A Persistent Spatial Semantic Representation for High-level Natural Language Instruction Execution

Anonymous Author(s)
Affiliation
Address
email

Abstract:
Natural language provides an accessible and expressive interface to specify long-term tasks for robotic agents. However, non-experts are likely to specify such tasks with high-level instructions, which abstract over specific robot actions through several layers of abstraction. We propose that key to bridging this gap between language and robot actions over long execution horizons are persistent representations. We propose a persistent spatial semantic representation method, and show how it enables building an agent that performs hierarchical reasoning to effectively execute long-term tasks. We evaluate our approach on the ALFRED benchmark and achieve state-of-the-art results, despite completely avoiding the commonly used step-by-step instructions. https://hlsm-alfred.github.io/

Keywords: vision and language, spatial representations

1 Introduction

Mobile manipulation in a home environment requires addressing multiple challenges, including exploration and making long-term inference about actions to perform. In addition to reasoning, robots require an accessible, yet sufficiently expressive interface to specify their tasks. Natural Language provides an intuitive mechanism for task specification, and coupled with advances in automated language understanding, is increasingly applied to embodied agents [e.g., 1–11].

In this paper, we study the problem of learning to map high-level natural language instructions to low-level mobile manipulation actions in an interactive 3D environment [12]. Existing work largely studies language tightly aligned to the robot actions, either using single-sentence instructions [e.g., 1, 2, 5, 9] or sequences of instructions [13–18]. In contrast, we focus on high-level instructions, which provide more efficient human-robot communication, but require long-horizon reasoning across layers of abstraction to generate actions not explicitly specified in the instruction.

Robust reasoning about manipulation goals from unrestricted high-level natural language instructions has a variety of open challenges. Consider the instruction *secure two discs in a bedroom safe* (Figure 1). The robot must first locate the *safe* in the *bedroom*. It then needs to distribute the actions entailed by *secure* to two objects (*two discs*), each requiring a distinct sequence of actions, but targeting the same *safe*. It is also required to map the verb *secure* to its action space. In parallel, the robot must address core mobile manipulation challenges, and often can only identify required actions as it observes and manipulates the world (e.g., if the *safe* needs to be opened).

We propose to construct and continually update a spatial semantic representation of the world from robot observations (Figure 2). Similar to widely used map representations [19–22], we retain the spatial properties of the environment, allowing the robot to navigate and reason about relations between objects, as required to accomplish its task. We propose the Hierarchical Language-conditioned Spatial Model (HLSM), a hierarchical approach that uses our spatial representation as a long-term memory to solve long-horizon tasks. HLSM consists of a high-level controller that generates subgoals, and a low-level controller that generates sequences of actions to accomplish them. In our example (Figure 1), the sequence of subgoals is ⟨pick up a CD, open the safe, put the CD in the safe, ...⟩, each requiring a sequence of actions. The spatial representation allows selecting subgoals that use previously observed objects outside of the agent's view, or to decide about needed exploration.

Submitted to the 5th Conference on Robot Learning (CoRL 2021). Do not distribute.

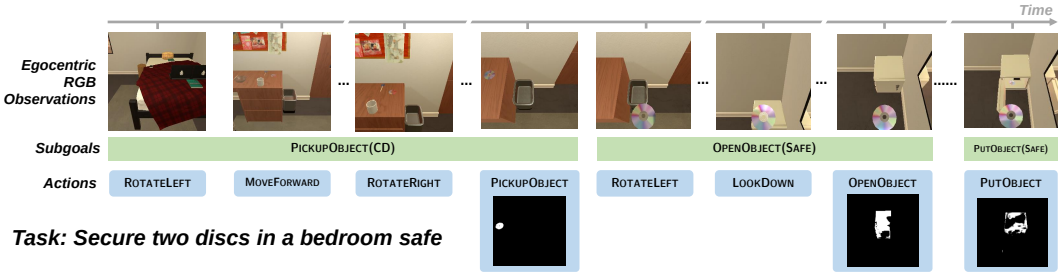

Figure 1: Illustration of the task and our hierarchical formulation. The agent receives a high-level natural language task. It needs to map RGB images to navigation and manipulation actions that accomplish the task.

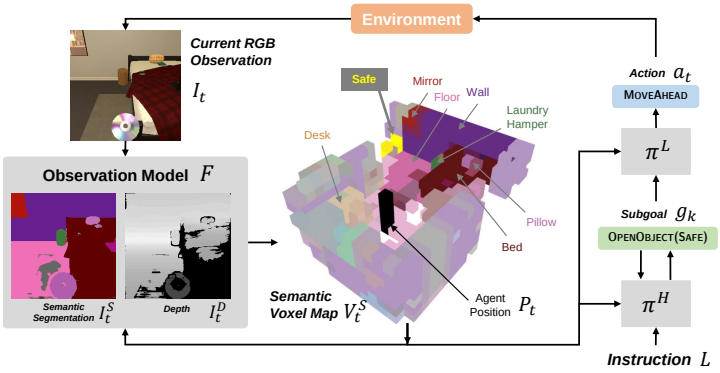

Figure 2: Model architecture consisting of an observation model, high-level controller (π^H), and low-level controller (π^L). The observation model updates the semantic voxel map state representation from RGB observations. π^H predicts the next subgoal given the instruction and the map. π^L outputs a sequence of actions to achieve the subgoal. The semantic voxel map is visualized in the middle with agent position illustrated as a black pillar, ans the current sugoal argument mask in yellow. Other colors are different segmentation classes. Saturated voxels are observed in the current timestep.

Input: Instr. L, Horizon H.
$g_{0,1,2\ldots}, \hat{s}_{-1} \leftarrow$ `null`
$k \leftarrow 1$
Observe initial o_0.
for $t = 0, 1, 2, \ldots H$ **do**
  $\hat{s}_t \leftarrow F(\hat{s}_{t-1}, o_t, g_k)$
  do
    if $g_k =$ `null` **then**
      $g_k \sim \pi^H(L, \hat{s}_t, \langle g_i \rangle_{i<k})$
      if $g_k = g_{\text{STOP}}$ **then**
        End episode
    $a_t \sim \pi^L(g_k, \hat{s}_t)$
    if $a_t = a_{\text{PASS}}$ **then**
      $k \leftarrow k + 1$
    if $a_t = a_{\text{FAIL}}$ **then**
      $g_k \leftarrow$ `null`
  while $a_t \in \{a_{\text{FAIL}}, a_{\text{PASS}}\}$
  Perform a_t, observe o_{t+1}
End episode

Algorithm 1: Execution Flow

We evaluate our approach on the ALFRED [12] benchmark and achieve state-of-the-art results without using the low-level instructions used by previous work [16–18], neither during training nor at test-time. This paper makes three key contributions: (a) a modular representation learning approach for the problem of mapping high-level natural language task descriptions to actions in a 3D environment; (b) a method for utilizing a spatial semantic representation within a hierarchical model for solving mobile manipulation tasks; and (c) state-of-the-art performance on the ALFRED benchmark, even outperforming all approaches that use detailed sequential instructions.

