

## Continuous Scene Representations for Embodied AI

Samir Yitzhak Gadre<sup>1\*</sup> Kiana Ehsani<sup>2</sup> Shuran Song<sup>1</sup> Roozbeh Mottaghi<sup>2,3</sup>  
<sup>1</sup> Columbia University <sup>2</sup> Allen Institute for AI <sup>3</sup> University of Washington  
[prior.allenai.org/projects/csr](http://prior.allenai.org/projects/csr)

### Abstract

We propose *Continuous Scene Representations (CSR)*, a scene representation constructed by an embodied agent navigating within a space, where objects and their relationships are modeled by continuous valued embeddings. Our method captures feature relationships between objects, composes them into a graph structure on-the-fly, and situates an embodied agent within the representation. Our key insight is to embed pair-wise relationships between objects in a latent space. This allows for a richer representation compared to discrete relations (e.g., [SUPPORT], [NEXT-TO]) commonly used for building scene representations. CSR can track objects as the agent moves in a scene, update the representation accordingly, and detect changes in room configurations. Using CSR, we outperform state-of-the-art approaches for the challenging downstream task of visual room rearrangement, without any task specific training. Moreover, we show the learned embeddings capture salient spatial details of the scene and show applicability to real world data. A summary video and code is available at [prior.allenai.org/projects/csr](http://prior.allenai.org/projects/csr).

### 1. Introduction

To operate within a scene, embodied agents require a comprehensive representation of their surroundings. Such perceptual understanding should not be limited to determining object identities, but should rather capture relationships between objects and between the agent and its surroundings. An expressive scene representation should also allow for downstream task completion of interactive tasks without additional training (i.e., zero-shot inference).

Towards building such representations, scene graphs [2, 29, 65] are a candidate as they provide a compact and explicit description of a scene. In a typical scene graph pipeline, it is common to define a set of relationship labels, use them to manually annotate connections between objects in frames, and train a model to infer target graphs. However,

\*Work done while SYG was a research intern at AI2. Correspondence to [sy@cs.columbia.edu](mailto:sy@cs.columbia.edu)

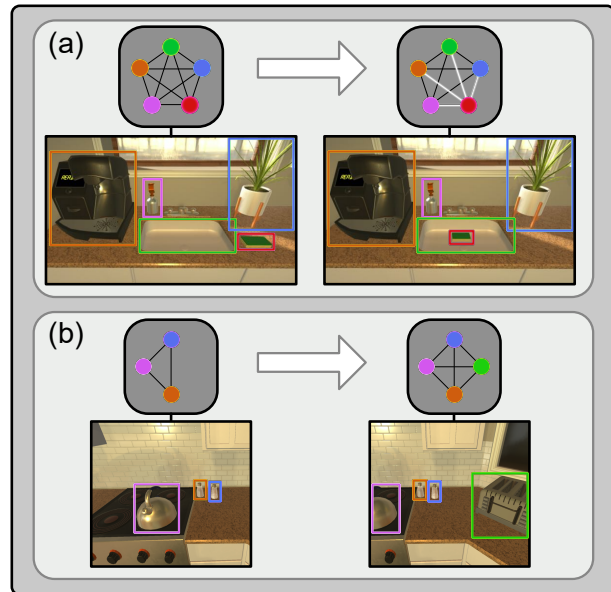


Figure 1. **Representing relationships.** Our work, Continuous Scene Representations, encodes object relationships into a graph with nodes and edges represented as feature vectors. **(a)** If an object is shuffled between two trajectories, different embeddings allow for change detection. **(b)** As the agent moves within a trajectory, new nodes and edges are populated.

once labels are defined, there are inherent limitations. Once trained, models are restricted to a fixed set of relationships.

Even if all relationships are modeled, how well can pre-defined discrete symbols represent the complex relationships between objects? Consider [SUPPORT], a common semantic relation used in the literature [32, 62, 64]. The relationship “table supports mug” indicates the mug is on the table, but it does not capture *where*. In practice, the agent behavior may depend on this information.

