Level*, driven by room detection and scene graph optimization computation costs, though remaining much lower than batch processing. Noticeable spikes in the runtime for *Hydra High-Level* (e.g., at 1400 seconds) correspond to the execution of the 3D scene graph optimization when new loop closures are added.

Table I reports timing breakdown for the incremental creation of each layer across scenes for a single trial. The object layer runtime is determined by the number of mesh vertices in the active window with object semantic class; hence the SidPac scenes have a higher computation cost than the other scenes. The room layer runtime is determined by the number of places (a combination of how complicated and large the scene is); this is why the Office has the largest computation cost for the rooms despite being smaller than the SidPac scenes.

While the timing results in Table I are obtained with a relatively powerful workstation, here we restate that Hydra can run in real-time on embedded computers commonly used in robotics applications. Towards this goal, we also measure timing statistics on an Nvidia Xavier NX for the uHumans2 Office scene. Hydra processes the objects in 75 ± 35 ms, the places in 33 ± 6 ms and rooms in 55 ± 41 ms. Note that our target “real-time” rate for these layers is keyframe rate (5 Hz). While there is still margin to optimize computation (see conclusions), these initial results stress the practicality and real-time capability of Hydra in building 3D scene graphs.

Loop Closure Ablation Study. Finally, we take a closer look at the quality of the loop closures candidates proposed by

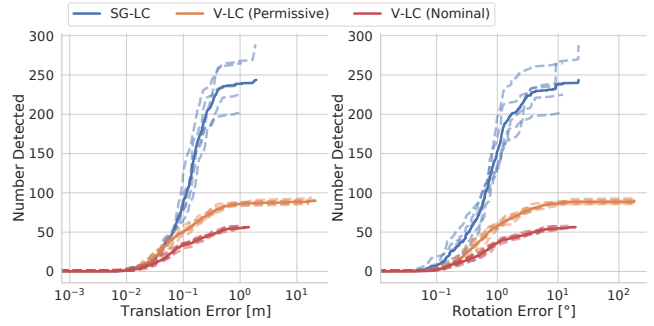


Fig. 10. Number of detected loop closures versus error of the estimated loop closure pose for three different loop closure detection configurations. Five individual trials and a trend-line are shown for each configuration.

our hierarchical loop closure detection approach, and compare it against traditional vision-based approaches on the Office scene. In particular, we compare our approach against a vision-based loop closure detection that uses DBoW2 for place recognition and ORB feature matching, as described in [50].

Fig. 10 shows the number of detected loop closures against the error of the registered solution (i.e., the relative pose between query and match computed by the geometric verification) for three different loop closure configurations: (i) “SG-LC”: the proposed scene graph loop closure detection, (ii) “V-LC (Nominal)”: a traditional vision-based loop closure detection with nominal parameters (same as Fig. 7), and (iii) “V-LC (Permissive)”: a vision-based loop closure detection with more permissive parameters (i.e., a decreased score threshold and less restrictive geometric verification settings). We report key parameters used in this evaluation in the appendix. As expected, making the vision-based detection parameters more permissive leads to more but lower-quality loop closures. On the other hand, the scene graph loop closure approach produces approximately twice as many loop closures within 10cm of error and 1 degree of error as the permissive vision-based approach. The proposed approach produces quantitatively and qualitatively better loop closures compared to both baselines.

VII. CONCLUSIONS

This paper introduces Hydra, a *real-time Spatial Perception System* that builds a 3D scene graph from sensor data in real-time. Hydra runs at sensor rate thanks to the combination of novel online algorithms and a highly parallelized perception architecture. Moreover, it allows building a *persistent* representation of the environment thanks to a novel approach for 3D scene graph optimization. While we believe the proposed approach constitutes a substantial step towards high-level 3D scene understanding for robotics, Hydra can be improved in many directions. First, some nodes in the reconstructed 3D scene graph are unlabeled (e.g., the algorithms in this paper are able to detect rooms, but are unable to label a given room as a “kitchen” or a “bedroom”); future work includes bridging Hydra with learning-based methods for 3D scene graph node labeling [60]. Second, it would be interesting to label nodes and edges of the 3D scene graph with a richer set of relations and affordances, building on [67]. Third, the

connections between our scene graph optimization approach and pose graph optimization offer opportunities to improve the efficiency of the optimization by leveraging recent advances in pose graph sparsification. Finally, the implications of using 3D scene graphs for prediction, planning, and decision-making are mostly unexplored (see [1, 46] for early examples), which opens further avenues for future work.