2 Related Work

Natural language has been extensively studied in robotics research, including with focus on instruction [1, 23], reference resolution [24], question generation [25–27], and dialogue [4, 28, 29]. Most work in this area has considered either synthetic instructions of relatively simple goals [7, 30–32], or natural language instructions where all intermediate steps are explained in detail [5, 12–14, 33–37]. In contrast, we focus on high-level instructions, which are more likely in home environments [38].

Representation of world state, action history, and language semantics plays a central role in robot systems and their algorithm design. Symbolic representations have been extensively studied for instruction following agents [1–4, 19, 20, 38–44]. While they simplify the symbol grounding problem and enable robustness, the ontologies on which they rely on are laborious to scale to new, unstructured environments and language. Representation learning presents an alternative by learning to map observations and language directly to actions [5, 8, 9, 11, 13, 33]. World state and language semantics are represented with vectors [13] or by memorizing past observations [8, 17]. Modelling improvements have enabled these approaches to achieve good performance on complex navigation tasks [7, 9, 11, 13, 14, 36], a success that has not yet translated to mobile manipulation [12, 45, 46].

We propose integrating a semantic voxel map state representation within a hierarchical representation learning system. Similar semantic 2D maps have been successfully used in instruction-

following tasks [7, 8, 47–49]. We extend these maps to 3D and show state-of-the-art results on a challenging mobile manipulation benchmark. Our map design is related to sparse metric, topological and semantic maps [10, 19–21, 50] that have enabled grounding symbolic instruction representations. Our map does not impose a topological structure or require reasoning about object instances, instead modelling a distribution over semantic classes for every voxel.

3 Problem Definition

Let $\mathcal{A}$ be the set of agent actions, and $\mathcal{S}$ the set of world states. Given a natural language instruction L and an initial state $s_0 \in \mathcal{S}$, the agent's goal is to generate an execution $\Xi = \langle s_0, a_0, s_1, a_1, \ldots, s_T, a_T \rangle$, where $a_t \in \mathcal{A}$ is an action taken by the agent at time t, $s_t \in \mathcal{S}$ is the state before taking a_t, and $s_{t+1} = \mathcal{T}(s_t, a_t)$ under environment dynamics $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$. The state s_t is defined by the environment layout and the poses and states of all objects and the agent. The agent does not have access to the state s_t, but only to an observation o_t. An observation $o_t = (I_t, P_t, v_t^S, L)$ includes a first-person RGB camera image I_t, the agent's pose P_t, a one-hot encoding of the object class the agent is holding v_t^S, and the instruction L.The task is considered successful if all goal-conditions corresponding to the task L are true at the final state s_T. Partial success is measured as the fraction of goal-conditions that have been achieved.

The ALFRED dataset includes sets of seen and unseen environments. The set of actions $\mathcal{A} = \mathcal{A}_{\text{nav}} \cup \mathcal{A}_{\text{int}}$ includes parameter-free navigation actions $\mathcal{A}_{\text{nav}} = \{\text{MoveAhead}, \text{RotateLeft}, \text{RotateRight}\}$ and interaction actions $\mathcal{A}_{\text{int}} = \{\text{Pickup}, \text{Put}, \text{ToggleOn}, \text{ToggleOff}, \text{Open}, \text{Close}, \text{Slice}\}$ parameterized by a binary mask that identifies the argument of the interaction in the agent's current first-person view. We compute P_t and v_t^S using dead-reckoning from RGB observations and actions.

4 Hierarchical Model with a Persistent Spatial Semantic Representation

We model the agent behavior with a policy π that maps an instruction L and the observation o_t at time t to an action a_t. The policy π is made of an *observation model* F and two controllers: a *high-level controller* π^H and a *low-level controller* π^L. The observation model builds a spatial *state representation* $\hat{s}_t$ that captures the cumulative agent knowledge of the world at time t. $\hat{s}_t$ is used by both π^H for high-level long-horizon task planning, and π^L for near-term reasoning, such as object search, navigation, collision avoidance, and manipulation. Figure 2 illustrates the policy.

The high-level controller π^H computes a probability over *subgoals*. A subgoal g is a tuple $(\texttt{type}, \texttt{arg}^C, \texttt{arg}^M)$, where $\texttt{type} \in \mathcal{A}_{\text{int}}$ is an interaction type (e.g., Open, Pickup), $\texttt{arg}^C$ is the semantic class of the interaction argument (e.g., Safe, CD), and $\texttt{arg}^M$ is a 3D mask identifying the location of the argument instance. In ALFRED, each interaction action in the set $\mathcal{A}_{\text{int}}$ corresponds to a subgoal type. When predicting the k-th subgoal at time t, π^H considers the instruction L, the current state representation $\hat{s}_t$, and the sequence of past subgoals $\langle g_i, \rangle_{i<k}$. During inference, we sample from π^H. Unlike $\arg\max$, sampling allows the agent to re-try the same or different subgoal incase of a potentially random failure (e.g., if a Mug was not found, pick up a Cup).

The low-level controller π^L is given the subgoal g_k as its goal specification at time t. At every timestep $j > t$, π^L maps the state representation $\hat{s}_j$ and subgoal g_k to an action a_j, until it outputs one of the stop actions: a_{PASS} or a_{FAIL} to indicate successful or failed subgoal completion.

Algorithm 1 describes the execution flow. At time $t = 0$ the initial observation o_0 is received. At each timesep, we update the state representation $\hat{s}_t$ (Line 5). If needed, we sample a new subgoal g_k from π^H (Line 8), and then sample an action a_t from π^L. If a_t is a_{PASS}, we increment subgoal counter k (Line 13). If it is a_{FAIL}, we discard the current subgoal k (Line 15). We repeat Lines 8–15 until an executable action a_t is sampled. We execute a_t, receive the next observation (Line 17), and proceed to the next timestep. The episode ends when the subgoal g_{STOP} is sampled (Line 10) or the horizon H is exceeded (Line 18).

