This article has been accepted for publication in the IEEE Robotics and Automation Letters. Please cite as: Jared Strader, Nathan Hughes, William Chen, Alberto Speranzon, Luca Carlone, "Indoor and Outdoor 3D Scene Graph Generation via Language-Enabled Spatial Ontologies", in *IEEE Robotics and Automation Letters*, 2024.

Indoor and Outdoor 3D Scene Graph Generation via Language-Enabled Spatial Ontologies

Jared Strader[®], Nathan Hughes[®], William Chen[®], Alberto Speranzon[®], and Luca Carlone[®]

Abstract—This paper proposes an approach to build 3D scene graphs in arbitrary indoor and outdoor environments. Such extension is challenging; the hierarchy of concepts that describe an outdoor environment is more complex than for indoors, and manually defining such hierarchy is time-consuming and does not scale. Furthermore, the lack of training data prevents the straightforward application of learning-based tools used in indoor settings. To address these challenges, we propose two novel extensions. First, we develop methods to build a spatial ontology defining concepts and relations relevant for indoor and outdoor robot operation. In particular, we use a Large Language Model (LLM) to build such an ontology, thus largely reducing the amount of manual effort required. Second, we leverage the spatial ontology for 3D scene graph construction using Logic Tensor Networks (LTN) to add logical rules, or axioms (e.g., "a beach contains sand"), which provide additional supervisory signals at training time thus reducing the need for labelled data, providing better predictions, and even allowing predicting concepts unseen at training time. We test our approach in a variety of datasets, including indoor, rural, and coastal environments, and show that it leads to a significant increase in the quality of the 3D scene graph generation with sparsely annotated data.

Index Terms—Semantic scene understanding, AI-based methods, 3D scene graphs, spatial ontologies

I. INTRODUCTION

3D scene graph [1], [2], [3], [4], [5] is a novel map representation that describes the environment as a hierarchical graph grounded in the physical world, where nodes represent spatial concepts and edges represent relationships between concepts. The last few years have seen significant progress in the development of systems for real-time construction of 3D

J. Strader, N. Hughes, and L. Carlone are with the Laboratory for Information & Decision Systems (LIDS), Massachusetts Institute of Technology, Cambridge, MA, USA, {jstrader, na26933,lcarlone}@mit.edu

W.Chen is with the Berkeley Artificial Intelligence Research (BAIR), University of California, Berkely, CA, USA, verityw@berkeley.edu

A Speranzon is with Lockheed Martin, Advanced Technology Labs, Eagan, MN, USA, alberto.speranzon@lmco.com

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



1

Fig. 1: 3D scene graph of a indoor-outdoor environment at West Point, NY, constructed using our approach. The hierarchical structure of the graph is based on a spatial ontology generated from Large Language Models, and grounded into a 3D map using Logic Tensor Networks. A satellite image is shown of the area overlaid with the constructed 3D scene graph.

scene graphs in *indoor* environments. These approaches typically combine traditional SLAM methods and novel geometric deep learning techniques, such as Graph Neural Networks (GNN) and transformers [5]. The resulting 3D scene graphs include multiple layers describing semantic concepts at different levels of abstraction (e.g., from objects to places, rooms, and buildings), which have been proven useful in robotics. For instance, the benefits of using 3D scene graphs have been demonstrated for path planning [2], task planning [6], [7], and specific naviation tasks such as object search [8], by exploiting the semantics and hierarchical structure of the representation. Although one can envision similar benefits using 3D scene graphs in outdoor environments by exploiting the semantics and hierarchy of the representation [9], [10], no work to date has focused on constructing 3D scene graphs in arbitrary outdoor settings.

Generalizing 3D scene graphs to outdoor environments poses three main challenges: (i) Indoor vs. outdoor semantics: although there is a clear hierarchy of semantic concepts for indoor environments such as objects, rooms, floors, and buildings, the concepts needed to describe a broad variety of outdoors scenes is less immediate; thus, it is undesirable for a human to manually define such set of labels for each application. (ii) Lack of training datasets: while there are established training datasets one can use for indoor scene graph generation, there is virtually no dataset that allows creating semantically rich 3D scene graphs in outdoor scenes. For instance, prior work [9] leverages annotations by OpenStreetMap (OSM), but OSM only provides a few annotated classes such as roads, highways, and buildings and does not include smaller objects; this is problematic since existing GNN-based methods [5] group nodes (e.g., objects)

Manuscript received: December, 8, 2023; Revised February, 15, 2024; Accepted March, 14, 2024.

This paper was recommended for publication by Editor Javier Civera upon evaluation of the Associate Editor and Reviewers' comments. This work is partially founded by the ARL DCIST program, by the ONR program "Low Cost Autonomous Navigation & Semantic Mapping in the Littorals", and by Lockheed Martin Corporation. This research was also partially 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.

into higher-level concepts (*e.g.*, rooms) to build hierarchical representations. (iii) *Reliability of learning-based methods:* GNN-based approaches can provide incorrect predictions, especially when trained on relatively small datasets or tested outside the training domain; thus, an approach to constrain GNN predictions using common-sense knowledge to improve generalization and accuracy across different types of scenes is desirable.

Contribution. We tackle the three challenges above and propose a neuro-symbolic approach to extend 3D scene graph construction to arbitrary environments. To address the first challenge, we construct a spatial ontology (Section IV) that describes spatial concepts and relations, which represents common-sense spatial knowledge (e.g., a "bathroom" contains a "shower"). While it is well-known that building a comprehensive ontology is difficult and time-consuming, we show that Large Language Models (LLM) can be used to automatically build an ontology, and we show they can produce a reasonable ontology without fine-tuning and with minimal human effort. To address the second and third challenge, we resort to a novel learning approach, namely Logic Tensor Networks (LTNs) [11] (Section V). In particular, we design a model consisting of a GNN and a multi-layer perceptron (MLP) wrapped inside an LTN [11] that ensures the prediction of the semantic labels in the 3D scene graph satisfy the relations encoded by the spatial ontology.

We evaluate our approach in a variety of datasets, including indoor, rural, and coastal environments (Section VI). Our results show that the proposed approach enhances performance with limited training data; for instance, the LTN boosts performance from 12.3% to 25.1% for the tested indoor scenes and 29.0% to 37.2% (on average) for the tested outdoor scenes when training the proposed model using only 0.1% of the training data. We report runtime to demonstrate the proposed approach is suitable for real-time applications (*e.g.*, between 0.027s and 0.207s per scene graph in our experiments). Furthermore, the LTN-based approach predicts labels unseen at training time by leveraging common-sense relations between objects and the location containing them encoded by the spatial ontology.