Beyond the inherent limitations of modeling relationships as discrete symbols, scene graphs used in the literature are typically static [28, 34, 61] (i.e., they represent a snapshot of a scene or a bundle of frames). In embodied settings, new objects are observed as an agent explores, and the scene representation should update on-the-fly. If an agent returns to a position, it should determine if objects have moved.

Inspired by these observations, we develop a scene representation that is more suitable for embodied AI tasks. We propose Continuous Scene Representations (CSR), a novel approach to construct a scene representation from egocentric RGB images as an embodied agent explores a scene. To address limitations of traditional scene graphs, our goal is to represent relations between objects as continuous vectors and update the representation on-the-fly as the agent moves. To enable downstream interactive task execution without additional training, we also present a simple strategy for planning with respect to the representation.

Constructing such a scene representation interactively poses various challenges. The graph should accommodate the new objects, and relationships with other objects should be inferred. A successful algorithm should also determine when detections correspond to the same object even in different views. To tackle these challenges, our idea is to learn *object relational embeddings* via a contrastive loss to represent the nodes and edges of a scene representation. In this scheme, modeling is not constrained to a pre-defined set of symbols. The agent maintains a memory of previously encountered embeddings. As the agent extracts new embeddings from egocentric observations, it compares them to the memory to determine which embeddings are new and which already exist. This can be used to detect changes in the scene and update the representation as shown in Fig. 1.

We perform experiments using the AI2-THOR [30] framework and the YCB-Video dataset [59]. We support experimentally that (1) *without any task specific training*, a simple planning approach that employs CSR as the underlying representation, outperforms a map-based, reinforcement and imitation learning baseline trained directly on the task of room rearrangement [4, 52]. (2) CSR is able to capture commonly used discrete relations that can be extracted via linear probes. (3) CSR effectively captures spatial relationship between objects within a scene. (4) Without any fine-tuning or hyperparameter tuning, a CSR trained on AI2-THOR is able to track objects over time in real world YCB-Video [59], which contains objects unseen during training.

## 2. Related Work

**Scene graphs.** There are various methods in the literature that use scene graphs to develop richer scene understanding models. It is common to build scene graphs for static images [14, 32, 34, 61, 62, 64]. One of the major issues of these approaches is the lack of temporal information (i.e., there is no mechanism to identify or track changes in a scene over time). To overcome this issue, other approaches create scene graphs from videos [13, 28, 33, 36, 45, 49]. These types of scene graphs capture temporal information; however, they are created using pre-recorded videos, hence, not suitable for embodied tasks that involve observations dependent on on-the-fly actions. The mentioned image-based

and video-based scene graphs capture only 2D relations between objects in a scene. However, others propose to create 3D scene graphs that encode 3D relationships between the components of a scene [2, 12, 20, 51, 66]. Our method is closer to the approaches that create scene graphs in embodied or physics based settings [6, 17, 31, 39, 67]. There are two important differences between our work and these approaches. First, we do not have a pre-defined set of relationships (e.g., [SUPPORT], [NEXT-TO]). Instead, we represent the relationships by a learned embedding. Second, our approach does not require object category information at inference.

**Mapping and localization.** Scene graphs are a form of abstraction for metric-semantic maps, so we discuss examples of mapping work. Classical mapping and localization work in the robotics literature [16, 19, 35] relies on low-level geometric features and does not encode semantic information. There are object-based maps that encode object-level semantics and the relationships [3, 7, 42]. However, they still operate based on the assumption of static environments. Dynamic SLAM approaches [41, 48, 60] have been proposed to handle moving objects in a scene. However, they primarily focus on table-top scenes. Recently, approaches have been designed to handle more complex dynamic scenes [40, 55]. There are few major differences with this work. (1) Our method produces an abstract sparse representation rather than a dense representation. (2) We create a continuous representation of the scene. (3) Methods such as [55] require a full sequence of observations as input, which is not suitable for embodied applications since the observations vary depending on the taken actions.