4.1 State Representation

The state representation $\hat{s}_t$ at time t captures the agent's current understanding of the state of the world, including the locations of objects observed and the agent's relation to them. The state representation is a tuple $(V_t^S, V_t^O, v_t^S, P_t)$. The semantic map $V_t^S \in [0, 1]^{W \times L \times H \times C}$ is a 3D voxel map

that for every position indicates which of the $c \in [1, C]$ object classes are present in the voxel. The observability map $V_t^O \in \{0,1\}^{W \times L \times H}$ is a 3D voxel map that indicates whether the corresponding position has been observed. The inventory vector $v_t^S \in \{0,1\}^C$ indicates which of the C object classes the agent is currently holding. The agent pose $P_t = (x, y, \omega_p, \omega_y)$ is specified by the 2D position (x, y), pitch angle ω_p, and yaw angle ω_y.

We also compute 2D *state affordance features* $\text{AFFORD}(\hat{s}_t) \in [0,1]^{7 \times W \times L}$ in a top-down view that represent each position with one or more of seven affordance classes {`pickable`, `receptacle`, `togglable`, `openable`, `ground`, `obstacle`, `observed`}. Each $[\text{AFFORD}(\hat{s}_t)]_{(\tau,x,y)} = 1.0$ if at least one of the voxels at position (x, y) has affordance class τ, otherwise it is zero. $\text{AFFORD}(\hat{s}_t)$ is suited for object class agnostic reasoning, for example predicting a pose to pick up an object. [1]

4.2 Observation Model

The observation model $F(\hat{s}_{t-1}, o_t, g_k)$ updates the state representation with new observations. It considers the current subgoal g_k to actively acquire information relevant to g_k. The computation of F consists of three steps: perception, projection, accumulation.

Perception Step We predict semantic segmentation I_t^S and depth map I_t^D from the RGB observation I_t. We use neural networks pre-trained in the ALFRED environment. The semantic segmentation $[I_t^S]_{(u,v)}$ is a distribution over C object classes at pixel (u, v). The depth map $[I_t^D]_{(u,v)}$ is a binned distribution over B bins.[2] We also heuristically compute a binary mask M_t^D that indicates which pixels have confident depth readings. We allow more confidence slack in pixels that correspond to the current subgoal argument $\texttt{arg}_t^C$ according to I_t^S. Appendix **??** provides further details. We use perception models based on the U-Net [51] architecture, but our framework supports other, potentially more powerful models as well (e.g. [52, 53]).

Projection Step We use a pinhole camera model to convert depth I_t^D and segmentation I_t^S to a point cloud that represents each image pixel (u, v) with a 3D position $(x, y, z) \in \mathbb{R}^{W \times L \times H}$ and a semantic distribution $[I_t^S]_{(u,v)}$. We use $\arg\max_B(I_t^D)$ to compute the 3D positions, and discard points at pixels (u, v) when the binary mask value is $[M_t^D]_{(u,v)} = 0$. We construct a discrete semantic voxel map $\hat{V}_t^S \in [0,1]^{W \times L \times H \times C}$, where W, L, and H are the width, height, and length. The value at each voxel $[\hat{V}_t^S]_{(x,y,z)}$ is the element-wise maximum of the segmentation distributions $[I_t^S]_{(u,v)}$ across all points (u, v) within the voxel. We additionally compute a binary observability map $\hat{V}_t^O \in \{0,1\}^{W \times L \times H}$ that indicates the voxels observed at time t. A voxel is observed if it contains points, or if a ray cast from the camera through the voxel centroid has expected depth greater than the distance from the camera to the centroid.

Accumulation Step We integrate $\hat{V}_t^S$ and $\hat{V}_t^O$ into a persistent state representation:

$$V_t^S = \hat{V}_t^S \times \hat{V}_t^S + V_{t-1}^S \times (1 - \hat{V}_t^O) \qquad V_t^O = \max(V_{t-1}^O, \hat{V}_t^O) \ . \tag{1}$$

This operation updates each voxel with the most recent semantic distribution, while retaining the values of all voxels not visible at time t. The output of the observation model is the spatial state representation $\hat{s}_t = (V_t^S, V_t^O, v_t^S, P_t)$. The inventory v_t^S and pose P_t are taken directly from o_t.

4.3 High-level Controller (π^H)

At timestep t, when invoked for the k-th time, the input to π^H is the instruction L, the sequence of past subgoals $\langle g_i \rangle_{i<k}$, and the current state representation $\hat{s}_t$. The output is the next subgoal $g_k = (\texttt{type}_k, \texttt{arg}_k^C, \texttt{arg}_k^M)$. Figure 3 illustrates the high-level controller architecture.

Input Encoding We encode the text L using a pre-trained BERT [54] model that we fine-tune during training. We use the CLS token embedding as the task embedding ϕ^L. We encode the state representation $\hat{s}_t$ to account for classes of all observed objects, and the object that the agent is holding: $\phi^s(\hat{s}_t) = [v_t^S; \max_{(x,y,z)}(V_t^S)]$, where $\max_{(x,y,z)}$ is a max-pooling operation over spatial dimensions and $[\cdot;\cdot]$ denotes concatenation. We compute the representations of previous subgoals as

[1] We assume a known mapping between object semantic classes and affordance classes.

[2] We use B uniformly spaced depth bins $\{0, \Delta_D, 2\Delta_D, \ldots, (B-1)\Delta_D\}$, where Δ_D is a depth resolution. We suggest Δ_D should be less than 50% of the voxel size.

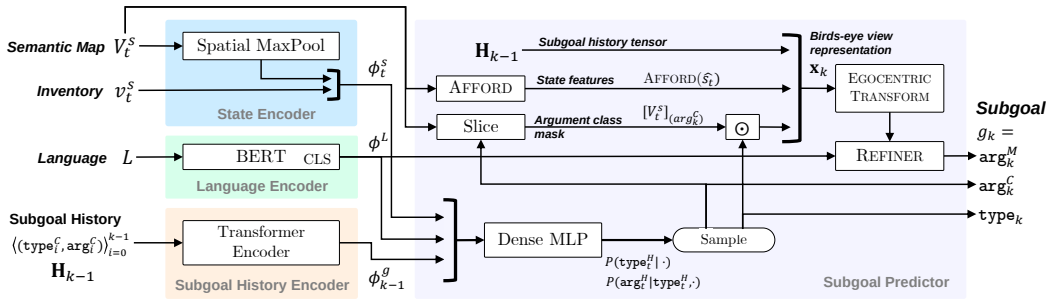

Figure 3: Illustration of the high-level controller π^H (Section 4.3).