II. RELATED WORK

3D Scene Graphs. The pioneering work of Armeni et al. [1] introduced the notion of 3D scene graphs as a hierarchical model of 3D environments. Following this, Rosinol et al. [2] extended this model to construct 3D scene graphs from sensor data while adding a richer hierarchy of layers. Other representative approaches also explore constructing 3D scene graphs from geometric data via GNNs [12] or probabilistic approaches [13]. More recently, several approaches explore the creation of 3D scene graphs in real-time in indoor environments [4], [3] or specific outdoor (*i.e.*, urban) environments [10]; in contrast, our work focuses on arbitrary indoor and outdoor environments. Most relevantly, Hughes et al. [5] discuss detecting room categories in 3D scene graphs using a fully-supervised GNN (which requires a significant amount of labeled training data), and Chen et al. [14] discuss using large language models for zero-shot room classification in 3D scene graphs. However, neither of these approaches applies to unstructured (i.e., outdoor) environments. Additionally, 3D scene graphs have started to find many applications in robotics and beyond. Recent works have successfully used 3D scene graphs to exploit the semantics and hierarchical structure for navigation [8], [15] as well as task and motion planning [6], [7]. Finally, 3D scene graphs also have nascent applications outside of robotics [16], [17].

Ontologies. To build and share a meaningful representation of a 3D environment, a robot needs a common "vocabulary" describing the concepts and relations observed by the robot. Such a vocabulary may be represented as an ontology. The precise definition of ontology varies between communities [18], and the term is often used interchangeably with knowledge graph and knowledge base. Typically, an ontology refers to the "schema" of data (i.e., abstract concepts and relations) whereas a knowledge graph is an instance of an ontology populated with data. We adopt the definition of knowledge base from [11] (more details in Section V). A significant effort has been made in the creation of common-sense ontologies and knowledge graphs [19], [20], and there has been a surge applying these to problems such as 2D scene graph generation [21], image classification [22], and visual question answering [23], to name a few. Most relevant to our work, Qiu et al. [24] and Zhang *et al.* [25] use ontologies to inform the generation of indoor 3D scene graphs, and Hsu et al. [26] discuss how to use LLMs to generate grounding predicates for object and their relationships. Constructing common-sense ontologies and knowledge graphs is an open problem; many authors have recently explored the role and use of LLMs to generate these representations. These works discuss either how to directly create common-sense representations by refining the LLM output [27] or evaluate the common-sense knowledge accuracy of LLMs [28]. The interested reader is referred to [29] for a discussion of challenges in this area.

III. PROBLEM STATEMENT

A 3D scene graph (e.g., Fig. 1) is a hierarchical graph [5] where nodes and edges represent spatial concepts and relations grounded in the physical world, and are grouped into layers corresponding to different levels of abstraction of the scene. In indoor environments, the layers from bottom to top consist of a metric-semantic 3D mesh, objects and agents, places (a topological representation of free space), and rooms (groupings of places describing rooms). While existing work on 3D scene graphs focus on grounding concepts for indoor environments, we extend the 3D scene graph model in [2], [3] to ground concepts in arbitrary environments (e.g., indoor and outdoor). Specifically, given the low-level layers (i.e., metric-semantic 3D mesh, objects and agents, and places) which have the same meaning in both indoor and outdoor environments and can be built as in [3], we construct the high-level layers (*i.e.*, regions layer generalizing the rooms layer) of a 3D scene graph based on a given set of spatial concepts (e.g., "kitchen" and "bathroom" for indoors or "road" and "field" for outdoors).

IV. LANGUAGE-ENABLED SPATIAL ONTOLOGIES

This section introduces the notion of *spatial ontology* and describes how to generate one using LLMs using two different approaches: text scoring and text completion.



Fig. 2: Text scoring approach for generating spatial ontology. The language model assigns a loss to the text string between low-level and high-level labels. The loss is rescaled using a softmax to assign weights to the edges, then the lowest weighted edges are pruned away.



Fig. 3: Text completion approach for generating spatial ontology. The language model is queried asking what low-level labels distinguish high-level labels. The response is parsed to generate the spatial ontology.

A. Spatial Ontology

A spatial ontology is a representation of spatial knowledge, which can be modeled as a graph where nodes are (abstract) spatial concepts and edges are spatial relations between the concepts. In this work, we consider a spatial ontology where each node corresponds to either low-level spatial concepts $\mathcal{L} =$ $\{\ell_1, \dots, \ell_n\}$ or high-level spatial concepts $\mathcal{H} = \{h_1, \dots, h_m\}$ where the n low-level and m high-level concepts correspond to semantics captured by the mesh layer (e.g., sand, rock) and locations captured by the places layer (e.g., beach), respectively. Furthermore, we only consider inclusion relations, which describe where low-level concepts may be located with respect to high-level concepts. For example, an edge between a low-level concept $\ell \in \mathcal{L}$ (e.g., "sink") and a high-level concept $h \in \mathcal{H}$ (e.g., "kitchen") means that ℓ is located within h (e.g., a "sink" is located in a "kitchen"). This implies that the spatial ontology is a bipartite graph with $|\mathcal{H}| \times |\mathcal{L}|$ biadjacency matrix Ω ¹ In contrast to a 3D scene graph, which is a grounded representation of the environment, a spatial ontology is independent from the specific environment observed by the robot. For example, the concept of "chair" in an ontology is not associated with a particular instance of a "chair", while in a 3D scene graph, the concept of "chair" is grounded in the scene the robot observes. In other words, an ontology is a graph representing all possible concepts and their relations (*i.e.*, it describes that chairs are typically in kitchens), while a 3D scene graph is a model of a specific environment (i.e., an instantiation of a spatial ontology with data).

B. Inferring Spatial Ontologies from LLMs

A spatial ontology may be defined manually by an expert, but such manual effort is time-consuming and does not scale well to large ontologies. While there is a history of using LLMs to evaluate relations between concepts, we focus on directly creating a spatial ontology, which can be leveraged for constructing 3D scene graphs. We describe two approaches (text scoring and text completion) to use LLMs to build spatial ontologies at scale, and while text completion is the preferred method used in our approach, we include text scoring for completeness since both generate reasonable ontologies based on our experiments (more in Section VI).

1) Text Scoring: In this approach, we initially assume edges exist between all pairs of low-level and high-level concepts. Weights are assigned to each edge in the ontology each representing if a low-level concept is located within a highlevel concept. The weights on the edges are assigned using an auto-regressive language model Λ , which is trained to assign probabilities to text strings. More precisely, Λ maps strings of tokens W to the log probability of the string written as $\Lambda(W) \approx \log p(W)$. The template g_{score} for generating the text string between the *i*th low-level concept $\ell_i \in \mathcal{L}$ and the *j*th high-level concept $h_j \in \mathcal{H}$ is given by

$$g_{score}(i,j) = \ell_i$$
 is often found in h_j .

The language model assigns a loss to each edge based on the associated text string, which corresponds to the log probability the text string will occur given the model parameters, then the weights for each edge are rescaled using a softmax function:

$$w_{ij} = \frac{\exp\{\Lambda(g_{score}(i,j))/K\}}{\sum_{k=1}^{|H|} \exp\{\Lambda(g_{score}(i,k))/K\}},$$
(1)

where $K \in (0, \infty]$ is a temperature parameter. The weights for the edges incident to each low-level concept are accumulated in descending order by weight until exceeding a threshold γ , then the edges not reached are pruned from the spatial ontology; formally, the *r* retained edges for the *i*th low-level concept satisfy $\sum_{j=1}^{r-1} w_{ij} \leq \gamma$ and $\sum_{j=1}^{r} w_{ij} > \gamma$ where γ is a user-defined threshold. Figure 2 provides an overview of the overall approach.