**Embodied tasks.** Recently, various embodied tasks such as navigation [5, 8, 37, 54, 56, 63], instruction following [1, 47], embodied question answering [15, 23], object manipulation [18, 21, 22, 46], and room rearrangement [4, 52] have been of interest. These tasks benefit substantially from some form of memory to store the state of the scene. Neural graph-based [9, 10, 43, 58] and dense representations [8, 25] have been proposed to construct the map of the environment. These methods mostly focus on creating occupancy maps of free space and obstacles or the reachability of different nodes. In our approach, we propose a scene representation that encodes object relations by feature embedding and evolves based on new observations and agent actions.

## 3. Scene Representation

Embodied AI requires an agent to have an understanding of the environment and to update its knowledge based on new observations. In this paper, our goal is to model scenes with our Continuous Scene Representations (CSR), which can then be used in downstream visual and embodied tasks (Fig. 2). We start by giving an overview of CSR (Sec. 3.1). Then we address learning and updating the scene represen-

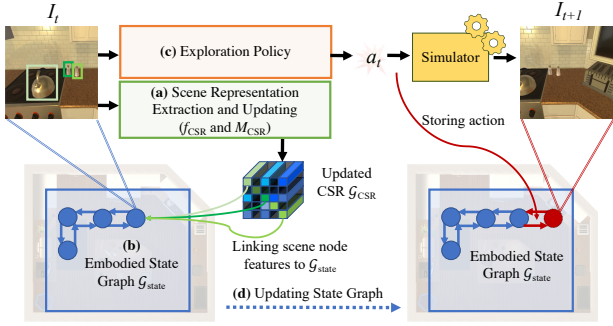


Figure 2. **Overview.** (a) Given an egocentric observation  $I_t$ , the scene representation module extracts nodes and edges to update the continuous scene representation (CSR). (b) References to observed nodes in the current image are stored in an embodied state graph, which links agent poses to CSR nodes that are seen from that pose. (c) An exploration policy infers an action to more comprehensively observe the scene and build a more complete CSR. (d) As the agent moves, the embodied state graph is updated with the action that leads to the next agent pose.

tation (Sec. 3.2). Finally, we describe strategies for linking CSR to agent embodiment (Sec. 3.3) and detecting changes in the scene (Sec. 3.4).

### 3.1. Overview

Given a navigating agent in a scene (e.g., a robot in a kitchen), our goal is to create a CSR of the room. A CSR is a graph where nodes represent objects and edges represent the relationships between them. Both nodes and edges are represented by continuous valued features. As the agent moves, the CSR updates to accommodate observed objects and relationships. The agent is situated within the scene representation to allow for planning.

**Features for nodes and edges.** To motivate continuous valued nodes and edges, consider a scene graph that models the [SUPPORT] relation. In a kitchen scene, the graph might include a directed edge from [STOVE] to [KETTLE], representing the “kettle is on the stove”. This representation throws away useful information like the appearance of the relationship or the relative positions of the objects. If instead, the representation was a continuous vector, it could preserve more information. Moreover, a continuous representation can implicitly encode many other relations (such as [NEXT-TO] and [INSIDE]), which could be recovered when necessary using light-weight, trainable linear projection heads. Hence, learning continuous features for relations could provide a flexible representation that could conceivably adapt based on an agent’s needs.

**Scene representation consistency.** To motivate consistency of the representation, consider again the example of the kettle. As an agent navigates in the scene, it should determine if a kettle is the same object from previous

timesteps to accurately model the underlying relationships. In the terminology of CSR, as the agent moves, it must decide if observed objects match with existing nodes or should spawn new nodes. Edges should be propagated accordingly. This is challenging, especially in our setting that uses only RGB images as input.

**Adding the embodiment.** CSR does not directly model the agent or its actions. Such information must be captured if an agent is to navigate to a particular CSR node. For example, given a node corresponding to [KETTLE], it would be useful to retrieve the sequence of actions to get to a position where the kettle is observable.