$\langle \textsc{Repr}(g_i)\rangle_{i=0}^{k-1}$, where $\textsc{Repr}(g_i)$ is the sum of a sinusoidal positional encoding [55] of index i and learned embeddings for $\texttt{type}_i$ and $\texttt{arg}_i^C$. We process this sequence with a two-layer Transformer autoregressive encoder [55] to compute $\langle \phi_i^g \rangle_{i=0}^{k-1}$. We take ϕ_{k-1}^g as the subgoal history embedding vector. We additionally encode the argument mask information $\texttt{arg}_i^M$ from the subgoal history in an integer-valued subgoal history tensor $\mathbf{H}_{k-1} \in \mathbb{N}^{K \times W \times L}$ where $[\mathbf{H}_{k-1}]_{(\tau,x,y)}$ is the number of times an interaction action type τ was performed at 2D position (x, y) in the birds-eye view:

$$[\mathbf{H}_{k-1}]_{(\tau,x,y)} = \sum_{\substack{i=0\ldots k-1 \\ \texttt{arg}_i^C = \tau}} \max_z([\texttt{arg}_i^M]_{(x,y,z)}) \quad . \tag{2}$$

Subgoal Prediction We concatenate the three representations $\mathbf{h}_{(t,k)} = [\phi^L; \phi_t^s; \phi_{k-1}^g]$. We use a densely connected multi-layer perceptron [56] to predict two distributions $P(\texttt{type}_k \mid \mathbf{h}_{(t,k)})$ and $P(\texttt{arg}_k^C \mid \texttt{type}_k, \mathbf{h}_{(t,k)})$, from which we sample a subgoal type $\texttt{type}_k$ and argument class $\texttt{arg}_k^C$.

The remaining component of the subgoal is the action argument mask $\texttt{arg}_k^M$. Let $[V_t^S]_{(\texttt{arg}_k^C)}$ be a voxel map that only retains the object information for objects of class $\texttt{arg}_k^C$ in the semantic map V_t^S. We refine it to identify a single object instance. We compute a birds-eye view representation:

$$\mathbf{x}_t = [\textsc{Afford}(\hat{s}_t);\ \mathbf{H}_{k-1};\ \max_z([V_t^S]_{(\texttt{arg}_k^C)}) \otimes \mathbb{1}_{\texttt{type}_k}] \tag{3}$$

where $\textsc{Afford}(\hat{s}_t)$ is a birds-eye view state affordance feature map (Section 4.1) and $\mathbb{1}_{\texttt{type}_k}$ is a one-hot encoding of $\texttt{type}_k$.[3] Finally, we compute the 3D argument mask $\texttt{arg}_k^M \in [0,1]^{W \times L \times H}$:

$$\texttt{arg}_k^M = \textsc{Refiner}(\textsc{EgoTransform}(\mathbf{x}_t, P_t), \phi^L) \quad , \tag{4}$$

where $\textsc{EgoTransform}(\mathbf{x}, P_t)$ transforms the map $\mathbf{x}$ to the agent egocentric pose P_t, $\textsc{Refiner}$ is a neural network based on the LingUNet architecture [35], and ϕ^L is the language embedding. The refined $\texttt{arg}_k^M$ is a $[0,1]$-valued 3D mask that identifies the instance of the interaction argument object. If the object is believed to be unobserved, then $\texttt{arg}_k^M$ contains all zeroes. The controller output is the subgoal $g_k = (\texttt{type}_k, \texttt{arg}_k^C, \texttt{arg}_k^M)$.

4.4 Low-level Controller (π^L)

The low-level controller π^L is conditioned on the most recent subgoal $g_k = (\texttt{type}_k, \texttt{arg}_k^C, \texttt{arg}_k^M)$. At time t, it maps the state representation $\hat{s}_t$ to an action a_t. It combines engineered and learned components. Appendix **??** provides the implementation details. The controller π^L invokes a set of procedures: `NavigateTo`, `SampleExplorationPosition`, `SampleInteractionPose`, and `InteractMask`. Their invokation follows a pre-specified execution flow across multiple timesteps. First, we perform a 360° rotation to observe the nearby environment. If no objects of type $\texttt{arg}_k^C$ are observed, we explore the environment by sampling a position $(x, y) = \texttt{SampleExplorationPosition}(\hat{s}_t)$, navigating there using the procedure $\texttt{NavigateTo}(x, y, \hat{s}_t)$, and performing a 360° rotation. We repeat exploration until a voxel in V_t^S contains the class $\texttt{arg}_k^C$ with >50% probability. To interact with an object, we sample an interaction pose $(x, y, \omega_y, \omega_p) = \texttt{SampleInteractionPose}(\hat{s}_t, g_k)$, invoke $\texttt{NavigateTo}(x, y, \hat{s}_t)$ to reach the position (x,y), and

[3] $\otimes$ denotes multiplication of a $X \times Y$ tensor with a K-dimensional vector to obtain a $K \times X \times Y$ tensor. $[\cdot;\cdot;\cdot]$ denotes channel-wise concatenation.

then rotate according to yaw and pitch angles (ω_y, ω_p). Finally, we generate the egocentric interaction mask $\texttt{mask}_t = \texttt{InteractMask}(\hat{s}_t, \texttt{arg}_k^M)$, and output the interaction action $(\texttt{type}_k, \texttt{mask}_t)$.

All procedures use the spatial representation $\hat{s}_t$. `NavigateTo` navigates to a goal position using a value iteration network (VIN) [57] that reasons over obstacle and observability maps from $\hat{s}_t$. `SampleExplorationPosition` samples positions on the boundary of observed space in $\hat{s}_t$. `SampleInteractionPose` uses a learned neural network NAVMODEL to predict a distributon of poses from which the interaction g_k will likely succeed. `InteractMask` uses the segmentation image I_t^S and the 3D argument mask $\texttt{arg}_t^M$ to compute the first-person mask of the target object.

5 Learning

The policy contains four learned models: the segmentation and depth networks, π^H, and the navigation model NAVMODEL used by π^L. We train all networks using supervised learning. We assume access to a training dataset $\mathcal{D} = \{(L^{(j)}, \Xi^{(j)})\}_{j=1}^{N_D}$ of high-level natural language instructions $L^{(j)}$ paired with demonstration execution $\Xi^{(j)}$ in a set of seen environments. Each execution $\Xi^{(j)}$ is a sequence of states and actions $\langle s_0^{(j)}, a_0^{(j)}, \ldots, s_T^{(j)}, a_T^{(j)} \rangle$. We denote N_P the total number of states in $\mathcal{D}$, and N_G the total number of subgoals.