2) Text Completion: While the approach described above generates reasonable ontologies (more in Section VI), the result is an ontology with edges for the most likely relations (e.g., a "wall" is more likely to be located in a "bathroom" than a "shower"); instead, the ontology should include edges that are useful for distinguishing between different high-level concepts based on low-level concepts (e.g., a "shower" is more likely to distinguish a "bathroom" from a "bedroom" than a "wall"). Here, we construct the spatial ontology by querying the language model for each high-level concept to generate the k edges connecting low-level concepts to each high-level concept. The template g_{comp} for generating the text string for the k edges associated with the *i*th high-level concept $h_i \in \mathcal{H}$ is given by

$$g_{comp}(i,k) =$$
 "Which $\{k\}$ items from $\{\mathcal{L}\}$ are most
likely to distinguish $\{h_i\}$ from $\{\mathcal{H}/h_i\}$. Answer
with a python list using exact strings in $\{\mathcal{L}\}$."

¹The adjacency matrix for a bipartite graph has the form $\begin{bmatrix} 0 & \Omega \\ \Omega^{\mathsf{T}} & 0 \end{bmatrix}$ where Ω is the biadjacency matrix. In our problem, the structure of the adjacency matrix formalizes the fact that the edges only relate low-level concepts (*e.g.*, objects) with high-level concepts (*e.g.*, locations).

After querying the language model, the response requires post-processing since the language model may not format the response as a python list or hallucinate concepts (*i.e.*, concepts not from \mathcal{L}). If the response is improperly formatted, the formatting is corrected in post-processing, and if hallucinated concepts are in the response, the text string used for the query is appended with "Do not respond with concepts in $\{\mathcal{L}'\}$ " where \mathcal{L}' is the set of concepts returned during this procedure not in the given set of concepts \mathcal{L} . The response generated by the language model is stochastic; thus, each query is repeated N times, and the top k edges generated are used in the spatial ontology; see Fig. 3.

V. LEVERAGING SPATIAL ONTOLOGIES FOR 3D SCENE GRAPH GENERATION

This section briefly reviews LTNs and shows how to generate 3D scene graphs using LTNs. Specifically, we show how to build a knowledge base from a spatial ontology, which is leveraged for constructing 3D scene graphs.

A. Logic Tensor Networks

A Logic Tensor Network (LTN) [11] is a neuro-symbolic framework that combines explicit symbolic knowledge expressed as a set of logical formulas with implicit subsymbolic knowledge (e.g., encoded within the parameters of a neural network). LTNs support tasks such as learning (making generalization from data), reasoning (verifying a consequence from facts), and query-answering (evaluating the truth value of a logical expression). To express the relationship between symbolic and subsymbolic knowledge, LTNs are built on Real Logic [30], which is a differentiable first-order language with truth values in [0, 1]. This allows converting Real Logic formulas (*e.g.*, a differentiable loss) into computational graphs that enable gradient-based optimization.

1) Syntax of Real Logic: Formally, Real Logic is defined on a first-order language \mathcal{F} with a signature² consisting of sets of symbols for constants, variables, functions, and predicates. A formula (or *axiom*) in \mathcal{F} is an expression constructed using symbols in the signature together with logical connectives, $\neg, \land, \lor, \rightarrow$, and quantifiers, \forall, \exists , following the same syntax³ as First-Order Logic (FOL). For example, we can construct a formula stating that "every *kitchen* contains a *sink*" as $\forall y(\texttt{Kitchen}(y) \rightarrow \exists x(\texttt{Sink}(x) \land \texttt{PartOf}(x, y))$, where x, y are variables and $\texttt{Sink}(\cdot)$, $\texttt{Kitchen}(\cdot)$, $\texttt{PartOf}(\cdot, \cdot)$ are predicates. In real-world scenarios, a formula may be partially true (*e.g.*, a *kitchen* may exist that does not contain a *sink*); thus, Real Logic adopts continuous truth values in [0, 1].

2) Semantics of Real Logic: The semantics of Real Logic are the rules for interpreting symbols with respect to the real world, thus explicitly connecting symbolic and subsymbolic knowledge. Intuitively, interpretations⁴ \mathcal{I} are the numeric values associated with each symbol, as well as how to numerically implement the operations applied to such values. In general,

constants and variables are interpreted as tensors, whereas operations include *functions*, which map tensors to tensors, and *predicates*, which map tensors to real numbers in [0, 1]. For example, a variable x may be interpreted as a feature vector associated to an object observed by the robot (*e.g., sink*), and the predicate Sink(·) may be interpreted as a binary classifier that takes a feature vector as input and returns if the input is a *sink* or not (*e.g.,* returning a truth value between 0 and 1). In Real Logic, a formula ϕ under an interpretation \mathcal{I} is satisfied if $\mathcal{I}(\phi) = 1$, but in practice, a formula ϕ is typically only partially satisfied (*i.e.,* $\mathcal{I}(\phi) \in [0, 1]$). Since predicates are partially true or false, there are various ways of defining connectives and quantifiers (*e.g.,* see [31]).⁵

3) Learning with LTNs: If we are given a set of formulas in Real Logic, LTNs allow learning the interpretation of symbols in the signature. For example, if constants are interpreted as embeddings, and functions and predicates are interpreted as neural networks, the interpretation \mathcal{I} depends on learnable parameters θ with a parametric interpretation $\mathcal{I}(\cdot|\theta)$. The symbols in the signature, logical formulas, and interpretations together are referred to as a *Real Logic knowledge base*. The optimization problem underlying LTNs for learning interpretations is a *satisfiability problem*, which yields the values of the parameters θ^* that maximize the truth values of the conjunction of a set of logical formulas \mathcal{K} in a Real Logic knowledge base. We can write this as:

$$\theta^* = \operatorname*{arg\,min}_{\theta \in \Theta} \left(1 - \operatorname{SatAgg} \mathcal{I}(\phi|\theta) \right) \tag{2}$$

where SatAgg : $[0,1]^* \rightarrow [0,1]$ is an aggregating operator over all the formulas $\phi \in \mathcal{K}$, typically implemented numerically as a generalized mean [11]. Since Real Logic interprets expressions as real-valued tensors, LTNs associate gradients to the expressions in the knowledge base and optimizes (2) via gradient-based optimization.

B. Formalizing 3D Scene Graph Generation in LTNs

This section presents our approach to create hierarchical 3D scene graphs with a spatial ontology using LTNs. Specifically, given the low-level layers of a 3D scene graph (*i.e.*, metric-semantic 3D mesh, objects and agents, and places), we construct the high-level layers based on a given spatial ontology and labeled examples (if available). We formulate this as a node classification problem on the places layer using a GNN for generating node embeddings and an MLP for predicting node labels. We leverage LTNs during training to constrain the model to make predictions consistent with the spatial ontology. The proposed architecture for learning and inference is given in Fig. 4. We train the model end-to-end to learn the parameters of both the GNN and MLP based on the satisfaction of a knowledge base, as described below.