DISCLAIMER

Research was sponsored by the United States Air Force Research Laboratory and the United States Air Force Artificial Intelligence Accelerator and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

ACKNOWLEDGMENTS

This work was partially funded by the AIA CRA FA8750-19-2-1000, ARL DCIST CRA W911NF-17-2-0181, and ONR RAIDER N00014-18-1-2828.

REFERENCES

- [1] C. Agia, J. Krishna Murthy and M. Khodeir, O. Miksik, V. Vineet, M. Mukadam, L. Paull, and F. Shkurti. Taskography: Evaluation robot task planning over large 3D scene graphs. In *Conference on Robot Learning (CoRL)*, 2021.
- [2] P. Antonante, V. Tzoumas, H. Yang, and L. Carlone. Outlier-robust estimation: Hardness, minimally tuned algorithms, and applications. *IEEE Trans. Robotics*, 38(1): 281–301, 2021. ([pdf](#)).
- [3] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5297–5307, 2016.
- [4] I. Armeni, Z. He, J. Gwak, A. Zamir, M. Fischer, J. Malik, and S. Savarese. 3D scene graph: A structure for unified semantics, 3D space, and camera. In *Intl. Conf. on Computer Vision (ICCV)*, pages 5664–5673, 2019.
- [5] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers, 2021.
- [6] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *Intl. Conf. on Computer Vision (ICCV)*, 2019.
- [7] Jose Luis Blanco and Pranjal Kumar Rai. nanoflann: a C++ header-only fork of FLANN, a library for nearest neighbor (NN) with kd-trees. <https://github.com/jlblancoc/nanoflann>, 2014.
- [8] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008, 2008.
- [9] Richard Bormann, Florian Jordan, Wenzhe Li, Joshua Hampp, and Martin Hagele. Room segmentation: Survey, implementation, and analysis. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1019–1026, 2016. doi: 10.1109/ICRA.2016.7487234.
- [10] S.L. Bowman, N. Atanasov, K. Daniilidis, and G.J. Pappas. Probabilistic data association for semantic SLAM. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1722–1729, 2017.
- [11] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J.J. Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Trans. Robotics*, 32(6):1309–1332, 2016. ISSN 1552-3098. doi: 10.1109/TRO.2016.2624754. arxiv preprint: 1606.05830, ([pdf](#)).
- [12] R. Chatila and J.-P. Laumond. Position referencing and consistent world modeling for mobile robots. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 138–145, 1985.
- [13] Angela Dai, Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Christian Theobalt. Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration. *ACM Transactions on Graphics (ToG)*, 36(4):1, 2017.
- [14] J. Dong, X. Fei, and S. Soatto. Visual-Inertial-Semantic scene representation for 3D object detection. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [15] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows, 2021.
- [16] Stephen Friedman, Hanna Pasula, and Dieter Fox. Voronoi random fields: Extracting the topological structure of indoor environments via place labeling. In *Intl. Joint Conf. on AI (IJCAI)*, page 2109–2114, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.
- [17] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski. Reconstructing building interiors from images. In *Intl. Conf. on Computer Vision (ICCV)*, 2009.
- [18] C. Galindo, A. Saffiotti, S. Coradeschi, P. Buschka, J.A. Fernández-Madrigal, and J. González. Multi-hierarchical semantic maps for mobile robotics. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 3492–3497, 2005.
- [19] Dorian Gálvez-López and J. D. Tardós. Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on Robotics*, 28(5):1188–1197, October 2012. ISSN 1552-3098. doi: 10.1109/TRO.2012.2197158.