We process $\mathcal{D}$ into three datasets. The perception dataset $\mathcal{D}^P = \{([I]^{(i)}, [I^D]^{(i)}, [I^S]^{(i)}\}_{i=1}^{N_P}$ includes RGB images $[I]^{(i)}$ with ground truth depth $[I^D]^{(i)}$ and segmentation $[I^S]^{(i)}$. The subgoal dataset $\mathcal{D}^g = \{(L^{(i)}, \hat{s}_t^{(i)}, \langle g_j^{(i)} \rangle_{j=0}^k)\}_{i=1}^{N_G}$ contains natural language instructions $L^{(i)}$, state representations $\hat{s}_t^{(i)}$ at the start of k-th subgoal execution, and sequences of the first k subgoals $\langle g_j^{(i)} \rangle_{j=0}^k$ extracted from $\Xi^{(j)}$. The navigation dataset $\mathcal{D}^N = \{(\hat{s}^{(i)}, g^{(i)}, P^{(i)})\}_{i=1}^{N_P}$ consists of state representations $\hat{s}^{(i)}$, subgoals $g^{(i)}$, and agent poses $P^{(i)}$ at the time of taking the interaction action corresponding to subgoal $g^{(i)}$. The state representations $\hat{s}^{(\cdot)}$ in datasets $\mathcal{D}^g$ and $\mathcal{D}^N$ are constructed using the observation model (Section 4.2), but using ground-truth depth and segmentation images.

We train the perception models on $\mathcal{D}^P$ and the π^H on $\mathcal{D}^g$ to predict the k-th subgoal by optimizing cross-entropy losses. We use $\mathcal{D}^N$ to train the navigation model NAVMODEL by optimizing a cross-entropy loss for positions and yaw angles, and an L2 loss for the pitch angle.

6 Experimental Setup

Environment, Data, and Evaluation We evaluate our approach on the ALFRED [12] benchmark. It contains 108 training scenes, 88/4 validation seen/unseen scenes, and 107/8 test seen/unseen scenes. There are 21,023 training tasks, 820/821 validation seen/unseen tasks, and 1533/1529 test seen/unseen tasks. Each task is specified with a high-level natural language instruction. The goal of the agent is to map raw RGB observations to actions to complete the task. ALFRED also provides detailed low-level step-by-step instructions, which simplify the reasoning process. We do not use these instructions for training or evaluation. We collect a training dataset of language-demonstration pairs for learning (Section 5). To extract subgoal sequences, we label each interaction action $a_t = (\texttt{type}_t, \texttt{mask}_t)$ and any preceding navigation actions with a single subgoal of $\texttt{type} = \texttt{type}_t$. We compute the subgoal argument class $\texttt{arg}^C$ and 3D mask $\texttt{arg}^M$ labels from the first-person mask $\texttt{mask}_t$, and ground truth segmentation and depth. Completing a task requires satisfying several goal conditions. Following the common evaluation [58, 59], we report two metrics. *Success rate* (SR) is the fraction of tasks for which all goal conditions were satisfied. *Goal condition rate* (GC) is the fraction of goal-conditions satisfied across all tasks.

Systems We compare our approach, the Hierarchical Language-conditioned Spatial Model (HLSM) to others on the ALFRED leaderboard that only use the high-level instructions. At the time of writing, the only such published approach is HiTUT [46], an approach that uses a flat BERT [54] architecture to model a hierarchical task structure without using a spatial representation. See Appendix **??** for a detailed comparison. We also compare to approaches that use the step-by-step instructions, which puts our method at a disadvantage.

Additionally, we perform ablations and study sensory oracles. To study the observation model, we compare to using sensory oracles for ground truth depth, ground truth segmentation, and both. We report high-level controller ablations that remove the subgoal encoder, language encoder, and state

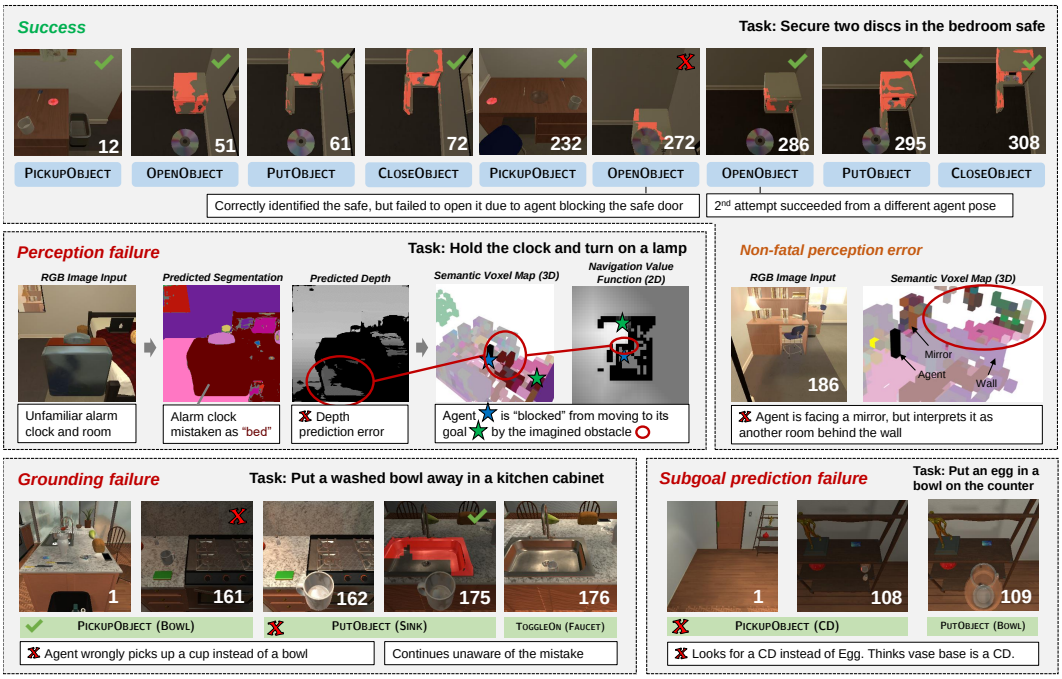

Figure 4: Qualitative results showcasing successes and failures of our approach. **Top row:** snapshots of every interaction action taken during a successful task. Action argument masks are overlaid in red over the RGB images. The white numbers are timesteps. **Middle-right**: illustration of a non-fatal perception error. **Middle-left**: illustration of a fatal perception error. The agent incorrectly interprets the reflection on the alarm clock as an obstacle, causing the agent (blue star) to believe that the path to the goal (green star) is blocked off. This is reflected in the navigation value function computed by the value iteration network (VIN) [57], where black cells are obstacles with value -1. White cell is the goal with value 1. **Bottom-left:** grounding failure. The agent wrongly picks up the cup instead of a bowl. Predicted subgoals are shown in green. **Bottom-right:** high-level controller and percepton failure. π^H predicts the wrong subgoal argument class (CD instead of EGG). The segmentation model then mistakes the vase for a CD.

representation encoder as used for predicting subgoal type $\texttt{type}_k$ and argument class $\texttt{arg}_k^C$, while still using the state representation $\hat{s}_t$ to predict the subgoal argument mask $\texttt{arg}_k^M$. We also study a low-level controller ablation that removes the exploration procedure.