1) Constants and Variables: While constants and variables are used in the same way in functions and predicates, the difference is that variables can appear in quantifiers (*i.e.*, a variable can have different instances with different values while constants are always a fixed value).

²The signature of a language is the set of symbols with meanings that vary with the interpretation (*i.e.*, non-logical symbols) and exclude connectives and quantifiers (*i.e.*, logical symbols).

³The syntax is the set of rules for constructing well-formed expressions.

⁴In [11], the authors use the terminology *grounding* instead of *interpretation*, the latter is standard in mathematical logic and preferred here.

⁵For example, in Łukasiewicz logic [31], $\mathcal{I}(\phi_1 \land \phi_2) = \max(0, \mathcal{I}(\phi_1) + \mathcal{I}(\phi_2) - 1)$ and $\mathcal{I}(\phi_1 \lor \phi_2) = \max(1, \mathcal{I}(\phi_1) + \mathcal{I}(\phi_2))$ where ϕ_1 and ϕ_2 are formulas.



(a) Learning Phase. In the learning phase, the place features are processed by the GNN to compute node embeddings, and the MLP processes the embeddings and returns the softmax of the logits. The MLP is wrapped in an LTN to compute the satisfaction of a Real Logic knowledge base using the data-driven axioms based on ground truth labels and ontology-driven axioms based on the aggregated semantics from the mesh layer.



(b) Inference Phase. In the inference phase, the LTN is dropped, and the model predicts the node labels based on the place features using the trained MLP and GNN from the learning phase.

Fig. 4: Proposed architecture for learning and inference.

Constants. The graph representing the places layer of a 3D scene graph is $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where \mathcal{V} is the set of place nodes and \mathcal{E} is the set of edges between place nodes. The label space is $\{y_1, \dots, y_m\}$, which corresponds to the *m* highlevel concepts \mathcal{H} in the spatial ontology; specifically, y_i is the one-hot vector for $h_i \in \mathcal{H}$. The spatial ontology is a constant Ω interpreted as the $|\mathcal{H}| \times |\mathcal{L}|$ biadjacency matrix (*e.g.*, built as described in Section IV), which defines the valid choices of a high-level concepts for a place node given the low-level concepts observed by the robot.

Variables. The GNN associates an embedding x_v to each place node $v \in \mathcal{V}$. This embedding is passed to an MLP to classify the high-level concept (*i.e.*, region) the node belongs to. At the same time, we attach a second (pre-computed) feature vector q_v to each place node $v \in \mathcal{V}$, describing the low-level concepts observed by the robot (*e.g.*, nearby objects). More in detail, each place node is connected to nearby nodes in the metric-semantic 3D mesh, called basis points (see [5]); the semantic labels of the basis points are one-hot vectors representing the low-level concepts \mathcal{L} in the ontology. We build q_v as the histogram of semantic labels of the basis points are captured by the basis points; thus, the object and agent layers are not required.

2) Predicates: The predicates take constants and variables as input and return a real number in [0, 1]. Here, we propose predicates, which are either data-driven (referred to as *equivalence* predicates) or ontology-driven (referred to as *inclusion* predicates). This allows learning the interpretations of predicates parameterized by the parameters θ of our model using a combination of labeled data and unlabeled data. The equivalence predicates rely only on the constants y_i (*i.e.*, the labeled data) while inclusion predicates rely only on the constant Ω (*i.e.*, only use the spatial ontology).

Equivalence. We propose $IsClassOf(\cdot, \cdot|\theta)$ as an equivalence predicate, which quantifies if a high-level concept predicted by the $MLP_{\theta}(\cdot)$ for a place node matches the ground truth label (if available at training time):

$$\mathcal{I}(\texttt{IsClassOf}|\theta): x_v^{(i)}, y_i \mapsto y_i^T \cdot \texttt{MLP}_{\theta}(x_v^{(i)})$$
 (3)

where $x_v^{(i)}$ denotes the embedding for a place node $v \in \mathcal{V}$ with label y_i . In words, the predicate measures the agreement between the predicted label $MLP_{\theta}(x_v^{(i)})$ and the ground truth label y_i , and represents the statement: "A place node with label y_i is classified as y_i ". Intuitively, this predicate plays a similar role as cross-entropy loss in supervised learning.

Inclusion. At training time, we encourage the model (*i.e.*, GNN+MLP) to predict labels for place nodes that match the spatial ontology based on the connected basis points. More formally, we build a formula stating "If the spatial ontology has connections valid with the basis points connected to a place node, the place node is classified with a label similar to the ones encoded in the spatial ontology". The first part of the formula is to circumvent missing data in the spatial ontology: for instance, the ontology may have disconnected low-level concepts, and we do not use the ontology in those cases. The second part encourages the prediction to agree with at least one of the edges in the spatial ontology. In order to formalize this formula, we need two predicates: $IsValid(\cdot, \cdot)$ and $IsSimilar(\cdot, \cdot, \cdot | \theta)$, which we describe below.

In order to assert the validity of the low-level concepts q_v attached to a node $v \in \mathcal{V}$ (*i.e.*, that low-level concepts connected to a place node are connected to high-level concepts in the ontology), we define a predicate IsValid(\cdot, \cdot):

$$\mathcal{I}(\text{IsValid}): \mathbf{\Omega}, q_v \mapsto \text{sum}(\mathbf{\Omega} \cdot \hat{q_v})$$
 (4)

where $\hat{\cdot}$ is an operator normalizing columns with respect to the ℓ_1 -norm. Intuitively, $\texttt{IsValid}(\Omega, q_v) = 0$ if none of the low-level concepts in the histogram q_v are connected to a high-level concept in the spatial ontology Ω .

Next, we define the predicate $\texttt{IsSimilar}(\cdot, \cdot, \cdot | \theta)$ to quantify the degree that high-level concepts predicted by the MLP are similar to the expected high-level concepts in the spatial ontology based on the low-level features q_v :

$$\mathcal{I}(\text{IsSimilar}|\theta): x_v, \mathbf{\Omega}, q_v \mapsto \text{MLP}_{\theta}(x_v)^T \cdot \hat{\mathbf{\Omega}} \cdot \hat{q_v}.$$
 (5)

Intuitively, $\hat{\Omega} \cdot \hat{q_v}$ represents the valid high-level concepts associated to low-level concepts in q_v (according to the ontology Ω), and the predicate computes the similarity between these high-level concepts and the ones predicted by the model, which enforces alignment between the model predictions $\mathbb{MLP}_{\theta}(x_v)$ and ontology predictions $\hat{\Omega} \cdot \hat{q_v}$.