**Scene representation notation.** A CSR is a graph  $\mathcal{G}_{CSR}(\mathcal{N}_{CSR}, \mathcal{E}_{CSR})$ , where  $\mathcal{N}_{CSR}, \mathcal{E}_{CSR}$  are the sets of nodes and edges, respectively. Each node and edge is a vector in  $\mathbb{R}^L$ ,  $L$  being the feature dimension. A local scene representation  $\mathcal{G}_{CSR}^t$  is the continuous scene representation constructed based on the observation of the agent at time  $t$ . There is a detection function that takes  $I_t \in \mathbb{R}^{H \times W \times 3}$  and extracts bounding box masks  $\{r_1, r_2, \dots, r_n\}$ ,  $r_i \in \{0, 1\}^{H \times W}$ , one for each object. Each mask  $r$  has value one within the box, else zero. A CSR encoder  $f_{CSR}(I_t, r_i, r_j) \mapsto y_{i,j} \in \mathbb{R}^L$  is introduced to map directed pairs of objects to their feature representations, where  $y_{i,j}$  is the continuous valued vector representing the feature between node  $i$  and  $j$ . If  $i = j$ ,  $y_{i,j} \in \mathcal{N}_{CSR}$ , otherwise  $y_{i,j} \in \mathcal{E}_{CSR}$ . A matching function  $M_{CSR}$ , aggregates the local observation into  $\mathcal{G}_{CSR}$  by taking an existing CSR and a local scene representation  $\mathcal{G}_{CSR}^t$  and updating  $\mathcal{G}_{CSR}$ .

**Embodied state graph notation.** An embodied state graph  $\mathcal{G}_{state}(\mathcal{S}, \mathcal{T})$ , contains states  $\mathcal{S}$  pointing to nodes in the CSR, with one state for every egocentric observation, and action transitions  $\mathcal{T}$  between states. The nodes  $\mathcal{S}$  keep track of references to the objects and relationships in  $\mathcal{G}_{CSR}$  that are seen by the agent at that state. A transition takes place via agent actions (e.g., MOVEFORWARD). Hence,  $\mathcal{G}_{state}$  stores information about how the agent should reach particular visual observations.

### 3.2. Creating the Scene Representation

In this section, we explain how to create a CSR or  $\mathcal{G}_{CSR}$ , which are used interchangeably. We first discuss building a local representation from the current egocentric view and learning node and edge representations. Since the agent explores over time, we then describe how the local information is aggregated.

**Building the local scene representation.** In conventional scene representations, nodes are represented by class labels (e.g., [MUG] and [SINK]), and edges by attributes relating classes (e.g., [UNDER] and [INSIDE]). Instead CSR uses continuous features for nodes and edges. To get box region proposals corresponding to nodes, a detector is employed on the current image  $I_t$ , which yields  $n$  detections. In prac-




Figure 3. **Building the Local Graph.** A detector extracts  $n$  boxes and all  $n^2$  pairs of boxes are batched and fed to the CSR encoder. This yields the CSR representation for this image. Node features are in green and edge features are in blue.

time we use a Faster-RCNN model; however, class labels are discarded. All  $n^2$  pairs of regions are mapped into a latent space, using the CSR encoder  $f_{\text{CSR}}$  as depicted in Fig. 3. A batch of  $n^2$  items is fed through  $f_{\text{CSR}}$ , with each item in the batch containing (1) a unique ordered pair of regions and (2)  $I_t$ . The result is  $n^2$  feature vectors, which form the local scene representation  $\mathcal{G}_{\text{CSR}}^t$  with  $n$  node features and  $n^2 - n$  edge features. While CSR feature extraction is quadratic in the number of detections, all computations are independent and hence parallelizable.

What is a good objective to learn the CSR encoder  $f_{\text{CSR}}$ ? Our high-level idea is to embed the same relationship from different view points to the same location in the embedding space, assuming objects do not move within a trajectory. Consider again the kettle on the stove. The relationship of the kettle to the stove can be observed from different angles, with a variety of objects in the background. However, the global characteristics of the relationship do not change across views. The same argument holds for nodes.