7 Results

Table 1 shows test and validation results. Our approach achieves state-of-the-art performance across both seen and unseen environments in the setting with only high-level instructions. We achieve 6.06% absolute (59.2% relative) improvement in SR on the test unseen split, and 6.7% absolute (36.4% relative) improvement in SR on the test seen split compared to HiTUT G-only.

Our approach performs competitively even when compared to approaches that also use the low-level step-by-step instructions. We achieve 0.86% absolute (5.6% relative) improvement in SR on the test unseen split compared to ABP [16].[4] On the test seen split, our approach performs reasonably well, however ABP [16], E.T. [17], and LWIT [18] perform better, reflecting stronger scene overfitting.

Tables 2 and 3 show development results. Ground truth depth (+ gt depth) provides 3.2%/1% absolute improvement in seen/unseen scenes. Ground truth segmentation (+ gt seg) provides 16.3%/19.3% absolute improvement. Using both (+ gt depth, gt seg) provides 19.1%/30.2% absolute improvement and narrows the seen/unseen gap from 13.2% to 2.1%.

We report high-level controller π^H input encoder ablations. The poor performance without the language encoder reflects task difficulty. Zeroing the input to the subgoal history encoder (but keeping position encodings) reduces performance by only 1.7%/1.8%, showing that knowing the index of the current subgoal in addition to the state representation is often sufficient. Not using the state representation for predicting subgoal type and argument class reduces SR by 1.7% in seen scenes, but increases SR by 1.1% in unseen scenes, achieving state-of-the-art validation-unseen SR

[4]The details and validation results of the ABP and LAV models were not published at the time of writing.

Method	Test				Validation			
	Seen		Unseen		Seen		Unseen	
	SR	GC	SR	GC	SR	GC	SR	GC
Low-level Sequential Instructions + High-level Goal Instruction								
SEQ2SEQ [12]	3.98	9.42	0.39	7.03	3.70	10.00	0.00	6.90
MOCA [45]	22.05	28.29	5.30	14.28	19.15	28.5	3.78	13.4
E.T. [17]	28.77	36.47	5.04	15.01	33.78	42.48	3.17	13.12
E.T. + synth. data [17]	38.42	45.44	8.57	18.6	46.59	52.82	7.32	20.87
LWIT [60]	30.92	45.44	9.42	20.91	33.70	43.10	9.70	23.10
HITUT[46]	21.27	29.97	13.87	20.31	25.24	34.85	12.44	23.71
ABP [16][4]	44.55	51.13	15.43	24.76	-	-	-	-
High-level Goal Instruction Only								
HITUT G-only[46]	18.41	25.27	10.23	20.27	13.63	21.11	11.12	17.89
LAV [61][4]	13.35	23.21	6.38	17.27	-	-	-	-
HLSM (Ours)	**25.11**	**35.79**	**16.29**	**27.24**	**25.00**	**36.37**	**11.80**	**24.70**

Table 1: Test results. Test seen/unseen and validation seen/unseen splits. Top section approaches use sequential step-by-step instructions. The bottom section uses only high-level instructions. Best results **using only high-level instructions** and using both types of instructions are highlighted.

Method	Validation			
	Seen		Unseen	
	SR	GC	SR	GC
HLSM	**25.0**	**36.4**	11.8	24.7
+ gt depth	28.2	36.7	12.8	26.2
+ gt depth, gt seg.	44.1	54.5	42.0	53.2
+ gt seg.	41.3	51.6	31.1	43.2
w/o language enc.	0.6	8.3	0.2	7.6
w/o subg. hist. enc.	23.3	31.8	10.0	22.8
w/o state repr enc.	23.8	34.3	**12.9**	**25.7**
w/o exploration	24.9	35.9	11.3	23.7

Table 2: Development results on validation split. Performance of our full approach, with perception oracles, a perception ablation, π^H ablations, and π^L ablations

Task Type	Validation			
	Seen		Unseen	
	SR	GC	SR	GC
Overall	25.0	36.4	11.8	24.7
Examine	37.2	55.9	22.5	45.7
Pick & Place	42.3	42.3	28.0	28.0
Stack & Place	8.7	23.1	1.8	7.0
Clean & Place	25.0	43.8	10.6	25.3
Cool & Place	11.9	29.9	3.7	22.9
Heat & Place	12.1	29.1	0.0	14.1
Pick 2 & Place	35.5	50.8	12.3	29.2

Table 3: Performance breakdown per task type on the validation split.

of 12.9% on the ALFRED benchmark. Therefore, predicting the sequence of subgoal types and argument classes (i.e., what to do) is at times possible without spatial reasoning, while grounding the subgoal (i.e., where to do it) requires spatial information. Removing random exploration from π^L reduces SR by 0.1%/0.5% on seen/unseen scenes, showing it is not very effective.

Figure 4 illustrates the model behavior, showing both successes and common failures. The main failures in valid unseen scenes are due to (1) perception errors that result in missing or extraneous obstacles or picking up wrong objects; (2) insufficiency of random exploration (e.g., not searching inside cabinets); (3) navigation model errors (e.g., blocking objects from opening); (4) subgoal prediction errors (e.g., picking up wrong objects); and (5) lack of state-aware multi-step planning and backtracking. More qualitative results are available in Appendix **??**.

8 Discussion and Limitations

We showed that a persistent spatial semantic representation enables a hierarchical model to achieve state-of-the-art performance on a challenging instruction-following mobile manipulation task. The main performance bottlenecks include long-horizon exploration, perception generalization to unseen environments, and low-level motion planning for continuous collision avoidance. In terms of learning, incorporating reinforcement learning to train π^H, π^L, and observation model F jointly could improve robustness. We defined the interface to π^L to be faithful to skills available on physical robots, but the exact implementation of π^L is not the focus of our work. Physical deployment would require changes to π^L, and study on robustness to errors in continuous environments, such as localization or motion uncertainty.

References

[1] S. Tellex, T. Kollar, S. Dickerson, M. R. Walter, A. Gopal Banerjee, S. Teller, and N. Roy. "Approaching the Symbol Grounding Problem with Probabilistic Graphical Models. *AI Magazine*, 2011.

[2] C. Matuszek, E. Herbst, L. Zettlemoyer, and D. Fox. Learning to parse natural language commands to a robot control system. In *ISER*, 2012.