3) Axioms: The axioms are the logical formulas constructed from constants, variables, and predicates representing statements. We distinguish the axioms by whether they use the equivalence or the inclusion predicates. The equivalence axioms ϕ_{Equiv} and inclusion axioms ϕ_{Incl} are:

$$\begin{split} \phi_{\text{Equiv}} &: \forall x_v^{(i)} \text{ IsClassOf}(x_v^{(i)}, y_i) \\ \phi_{\text{Incl}} &: \forall x_v, q_v \text{ IsValid}(\mathbf{\Omega}, q_v) \rightarrow \text{IsSimilar}(x_v, \mathbf{\Omega}, q_v) \end{split}$$

where diagonal quantification is used for both ϕ_{Equiv} and ϕ_{Incl} as described in [11]. Here, $x_v^{(i)}$ corresponds to place nodes

with label y_i , and each place node $v \in \mathcal{V}$ is assigned features x_v and q_v ; thus, we only want to construct statements between x_v and q_v for each $v \in \mathcal{V}$. At training time, the equivalence axioms promote the model to make predictions that agree with the annotated labels (if available), while the inclusion axioms encourage the model to make prediction that agree with the spatial ontology based on the semantics of the connected basis points. In general, the equivalence axioms may contradict the inclusive axioms if the annotated labels disagree with the spatial ontology; in such cases, both types of axioms will influence the predictions.

VI. EXPERIMENTS

In this section, we describe experiments conducted to evaluate our approach across a variety of environments where annotated data is limited or missing completely.

Experimental Setup. We utilize three datasets: Matterport3D (MP3D) [32], West Point, and Castle Island. For all datasets, we use Hydra [5] to generate 3D scene graphs. The MP3D [32] dataset is an RGB-D dataset consisting of 90 indoor scenes. We use the Habitat simulator [33] to render the visual and human-annotated semantic data for the scenes. The provided label space has 40 semantic labels (i.e., mpcat40 [32]), of which we discard the labels gym_equipment and beam. For our experiments, we use the official train, test, and validation splits of the scenes [32]. We collect two outdoor datasets. The first, West Point, consists of two robot trajectories where a RealSense D455 was mounted on a Clearpath Jackal rover. These trajectories explore areas mostly consisting of roads, fields, and buildings. The second, Castle Island, was collected from Pleasure Bay beach in Boston via a hand-held sensor payload, which consists of two trajectories exploring areas around the beach and nearby parking lot. The trajectories from both datasets are roughly a half kilometer to a kilometer in length. For both datasets, we use OneFormer [34] to provide semantics (*i.e.*, low-level concepts) with the ADE20k [35] label space. We use LOCUS [36] with a Velodyne mounted on the Jackal to provide trajectory estimates for the West Point dataset, and Kimera-VIO [37] for trajectory estimates of Castle Island.

Training. We train the model using different axiom configurations that include (i) only the data-driven axioms ϕ_{Equiv} (label: "SAT (Equiv.)"), (ii) only the ontology-driven axioms ϕ_{Incl} (label: "SAT (Incl.)"), and (iii) both the data-driven and ontology-driven axioms (label: "SAT (Both)"). As a baseline, we train a GNN using cross-entropy loss (label: "Cross-Entropy"), which is expected to be similar to the model trained with the purely data-driven axioms. The node features for the places are the concatenation of the position and the word2vec [38] encoding of the aggregated semantic labels of the child nodes in the mesh layer. For the GNN, we use a graph attention network (GAT) [39] and tune the hyper-parameters using grid search. We use text completion (Section IV-B2) for generating the ontology, since it generates ontologies that agree (with respect to precision) with human judgement more closely than text scoring (Section IV-B1); see Table I. For the connectives and quantifiers, we use a product configuration except for the implication operator where we use the Goguen operator [31]. We use stochastic gradient descent with the Adam optimizer, and we run the optimization

Ontology	Accuracy	Precision	Recall
Text Completion $(k=1)$	0.50	1.0	0.14
Text Completion $(k=2)$	0.52	0.83	0.22
Text Completion $(k=3)$	0.65	1.0	0.40
Text Completion $(k=4)$	0.68	0.85	0.54
Text Completion $(k=5)$	0.71	0.86	0.59
Text Scoring	0.81	0.77	0.95

TABLE I: Evaluation of ontologies produced by text scoring and text completion against an ontology created using human judgement for the MP3D label space with a temperature K=10 for text scoring. Best in bold.

Ontology	MP3D	West Point	Castle Island
k = 1 $k = 2$ $k = 3$ $k = 4$ $k = 5$ Custom	$14.0 \pm 2.0 \\ 16.7 \pm 3.6 \\ 23.1 \pm 0.2 \\ 2.7 \pm 0.0 \\ 12.3 \pm 0.0 \\ 27.7 \pm 0.4 \\ 12.3 \pm 0.0 \\ 27.7 \pm 0.4 \\ 12.3 \pm 0.0 \\ 12.3$	37.8 ± 0.3 11.9 ± 14.7 42.7 ± 0.2 39.9 ± 4.7 38.8 ± 12.0 27.7 ± 1.2	$16.1 \pm 3.6 \\ 28.7 \pm 2.2 \\ 24.5 \pm 0.6 \\ 28.7 \pm 8.5 \\ 19.7 \pm 7.1 \\ 22.7 \pm 1.8 \\ 1.8 \\ 22.7 \pm 1.8 \\ 1.8 \\ 22.7 \pm 1.8 \\ 22.7 \pm 1.8 \\ 22.7 \pm 1.8 \\ 23.7 \\ 23.7 \\ 23.7 \\ 23.7 \\ 23.7 \\ 24.7 \\ 24.7 \\ 25.7 $

TABLE II: Comparison of mean accuracy achieved by the "SAT (Incl.)" loss configuration across the three datasets between the LLM generated ontologies (k = 1, ..., 5) versus a hand-generated ontology for the same low-level label space and high-level concepts. Best in bold.

until the loss has converged (*e.g.*, changes less than 1e-06 for 10 epochs) or 1000 epochs is reached. The model with the best validation accuracy is retained for testing. Each model is trained 10 times using a random initialization of the model parameters, and the statistics are reported in the results. We implement the architectures using PyTorch Geometric [40] and LTNTorch [11].

Results. We conduct an ablation study comparing the crossentropy baseline and the different axioms configurations for the LTN. The accuracy (i.e., percentage of correctly classified places) for all datasets is reported in Fig. 5, and the average runtime for testing per scene graph on a CPU (Intel i7-11800 at 2.3GHz) is 0.027s for MP3D, 0.067s for West Point, and 0.207s for Castle Island. For each dataset, we evaluate the performance for 100%, 10%, 1%, 0.4%, 0.2%, and 0.1% of the training data. We use the train-test-validation splits for MP3D as described in [32], and we mask these percentages of places from the training scene to simulate a scenario where annotated data is limited. For West Point and Castle Island, we train in a semi-supervised manner where we combine both graphs (70%training, 15% validation and testing), as the datasets are too small to generalize from a single training scene to a single testing scene.