- [20] Sourav Garg and Michael Milford. Seqnet: Learning descriptors for sequence-based hierarchical place recognition. *IEEE Robotics and Automation Letters*, 6(3): 4305–4312, 2021. doi: 10.1109/LRA.2021.3067633.
- [21] Abel Gawel, Carlo Del Don, Roland Siegwart, Juan Nieto, and Cesar Cadena. X-view: Graph-based semantic multi-view localization. *IEEE Robotics and Automation Letters*, 3(3):1687–1694, 2018. doi: 10.1109/LRA.2018.2801879.
- [22] N. Gothoskar, M. Cusumano-Towner, B. Zinberg, M. Ghavamizadeh, F. Pollok, A. Garrett, J.B. Tenenbaum, D. Gutfreund, and V.K. Mansinghka. 3DP3: 3D scene perception via probabilistic programming. In *ArXiv preprint: 2111.00312*, 2021.
- [23] M. Grinvald, F. Furrer, T. Novkovic, J. J. Chung, C. Cadena, R. Siegwart, and J. Nieto. Volumetric Instance-Aware Semantic Mapping and 3D Object Discovery. *IEEE Robotics and Automation Letters*, 4(3):3037–3044, 2019.
- [24] Shihua Huang, Zhichao Lu, Ran Cheng, and Cheng He. FaPN: Feature-aligned pyramid network for dense image prediction. In *International Conference on Computer Vision (ICCV)*, 2021.
- [25] Gregory Izatt and Russ Tedrake. Scene understanding and distribution modeling with mixed-integer scene parsing. In *technical report (under review)*, 2021.
- [26] U. Kim, J. Park, T. Song, and J. Kim. 3-D scene graph: A sparse and semantic representation of physical environments for intelligent agents. *IEEE Trans. Cybern.*, PP:1–13, Aug. 2019.
- [27] Alexander Kleiner, Rodrigo Baravalle, Andreas Kolling, Pablo Pilotti, and Mario Munich. A solution to room-by-room coverage for autonomous cleaning robots. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5346–5352, 2017. doi: 10.1109/IROS.2017.8206429.
- [28] B.J. Kuipers. Modeling spatial knowledge. *Cognitive Science*, 2:129–153, 1978.
- [29] B.J. Kuipers. The Spatial Semantic Hierarchy. *Artificial Intelligence*, 119:191–233, 2000.
- [30] Boris Lau, Christoph Sprunk, and Wolfram Burgard. Efficient grid-based spatial representations for robot navigation in dynamic environments. *Robot. Auton. Syst.*, 61(10):1116–1130, October 2013. ISSN 0921-8890. doi: 10.1016/j.robot.2012.08.010. URL <https://doi.org/10.1016/j.robot.2012.08.010>.
- [31] C. Li, H. Xiao, K. Tateno, F. Tombari, N. Navab, and G. D. Hager. Incremental scene understanding on dense SLAM. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 574–581, 2016.
- [32] K. Lianos, J. Schönberger, M. Pollefeys, and T. Sattler. Vso: Visual semantic odometry. In *European Conf. on Computer Vision (ECCV)*, pages 246–263, 2018.
- [33] Shiqi Lin, Jikai Wang, Meng Xu, Hao Zhao, and Zonghai Chen. Topology aware object-level semantic mapping towards more robust loop closure. *IEEE Robotics and Automation Letters*, 6(4):7041–7048, 2021. doi: 10.1109/LRA.2021.3097242.
- [34] Chen Liu, Jiaye Wu, and Yasutaka Furukawa. FloorNet: A unified framework for floorplan reconstruction from 3d scans. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [35] Yu Liu, Yvan Petillot, David Lane, and Sen Wang. Global Localization with Object-Level Semantics and Topology. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 4909–4915, May 2019. doi: 10.1109/ICRA.2019.8794475. ISSN: 2577-087X.
- [36] S. Lowry, N. Sünderhauf, P. Newman, J. Leonard, D. Cox, P. Corke, and M. Milford. Visual place recognition: A survey. *IEEE Trans. Robotics*, 32(1):1–19, 2016.
- [37] R. Lukierski, S. Leutenegger, and A. J. Davison. Room layout estimation from rapid omnidirectional exploration. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 6315–6322, 2017.
- [38] J. McCormac, A. Handa, A. J. Davison, and S. Leutenegger. SemanticFusion: Dense 3D Semantic Mapping with Convolutional Neural Networks. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2017.
- [39] J. McCormac, R. Clark, M. Bloesch, A.J. Davison, and S. Leutenegger. Fusion++: Volumetric object-level SLAM. In *Intl. Conf. on 3D Vision (3DV)*, pages 32–41, 2018.
- [40] G. Narita, T. Seno, T. Ishikawa, and Y. Kaji. Panopticfusion: Online volumetric semantic mapping at the level of stuff and things. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [41] L. Nicholson, M. Milford, and N. Sünderhauf. Quadric-SLAM: Dual quadrics from object detections as landmarks in object-oriented SLAM. *IEEE Robotics and Automation Letters*, 4:1–8, 2018.
- [42] Kyel Ok, Katherine Liu, and Nicholas Roy. Hierarchical object map estimation for efficient and robust navigation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1132–1139, 2021. doi: 10.1109/ICRA48506.2021.9561225.
- [43] Helen Oleynikova, Zachary Taylor, Marius Fehr, Roland Siegwart, and Juan Nieto. Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 1366–1373. IEEE, 2017.
- [44] Helen Oleynikova, Zachary Taylor, Roland Siegwart, and Juan Nieto. Sparse 3D topological graphs for micro-aerial vehicle planning. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2018.
- [45] Cao Qin, Yunzhou Zhang, Yingda Liu, and Guanghao Lv. Semantic loop closure detection based on graph matching in multi-objects scenes. *Journal of Visual Communication and Image Representation*, 76:103072, 2021.
- [46] Z. Ravichandran, L. Peng, N. Hughes, J.D. Griffith, and L. Carlone. Hierarchical representations and explicit memory: Learning effective navigation policies on 3D