[3] D. K. Misra, J. Sung, K. Lee, and A. Saxena. Tell me dave: Context-sensitive grounding of natural language to mobile manipulation instructions. In *RSS*, 2014.

[4] J. Thomason, S. Zhang, R. J. Mooney, and P. Stone. Learning to interpret natural language commands through human-robot dialog. In *IJCAI*, 2015.

[5] D. Misra, J. Langford, and Y. Artzi. Mapping instructions and visual observations to actions with reinforcement learning. In *EMNLP*, 2017.

[6] D. Nyga, S. Roy, R. Paul, D. Park, M. Pomarlan, M. Beetz, and N. Roy. Grounding robot plans from natural language instructions with incomplete world knowledge. In *CoRL*, 2018.

[7] V. Blukis, N. Brukhim, A. Bennet, R. Knepper, and Y. Artzi. Following high-level navigation instructions on a simulated quadcopter with imitation learning. In *RSS*, 2018.

[8] V. Blukis, D. Misra, R. A. Knepper, and Y. Artzi. Mapping navigation instructions to continuous control actions with position-visitation prediction. In *CoRL*, 2018.

[9] V. Blukis, Y. Terme, E. Niklasson, R. A. Knepper, and Y. Artzi. Learning to map natural language instructions to physical quadcopter control using simulated flight. In *CoRL*, 2019.

[10] S. Patki, E. Fahnestock, T. M. Howard, and M. R. Walter. Language-guided semantic mapping and mobile manipulation in partially observable environments. In *CoRL*, 2019.

[11] V. Blukis, R. A. Knepper, and Y. Artzi. Few-shot object grounding and mapping for natural language robot instruction following. In *CoRL*, 2020.

[12] M. Shridhar, J. Thomason, D. Gordon, Y. Bisk, W. Han, R. Mottaghi, L. Zettlemoyer, and D. Fox. Alfred: A benchmark for interpreting grounded instructions for everyday tasks. In *CVPR*, 2020.

[13] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. van den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *CVPR*, 2018.

[14] H. Tan, L. Yu, and M. Bansal. Learning to navigate unseen environments: Back translation with environmental dropout. In *NAACL-HLT*, 2019.

[15] P. Anderson, A. Shrivastava, J. Truong, A. Majumdar, D. Parikh, D. Batra, and S. Lee. Sim-to-real transfer for vision-and-language navigation. In *CoRL*, 2020.

[16] B. Kim, S. Bhambri, K. P. Singh, R. Mottaghi, and J. Choi. Abp, alfred leaderboard, may 10th 2021. https://leaderboard.allenai.org/alfred/submission/c2t70j37q4q5ci4so89g. Accessed: June 16th, 2021.

[17] A. Pashevich, C. Schmid, and C. Sun. Episodic Transformer for Vision-and-Language Navigation, 2021.

[18] Anonymous. Lwit, alfred leaderboard, may 10th 2021. https://leaderboard.allenai.org/alfred/submission/bvppcin94ro4j7j0jqlg. Accessed: May 25th, 2021.

[19] M. R. Walter, S. Hemachandra, B. Homberg, S. Tellex, and S. Teller. Learning Semantic Maps from Natural Language Descriptions. In *RSS*, 2013.

[20] S. Hemachandra, F. Duvallet, T. M. Howard, N. Roy, A. Stentz, and M. R. Walter. Learning models for following natural language directions in unknown environments. In *ICRA*, 2015.

[21] S. Patki, A. F. Daniele, M. R. Walter, and T. M. Howard. Inferring compact representations for efficient natural language understanding of robot instructions. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 6926–6933. IEEE, 2019.

[22] I. Kostavelis and A. Gasteratos. Semantic mapping for mobile robotics tasks: A survey. *Robotics and Autonomous Systems*, 2015.

[23] T. Kollar, S. Tellex, D. Roy, and N. Roy. Toward understanding natural language directions. In *HRI*, 2010.

[24] C. Matuszek, N. FitzGerald, L. Zettlemoyer, L. Bo, and D. Fox. A Joint Model of Language and Perception for Grounded Attribute Learning. In *ICML*, 2012.

[25] S. Tellex, R. Knepper, A. Li, D. Rus, and N. Roy. Asking for help using inverse semantics. In *RSS*, 2014.

[26] R. A. Knepper, S. Tellex, A. Li, N. Roy, and D. Rus. Recovering from Failure by Asking for Help. *Autonomous Robots*, 2015.

[27] Z. Gong and Y. Zhang. Temporal spatial inverse semantics for robots communicating with humans. In *ICRA*, 2018.

[28] T. Brick and M. Scheutz. Incremental natural language processing for hri. In *HRI*, 2007.

[29] S. Tellex, P. Thaker, R. Deits, D. Simeonov, T. Kollar, and N. Roy. Toward information theoretic human-robot dialog. *Robotics*, 2013.

[30] K. M. Hermann, F. Hill, S. Green, F. Wang, R. Faulkner, H. Soyer, D. Szepesvari, W. Czarnecki, M. Jaderberg, D. Teplyashin, et al. Grounded language learning in a simulated 3d world. *arXiv preprint arXiv:1706.06551*, 2017.

[31] D. S. Chaplot, K. M. Sathyendra, R. K. Pasumarthi, D. Rajagopal, and R. Salakhutdinov. Gated-attention architectures for task-oriented language grounding. *AAAI*, 2018.

[32] C. Paxton, Y. Bisk, J. Thomason, A. Byravan, and D. Fox. Prospection: Interpretable plans from language by predicting the future. In *ICRA*, 2019.

[33] A. Suhr and Y. Artzi. Situated mapping of sequential instructions to actions with single-step reward observation. In *ACL*, 2018.

[34] D. Fried, R. Hu, V. Cirik, A. Rohrbach, J. Andreas, L.-P. Morency, T. Berg-Kirkpatrick, K. Saenko, D. Klein, and T. Darrell. Speaker-follower models for vision-and-language navigation. In *NeurIPS*, 2018.

[35] D. Misra, A. Bennett, V. Blukis, E. Niklasson, M. Shatkin, and Y. Artzi. Mapping instructions to actions in 3D environments with visual goal prediction. In *EMNLP*, 2018.

[36] C.-Y. Ma, J. Lu, Z. Wu, G. AlRegib, Z. Kira, R. Socher, and C. Xiong. Self-monitoring navigation agent via auxiliary progress estimation. In *ICLR*, 2019.

[37] H. Chen, A. Suhr, D. Misra, and Y. Artzi. Touchdown: Natural language navigation and spatial reasoning in visual street environments. In *CVPR*, 2019.

[38] D. K. Misra, K. Tao, P. Liang, and A. Saxena. Environment-driven lexicon induction for high-level instructions. In *ACL*, 2015.