The data-driven approaches (our approach, the GNN baseline, and our data-driven axiom variant) show similar performance when trained on the full MP3D test dataset; see Fig. 5a. However, as less data is used during training, the purely data-driven approaches tend towards performance comparable with randomly guessing labels, while the proposed approach does not fall below the purely ontology-driven approach. The same trend is visible in Fig. 5b for the West Point dataset. The data-driven approaches in Fig. 5c for the Castle Island dataset do not have a significant gap in performance compared to the ontology-driven approaches. This is due to a lack in distinguishing low-level concepts for the high-level concepts in the scene. We choose k by generating ontologies with multiple values $(k = \{1, 2, 3, 4, 5\})$, and the k with the highest test accuracy is used in the ablations; alternatively, k may be chosen by evaluating against a smaller ontology created based on human judgement. We compare these generated ontologies



Fig. 5: (a) Results of the predicate ablation for the MP3D dataset. (b) Results of the predicate ablation for the West Point dataset. (c) Results of the predicate ablation for the Castle Island dataset. For all experiments, we run 10 trials of each loss configuration for each percentage of the training data. The mean of the 10 trials is shown as a dashed line, and the shaded area shows the standard deviation.



(d) West Point

(e) Castle Island

(f) Outdoor Color Legend

Fig. 6: (a) The "q9vSo1VnCiC" scene from MP3D and (b) classified place nodes for the "q9vSo1VnCiC" scene along with (c) associated color legend for node labels. (d) Classified place nodes for one of the West Point trajectories. (e) Classified place nodes for on of the Castle Island trajectories. (f) Associated color legend for both West Point and Castle Island. The ground truth regions are drawn as shaded polygons, while the place nodes are all drawn as circles.

	% of Data Kept for MP3D w/ Partial Labels			
Loss	100%	0.4%	0.2%	0.1%
CrossEntropy SAT (Equiv.) SAT (Incl.) SAT (Both)	$\begin{array}{c} 24.1 \pm 2.4 \\ 24.7 \pm 1.4 \\ 22.7 \pm 0.6 \\ \textbf{32.5} \pm \textbf{1.2} \end{array}$	$\begin{array}{c} 16.7\pm 5.6\\ 15.8\pm 4.0\\ \textbf{22.7}\pm \textbf{0.6}\\ 20.6\pm 1.4 \end{array}$	$\begin{array}{c} 12.7 \pm 4.7 \\ 12.8 \pm 5.5 \\ 22.7 \pm 0.6 \\ \textbf{22.8} \pm \textbf{1.9} \end{array}$	$\begin{array}{c} 12.5\pm5.1\\ 11.7\pm5.1\\ 22.7\pm0.6\\ \textbf{23.3}\pm\textbf{1.4} \end{array}$

TABLE III: Accuracy for different holdout percentages of training data with labels masked for *bathroom*, *kitchen*, *hallway*, *bedroom*, *living room*, and *family room*. No data held out during testing. Best in bold.

with a hand-generated ontology and report the results in Table II.

Finally, a 3D scene graph obtained by our method for each dataset is shown in Fig. 6. In general, the high-level concepts generated by our approach are plausible for the associated scenes. For certain scenes, our approach struggles to distinguish similar high-level concepts. This is apparent for the beach (*e.g.*, ground truth labels "backshore" are classified as "beach" instead). Additionally, low-level concepts that contradict the spatial ontology cause our model to misclassify certain regions. This is clear in the MP3D scene shown in Fig. 6b; the lounge in the top-right of the building is misclassified as a bedroom due to the presence of a day-bed in the room, which is labeled as "bed" in the ground truth semantics of the MP3D dataset. Lastly, we show in the MP3D dataset that the ontology-driven models allow predicting labels unseen at training time. This scenario is simulated by masking all place nodes for the following set of labels: *bathroom, kitchen, hallway, bedroom, living room,* and *family room.* In this scenario, the ontology-driven models significantly outperform the purely data-driven models as shown in Table III since labels for the removed classes are missing during training.

VII. CONCLUSION

We propose the first approach to build hierarchical 3D scene graphs of arbitrary (indoor and outdoor) environments. Towards this goal, we leverage Large Language Models to build a *spatial ontology* that defines the set of concepts and relations relevant for robot operation. We show how to use such a spatial ontology for 3D scene graph construction using *Logic Tensor Networks* (LTN), which allow adding logical rules, or *axioms (e.g.,* "a beach contains sand"), during training and prediction. We test our approach in a variety of datasets, including indoor, rural, and coastal environments, and show that it leads to a significant increase in the quality of the 3D scene graph generation with sparsely annotated data, and it produces reasonable results even in a zero-shot setting. This

work opens several avenues for future research, including constructing 3D scene graphs with multiple high-level layers, incorporating different types of relations in the ontology, and more powerful feature representations.

REFERENCES

- 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)*, 2019, pp. 5664–5673. 1, 2
- [2] 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*, vol. 40, no. 12–14, pp. 1510–1546, 2021. 1, 2
- [3] N. Hughes, Y. Chang, and L. Carlone, "Hydra: a real-time spatial perception engine for 3D scene graph construction and optimization," in *Robotics: Science and Systems (RSS)*, 2022. 1, 2
- [4] 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, pp. 7515–7525. 1, 2
- [5] N. Hughes, Y. Chang, S. Hu, R. Talak, R. Abdulhai, J. Strader, and L. Carlone, "Foundations of spatial perception for robotics: Hierarchical representations and real-time systems," *Intl. J. of Robotics Research*, 2024. 1, 2, 5, 6
- [6] C. Agia, K. M. Jatavallabhula, M. Khodeir, O. Miksik, V. Vineet, M. Mukadam, L. Paull, and F. Shkurti, "Taskography: Evaluating robot task planning over large 3D scene graphs," in *Conference on Robot Learning (CoRL)*, 2022, pp. 46–58. 1, 2
- [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 task planning," in *Conference on Robot Learning* (*CoRL*), 2023, pp. 23–72. 1, 2
- [8] Z. Ravichandran, L. Peng, N. Hughes, J. 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, pp. 9272– 9279. 1, 2
- [9] M. Berg, G. Konidaris, and S. Tellex, "Using Language to Generate State Abstractions for Long-Range Planning in Outdoor Environments," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2022, pp. 1888– 1895. 1
- [10] E. Greve, M. Büchner, N. Vödisch, W. Burgard, and A. Valada, "Collaborative dynamic 3d scene graphs for automated driving," *arXiv preprint:* 2309.06635, 2023. 1, 2
- [11] S. Badreddine, A. d'Avila Garcez, L. Serafini, and M. Spranger, "Logic tensor networks," *Artificial Intelligence*, vol. 303, p. 103649, 2022. 2, 4, 5, 6, 9
- [12] J. Wald, H. Dhamo, N. Navab, and F. Tombari, "Learning 3D semantic scene graphs from 3D indoor reconstructions," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 3961– 3970. 2
- [13] N. Gothoskar, M. Cusumano-Towner, B. Zinberg, M. Ghavamizadeh, F. Pollok, A. Garrett, J. Tenenbaum, D. Gutfreund, and V. Mansinghka, "3DP3: 3D scene perception via probabilistic programming," in *arXiv* preprint: 2111.00312, 2021. 2
- [14] W. Chen, S. Hu, R. Talak, and L. Carlone, "Leveraging large (visual) language models for robot 3D scene understanding," arXiv preprint: 2209.05629, 2022. 2
- [15] Z. Seymour, N. Chowdhury Mithun, H.-P. Chiu, S. Samarasekera, and R. Kumar, "Graphmapper: Efficient visual navigation by scene graph generation," in *Intl. Conf. on Pattern Recognition (ICPR)*, 2022, pp. 4146–4153. 2
- [16] H. Dhamo, F. Manhardt, N. Navab, and F. Tombari, "Graph-to-3d: Endto-end generation and manipulation of 3d scenes using scene graphs," in *Intl. Conf. on Computer Vision (ICCV)*, 2021, pp. 16352–16361. 2
- [17] T. Tahara, T. Seno, G. Narita, and T. Ishikawa, "Retargetable AR: Context-aware augmented reality in indoor scenes based on 3d scene graph," in *IEEE and ACM Intl. Sym. on Mixed and Augmented Reality* (ISMAR), 2020, pp. 249–255. 2
- [18] T. R. Gruber, "Toward principles for the design of ontologies used for knowledge sharing?" *Intl. J. of Human-Computer Studies*, vol. 43, no. 5-6, pp. 907–928, 1995. 2
- [19] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. Ives, "DBpedia: A nucleus for a web of open data," in *Semantic Web*, 2007, pp. 722–735. 2