- scene graphs using graph neural networks. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2022. (pdf).
- [47] V. Reijgwart, A. Millane, H. Oleynikova, R. Siegwart, C. Cadena, and J. Nieto. Voxgraph: Globally consistent, volumetric mapping using signed distance function submaps. *IEEE Robotics and Automation Letters*, 2020.
- [48] A. Rosinol, M. Abate, Y. Chang, and L. Carlone. Kimera: an open-source library for real-time metric-semantic localization and mapping. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2020. doi: 10.1109/ICRA40945.2020.9196885. arXiv preprint arXiv:1910.02490, (video), (code), (pdf).
- [49] A. Rosinol, A. Gupta, M. Abate, J. Shi, and L. Carlone. 3D dynamic scene graphs: Actionable spatial perception with places, objects, and humans. In *Robotics: Science and Systems (RSS)*, 2020. doi: 10.15607/RSS.2020.XVI.079. URL <http://news.mit.edu/2020/robots-spatial-perception-0715>. (pdf), (media), (video).
- [50] A. Rosinol, A. Violette, M. Abate, N. Hughes, Y. Chang, J. Shi, A. Gupta, and L. Carlone. Kimera: from SLAM to spatial perception with 3D dynamic scene graphs. *Intl. J. of Robotics Research*, 40(12–14):1510–1546, 2021. arXiv preprint arXiv:2101.06894, (pdf).
- [51] R. Rosu, J. Quenzel, and S. Behnke. Semi-supervised semantic mapping through label propagation with semantic texture meshes. *Intl. J. of Computer Vision*, 06 2019.
- [52] Jose-Raul Ruiz-Sarmiento, Cipriano Galindo, and Javier Gonzalez-Jimenez. Building multiversal semantic maps for mobile robot operation. *Knowledge-Based Systems*, 119:257–272, 2017.
- [53] R. F. Salas-Moreno, R. A. Newcombe, H. Strasdat, P. H. J. Kelly, and A. J. Davison. SLAM++: Simultaneous localisation and mapping at the level of objects. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [54] Lukas Schmid, Jeffrey Delmerico, Johannes Schönberger, Juan Nieto, Marc Pollefeys, Roland Siegwart, and Cesar Cadena. Panoptic multi-tdsfs: a flexible representation for online multi-resolution volumetric mapping and long-term dynamic scene consistency. *arXiv preprint arXiv:2109.10165*, 2021.
- [55] Stefan Schubert, Peer Neubert, and Peter Protzel. Fast and memory efficient graph optimization via ICM for visual place recognition. In *Proc. of Robotics: Science and Systems (RSS)*, 2021.
- [56] M. Shan, Q. Feng, and N. Atanasov. Object residual constrained visual-inertial odometry. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 5104–5111, 2020.
- [57] Sinisa Stekovic, Mahdi Rad, Friedrich Fraundorfer, and Vincent Lepetit. MonteFloor: Extending MCTS for reconstructing accurate large-scale floor plans, 2021.
- [58] J. Stückler and S. Behnke. Multi-resolution surfel maps for efficient dense 3d modeling and tracking. *J. Vis. Commun. Image Represent.*, 25(1):137–147, 2014.
- [59] R. Sumner, J. Schmid, and M. Pauly. Embedded deformation for shape manipulation. *ACM SIGGRAPH 2007 papers on - SIGGRAPH '07*, 2007. doi: 10.1145/1275808.1276478.
- [60] R. Talak, S. Hu, L. Peng, and L. Carlone. Neural trees for learning on graphs. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021. (pdf).
- [61] K. Tateno, F. Tombari, and N. Navab. Real-time and scalable incremental segmentation on dense SLAM. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 4465–4472, 2015.
- [62] S. Thrun. Robotic mapping: a survey. In *Exploring artificial intelligence in the new millennium*, pages 1–35. Morgan Kaufmann, Inc., 2003.
- [63] Johanna Wald, Helisa Dhama, Nassir Navab, and Federico Tombari. Learning 3D semantic scene graphs from 3D indoor reconstructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3961–3970, 2020.
- [64] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(10):3349–3364, 2021. doi: 10.1109/TPAMI.2020.2983686.
- [65] T. Whelan, M. Kaess, H. Johannsson, M. Fallon, J. Leonard, and J. McDonald. Real-time large-scale dense RGB-D SLAM with volumetric fusion. *Intl. J. of Robotics Research*, 34(4–5):598–626, 2015.
- [66] T. Whelan, R.F. Salas-Moreno, B. Glocker, A. J. Davison, and S. Leutenegger. ElasticFusion: Real-Time Dense SLAM and Light Source Estimation. 2016.
- [67] S. Wu, J. Wald, K. Tateno, N. Navab, and F. Tombari. SceneGraphFusion: Incremental 3D scene graph prediction from RGB-D sequences. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [68] B. Xu, W. Li, D. Tzoumanikas, M. Bloesch, A. Davison, and S. Leutenegger. MID-Fusion: Octree-based object-level multi-instance dynamic SLAM. pages 5231–5237, 2019.
- [69] H. Yang, J. Shi, and L. Carlone. TEASER: Fast and Certifiable Point Cloud Registration. *IEEE Trans. Robotics*, 37(2):314–333, 2020. extended arXiv version 2001.07715 (pdf).
- [70] H. Zender, O. Martínez Mozos, P. Jensfelt, G.-J.M. Kruijff, and W. Burgard. Conceptual spatial representations for indoor mobile robots. *Robotics and Autonomous Systems*, 56(6):493–502, 2008. From Sensors to Human Spatial Concepts.
- [71] Tian Zheng, Guoqing Zhang, Lei Han, Lan Xu, and Lu Fang. Building fusion: Semantic-aware structural building-scale 3d reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020. doi: 10.1109/TPAMI.2020.3042881.
- [72] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela

Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

APPENDIX

Loop Closure Ablation Parameters. Table II reports key parameters used for “V-LC (Permissive)” and “V-LC (Nominal)” in the loop closure ablation study in Fig. 10. For the meaning of each parameter we refer the reader to the open-source Kimera implementation released by [50].

Parameter	V-LC (Permissive)	V-LC (Nominal)
L1 Score Threshold	0.05	0.4
Minimum NSS	0.005	0.05
Min. RANSAC Correspondences	12	15
5pt RANSAC Inlier Threshold	0.01	0.001
Lowe Matching Ratio	0.8	0.9

TABLE II
VISUAL LOOP CLOSURE PARAMETERS

Table III reports key parameters used for “SG-LC” in the ablation study in Fig. 10. All scene graph descriptors were computed with a radius of 13 meters. SG-LC does not use NSS (Normalized Similarity Scoring) to filter out matches.

Parameter	SG-LC
Agent L1 Match Threshold	0.01
Object L1 Match Threshold	0.3
Places L1 Match Threshold	0.5
Min. RANSAC Correspondences	15
5pt RANSAC Inlier Threshold	0.001
Object L1 Registration Threshold	0.8
Object Minimum Inliers	5
TEASER Noise Bound [m]	0.1

TABLE III
SCENE GRAPH LOOP CLOSURE PARAMETERS