[39] M. MacMahon, B. Stankiewicz, and B. Kuipers. Walk the talk: Connecting language, knowledge, and action in route instructions. In *AAAI*, 2006.

[40] S. R. K. Branavan, L. S. Zettlemoyer, and R. Barzilay. Reading between the lines: Learning to map high-level instructions to commands. In *ACL*, 2010.

[41] S. Tellex, T. Kollar, S. Dickerson, M. Walter, A. Banerjee, S. Teller, and N. Roy. Understanding natural language commands for robotic navigation and mobile manipulation. In *AAAI*, 2011.

[42] F. Duvallet, T. Kollar, and A. Stentz. Imitation learning for natural language direction following through unknown environments. In *ICRA*, 2013.

[43] Y. Artzi and L. Zettlemoyer. Weakly supervised learning of semantic parsers for mapping instructions to actions. *TACL*, 2013.

[44] E. C. Williams, N. Gopalan, M. Rhee, and S. Tellex. Learning to parse natural language to grounded reward functions with weak supervision. In *ICRA*, 2018.

[45] K. P. Singh, S. Bhambri, B. Kim, R. Mottaghi, and J. Choi. Moca: A modular object-centric approach for interactive instruction following. *arXiv preprint arXiv:2012.03208*, 2020.

[46] Y. Zhang and J. Chai. Hierarchical task learning from language instructions with unified transformers and self-monitoring. *ACL Findings*, 2021.

[47] D. Gordon, A. Kembhavi, M. Rastegari, J. Redmon, D. Fox, and A. Farhadi. Iqa: Visual question answering in interactive environments. In *CVPR*, 2018.

[48] P. Anderson, A. Shrivastava, D. Parikh, D. Batra, and S. Lee. Chasing ghosts: Instruction following as bayesian state tracking. In *NeurIPS*, 2019.

[49] H. Saha, F. Fotouhi, Q. Liu, and S. Sarkar. A modular vision language navigation and manipulation framework for long horizon compositional tasks in indoor environment. *arXiv preprint arXiv:2101.07891*, 2021.

[50] S. Hemachandra, M. R. Walter, S. Tellex, and S. Teller. Learning spatial-semantic representations from natural language descriptions and scene classifications. In *ICRA*, 2014.

[51] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.

[52] J. McCormac, A. Handa, S. Leutenegger, and A. J. Davison. Scenenet rgb-d: Can 5m synthetic images beat generic imagenet pre-training on indoor segmentation? In *ICCV*, 2017.

[53] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick. Detectron2. https://github.com/facebookresearch/detectron2, 2019.

[54] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019.

[55] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS*, 2017.

[56] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.

[57] A. Tamar, Y. Wu, G. Thomas, S. Levine, and P. Abbeel. Value iteration networks. In *NeurIPS*, 2016.

[58] M. Shridhar and D. Hsu. Interactive visual grounding of referring expressions for human-robot interaction. In *RSS*, 2018.

[59] M. Shridhar, J. Thomason, D. Gordon, Y. Bisk, W. Han, R. Mottaghi, L. Zettlemoyer, and D. Fox. Alfred leaderboard. https://leaderboard.allenai.org/alfred/. Accessed: June 16th, 2021.

[60] V.-Q. Nguyen, M. Suganuma, and T. Okatani. Look wide and interpret twice: Improving performance on interactive instruction-following tasks. In *IJCAI*, 2021.

[61] K. Nottingham, L. Liang, D. Shin, C. C. Fowlkes, R. Fox, and S. Singh. Lav, alfred leaderboard, may 10th 2021. https://leaderboard.allenai.org/alfred/submission/c2cm7eranqs9puf9uvjg. Accessed: May 25th, 2021.

[62] H.-T. L. Chiang, A. Faust, M. Fiser, and A. Francis. Learning navigation behaviors end-to-end with autorl. *RA-L*, 2019.

[63] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox. Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes. *arXiv preprint arXiv:2103.14127*, 2021.

[64] D. Meagher. Geometric modeling using octree encoding. *Computer graphics and image processing*, 1982.

[65] T. Takikawa, J. Litalien, K. Yin, K. Kreis, C. Loop, D. Nowrouzezahrai, A. Jacobson, M. McGuire, and S. Fidler. Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In *CVPR*, 2021.

[66] H. P. Grice. Logic and conversation. In *Speech acts*. 1975.

[67] J. Roh, C. Paxton, A. Pronobis, A. Farhadi, and D. Fox. Conditional driving from natural language instructions. In *CoRL*, pages 540–551, 2020.

[68] C. Matuszek, D. Fox, and K. Koscher. Following directions using statistical machine translation. In *HRI*, 2010.

[69] N. Gopalan, D. Arumugam, L. L. Wong, and S. Tellex. Sequence-to-sequence language grounding of non-markovian task specifications. In *RSS*, 2018.

[70] D. Bahdanau, F. Hill, J. Leike, E. Hughes, A. Hosseini, P. Kohli, and E. Grefenstette. Learning to understand goal specifications by modelling reward. In *ICLR*, 2018.

[71] P. Goyal, S. Niekum, and R. J. Mooney. Pixl2r: Guiding reinforcement learning using natural language by mapping pixels to rewards. *CoRL*, 2020.

[72] J. Krantz, E. Wijmans, A. Majumdar, D. Batra, and S. Lee. Beyond the nav-graph: Vision-and-language navigation in continuous environments. *arXiv preprint arXiv:2004.02857*, 2020.

[73] V. Jain, G. Magalhaes, A. Ku, A. Vaswani, E. Ie, and J. Baldridge. Stay on the path: Instruction fidelity in vision-and-language navigation. In *ACL*, 2019.

[74] A. Ku, P. Anderson, R. Patel, E. Ie, and J. Baldridge. Room-across-room: Multilingual vision-and-language navigation with dense spatiotemporal grounding. In *EMNLP*, 2020.

[75] M. Persson, T. Duckett, C. Valgren, and A. Lilienthal. Probabilistic semantic mapping with a virtual sensor for building/nature detection. In *Computational Intelligence in Robotics and Automation*, 2007.

[76] H. Zender, O. M. Mozos, P. Jensfelt, G.-J. Kruijff, and W. Burgard. Conceptual spatial representations for indoor mobile robots. *Robotics and Autonomous Systems*, 2008.

[77] A. Pronobis, O. Martinez Mozos, B. Caputo, and P. Jensfelt. Multi-modal semantic place classification. *IJRR*, 2010.

[78] A. Pronobis. *Semantic mapping with mobile robots*. PhD thesis, KTH Royal Institute of Technology, 2011.