- [20] R. Speer, J. Chin, and C. Havasi, "ConceptNet 5.5: An open multilingual graph of general knowledge," in AAAI Conference on Artificial Intelligence, 2017, pp. 4444–4451. 2
- [21] A. Zareian, S. Karaman, and S.-F. Chang, "Bridging knowledge graphs to generate scene graphs," in *European Conf. on Computer Vision* (ECCV), 2020, pp. 606–623. 2
- [22] Y. Movshovitz-Attias, Q. Yu, M. C. Stumpe, V. Shet, S. Arnoud, and L. Yatziv, "Ontological supervision for fine grained classification of street view storefronts," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 1693–1702. 2
- [23] W. Zheng, L. Yin, X. Chen, Z. Ma, S. Liu, and B. Yang, "Knowledge base graph embedding module design for visual question answering model," *Pattern Recognition*, vol. 120, p. 108153, 2021. 2
- [24] Y. Qiu and H. I. Christensen, "3D scene graph prediction on point clouds using knowledge graphs," arXiv preprint: 2308.06719, 2023. 2
- [25] S. Zhang, S. Li, A. Hao, and H. Qin, "Knowledge-inspired 3D scene graph prediction in point cloud," in Advances in Neural Information Processing Systems (NIPS), vol. 34, 2021, pp. 18620–18632.
- [26] J. Hsu, J. Mao, and J. Wu, "NS3D: Neuro-symbolic grounding of 3D objects and relations," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 2614– 2623. 2
- [27] C. Chen, K. Lin, and D. Klein, "Constructing taxonomies from pretrained language models," in *Conference of the North American Chapter* of the Association for Computational Linguistics: Human Language Technologies, 2021, pp. 4687–4700. 2
- [28] K. Sun, Y. E. Xu, H. Zha, Y. Liu, and X. L. Dong, "Head-to-Tail: How knowledgeable are large language models (LLM)? A.K.A. will LLMs replace knowledge graphs?" arXiv preprint: 2308.10168, 2023. 2
- [29] J. Z. Pan, S. Razniewski, J.-C. Kalo, S. Singhania, J. Chen, S. Dietze, H. Jabeen, J. Omeliyanenko, W. Zhang, M. Lissandrini *et al.*, "Large language models and knowledge graphs: Opportunities and challenges," *arXiv preprint: 2308.06374*, 2023. 2
- [30] R. Fagin, R. Riegel, and A. Gray, "Foundations of reasoning with uncertainty via real-valued logics," arXiv preprint: 2008.02429, 2020. 4
- [31] E. van Krieken, E. Acar, and F. van Harmelen, "Analyzing differentiable fuzzy logic operators," *Artificial Intelligence*, vol. 302, p. 103602, 2022.
 4, 6
- [32] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3D: Learning from RGB-D data in indoor environments," *Intl. Conf. on 3D Vision*, pp. 667–676, 2017. 6
- [33] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, D. Parikh, and D. Batra, "Habitat: A platform for embodied AI research," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 9339– 9347. 6
- [34] J. Jain, J. Li, M. Chiu, A. Hassani, N. Orlov, and H. Shi, "OneFormer: One Transformer to Rule Universal Image Segmentation," in *IEEE Conf.* on Computer Vision and Pattern Recognition (CVPR), 2023, pp. 2989– 2998. 6
- [35] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, "Scene parsing through ADE20K dataset," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5122–5130. 6
- [36] A. Reinke, M. Palieri, B. Morrell, Y. Chang, K. Ebadi, L. Carlone, and A. Agha-mohammadi, "LOCUS 2.0: Robust and computationally efficient lidar odometry for real-time underground 3D mapping," *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 4, pp. 9043–9050, 2022. 6
- [37] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an opensource library for real-time metric-semantic localization and mapping," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2020, pp. 1689– 1696. 6
- [38] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprints: 1301.3781*, 2013. 6
- [39] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lió, and Y. Bengio, "Graph attention networks," in *arXiv preprint: 1710.10903*, 2018. 6
- [40] M. Fey and J. E. Lenssen, "Fast graph representation learning with Py-Torch Geometric," in *Intl. Conf. on Learning Representations Workshop* on Representation Learning on Graphs and Manifolds, 2019. 6, 9
- [41] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lió, and Y. Bengio, "Graph attention networks," in *Intl. Conf. on Learning Representations (ICLR)*, 2018. 9

A. Details for Training

Here, we provide additional details on training the proposed model.

Implementation. We implement the learning architectures using PyTorch Geometric [40] version 2.3.1 built with CUDA verion 11.7 and LTNTorch [11] version 1.0.1, and we use GPT-4 for generating the spatial ontology. For the GNN, we use GAT [41] for the message passing layers, and we use ReLU activation functions and dropout between message passing layers where dropout is enabled for training but disabled for testing. After the messaging passing layers in the GNN, the hidden states are passed to the MLP with a linear layer followed by ReLU activation functions followed by another linear layer, then a softmax layer is applied for predicting the node labels. The input and hidden dimensions for the MLP are the same as the hidden dimensions for the message passing layers in the GNN.

Hyper-Parameter Tuning. We provide details on the method applied to tune the hyper-parameters of our model. Specifically, we search over the following sets:

- Message passing iterations: [1,2,3]
- Message passing hidden dimensions: [16, 32, 64]
- Number of attention heads: [1, 2, 4]
- Learning rate: [0.01, 0.001, 0.0001]
- Dropout probability: [0.1, 0.25, 0.5]
- Weight decay: [0.005, 0.0005, 0.00005]

In order to provide fair comparison between the different axioms configurations, a grid search was performed over these hyper-parameters for the Cross-Entropy configuration as well as each of the LTN configurations. For the Cross-Entropy configuration and the LTN configurations, the number of convolutional layers, the learning rate, and the weight decay made the most impact on the accuracy of the model (with weight decay making less impact than convolutional layers and learning rate). Thus, we tune the hyper-parameters for the message passing iterations while keeping the remaining parameters constant at the median values. The learning rate and weight decay are next tuned by searching over all combinations of these values, and the grid search is performed over the dropout rate and the number of attention heads. The grid search described above was performed with 1000 epochs using the MP3D dataset, and we repeated this process for the learning rate and weight decay for both the West Point and Castle Island datasets, and optimal values of hyper-parameters remained the same as for the MP3D dataset. The optimal hyper-parameters based on the grid search are the same for each of the Cross-Entropy and different axioms configuration; see Table IV for the selected values of the hyper-parameters.

Lavers	3
Hidden Dimension	32
Learning Rate	0.001
Dropout	0.25
L_2 Regularization	5e-05
Attention Heads	4

TABLE IV: Hyper-parameters for training the proposed model where the same hyper-parameters are used for the Cross-Entropy configuration and the LTN configurations.

Additionally, we must select the fuzzy logical operators used for the axioms in the LTN along with the hyper-parameters associated with these operators. For the implication operator, we use the Goguen operator instead of the more commonly used Reichenbach operator, as the Goguen operator provided better accuracy. Additionally, we use the *p*-mean error for the aggregating operator SatAgg(\cdot) as described in [11], and we

use p=2 for the data-driven axioms (*i.e.*, equivalence axioms) and p=4 for the ontology-driven axioms (*i.e.*, inclusive axioms). The model was trained with p = [2, 4] for each axiom configuration, and the optimal value was selected; however, the difference in accuracy was insignificant.

Additionally, for the experiments where only a percentage of annotations for the data is kept during training, the annotations kept are selected randomly; furthermore, we seed the random number generator identically between the experiments for the Cross-Entropy configuration and the different LTN configurations, so the results are consistent and reproducible between experiments. Furthermore, we use a different seed for each of the Monte-Carlo trials in the experiments (*e.g.*, trial 1 is seeded as 1, trial 2 is seeded as 2, and so on for the Cross-Entropy and LTN configurations).

Runtime. The model is trained using an AMD Ryzen Threadripper PRO 5955WX for the CPU and an NVIDIA GeForce RTX 4090 for the GPU. The runtime for training and testing the proposed model is provided in Table V. Note, the number of parameters of the model does not change between the Cross-Entropy and LTN configurations; thus, the runtime during testing is not affected by the configuration used during training.

Loss	MP3D	West Point	Castle Island
Cross-Entropy	0.018s	0.002s	0.005s
SAT (Equiv.)	0.035s	0.003s	0.007s
SAT (Incl.)	0.027s	0.003s	0.006s
SAT (Both.)	0.038s	0.003s	0.008s
Testing (CPU):	0.027s	0.067s	0.207s

TABLE V: Comparison of runtime for the proposed model for training and testing. The reported time is averaged across all epochs and all graphs for training and across all graphs for testing. The average number of (nodes, edges) for the graphs in each dataset are (1283, 7407) for MP3D, (3930, 15057) for West Point, and (11059, 69784) for Castle Island.

B. Evaluation of Spatial Ontology

Here, we provide additional details regarding the evaluation of the generated ontologies using text scoring and text completion against a ground truth ontology based on human judgement. Specifically, we annotated a random sample of relations between low-level and high-level concepts (e.g., a "sink" is found in a "kitchen") as being "likely", "sometimes", and "rarely". We sample 38 pairs where an edge exists between 22 of the pairs (annotated as "likely") and an edge does not exist between 16 of the pairs (annotated as "sometimes" or "rarely"). Next, we compute the accuracy, precision, and recall for the spatial ontologies generated by the text scoring approach (using a temperature of K = 10) and the text completion approach (using $k = \{1, 2, 3, 4, 5\}$) against the human annotated relations; see Table I. For the MP3D dataset, there are 31 high-level and 39 low-level concepts, and for the outdoor datasets, there are 15 high-level and 150 low-level concepts. For the text completion approach, the number of edges is $k|\mathcal{H}|$, which is 78 (k = 3) for MP3D, 42 for West

Point (k = 3), and 28 for Castle Island (k = 2). Note, we do not include rooms labeled unknown when generating the spatial ontology. The number of edges is 99 for the MP3D dataset using the text scoring approach, and the text scoring approach was not applied to the outdoor datasets.

While the accuracy and recall for the text scoring approach significantly outperforms the text completion approach, the precision is significantly better for the text completion approach. Since the ontology is generally sparse, the precision of the spatial ontology is a better performance indicator for the LTN, which is consistent with the results. Recall, the value of k=3 provides the best accuracy for the LTN for the MP3D dataset as shown in Table II, which is also the spatial ontology tied with k=1 for the highest precision (where k=3 has higher accuracy and recall than k=1). Furthermore, this can be observed by comparing the accuracy for grounding the high-level concepts in 3D scene graphs using the text completion versus the text scoring approaches; see Table VI.

In contrast to the text scoring approach, the text completion approach requires choosing a value of k, which restricts the number of edges incident to each low-level concept. The

number of edges incident to each low-level concept in the text scoring approach varies between low-level concepts, which could be preferrable; however, many incident edges to a particular low-level concept indicates the concept may not be useful for distinguishing high-level concepts (*e.g.*, "wall" is common to most rooms in a building making it a poor indicator for a room type).

	SAT (Incl.)		SAT (Both)	
% Kept	Text Scoring	Text Completion	Text Scoring	Text Completion
$\begin{array}{c} 100\% \\ 10\% \\ 1\% \\ 0.4\% \\ 0.2\% \\ 0.1\% \end{array}$	$\begin{array}{c} 12.1 \pm 4.6 \\ 12.1 \pm 4.6 \end{array}$	$\begin{array}{c} \mathbf{23.1 \pm 0.2} \\ \mathbf{23.1 \pm 0.2} \end{array}$	$52.7 \pm 4.6 \\ 50.7 \pm 1.6 \\ 27.3 \pm 2.4 \\ 19.8 \pm 2.8 \\ 18.0 \pm 1.1 \\ 16.7 \pm 1.0$	$52.9 \pm 1.8 \\ 53.2 \pm 1.1 \\ 35.9 \pm 3.3 \\ 29.8 \pm 3.2 \\ 26.8 \pm 2.3 \\ 25.1 \pm 2.2$

TABLE VI: Comparison of accuracy of the proposed model between the "text scoring" and "text completion" methods for generating the spatial ontology for different percentages of data kept during training. Note, the predicates for SAT (Incl.) do not use labeled data; thus, the performance is identical across data kept during training. No data held out during testing. Best in bold.