To learn features capable of registering relations from different viewpoints, we take inspiration from the momentum contrast approaches [11, 26]. Given two views of the same pair of objects, a representation is extracted from the first view via  $f_{\text{CSR}}$  and from the second view via a momentum encoder, which has the same architecture and initialization as  $f_{\text{CSR}}$ , but is updated via a slow moving average of  $f_{\text{CSR}}$  weights instead of by gradient updates. InfoNCE loss [50] is used to enforce similarity of relations in feature space,  $\mathcal{L}_{\text{CSR}} = -\log \frac{\exp(\text{CosSim}(q, k^+)/\tau)}{\sum_{i=0}^K \exp(\text{CosSim}(q, k^i)/\tau)}$ . Here  $q$  is the feature representation of a node or edge outputted by  $f_{\text{CSR}}$ ,  $k^+$  is the momentum encoder representation for the second view of the same relationship, each  $k^i$  is a negative drawn from a bank of  $K$  other momentum encoded relationships,  $\tau$  is a softmax temperature scaling parameter, and CosSim is the dot product between L2 normalized features. Notably this objective encourages learning features that (1) have high CosSim when they represent that same relationship, (2) are multi-view consistent in the feature space, and (3) are different from those of other relationships. See Fig. 4

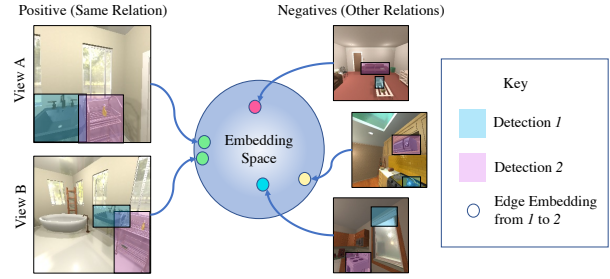


Figure 4. **Embedding Space.** The loss encourages the same relationship observed from different views to be close in feature space and far from other relationships.

for visualization of these attributes.

**From local to global scene representation.** Considering the agent is navigating over time, it is not sufficient to capture information from a single view. As the agent traverses, it is necessary to determine if node representations are new objects or correspond to previously seen objects. To handle such complexity, we design a way to update nodes and edges based on the current local observation.

Our insight is to use a matching function  $M_{\text{CSR}}$  between the node embeddings from the current local scene representation  $\mathcal{G}_{\text{CSR}}^t$  and  $\mathcal{G}_{\text{CSR}}$  to update  $\mathcal{G}_{\text{CSR}}$ . At the first timestep  $\mathcal{G}_{\text{CSR}} \leftarrow \mathcal{G}_{\text{CSR}}^t$ . For subsequent timesteps, node features corresponding to the same object should have high CosSim. As shown in Fig. 5,  $M_{\text{CSR}}$  first computes CosSim between pairs of features in  $\mathcal{N}_{\text{CSR}}$  and  $\mathcal{N}_{\text{CSR}}^t$ , which yields a score matrix. Next a maximal linear score assignment (Hungarian match) creates an initial matching of nodes. If the score of a match exceeds a threshold, the nodes are considered a true match (more details in Sec. 4.2). The matched local node features are then averaged with their matches to update the CSR nodes. For edges, if a new relationship is observed it is added otherwise it is averaged into the representation in an analogous fashion to the nodes. Unmatched nodes and corresponding edges from  $\mathcal{G}_{\text{CSR}}^t$  are then directly added as new nodes and relations.

### 3.3. Linking CSR to Agent States for Planning

We have discussed how to create our representation, which captures object and their relationships in the scene. How should this representation be used to solve embodied tasks? While using the representation as additional input for downstream training is an option, we present a simple alternative. Namely, we discuss the creation of an embodied state graph  $\mathcal{G}_{\text{state}}$  to enable planning within the CSR. As the agent traverses, we keep track of (1) scene nodes that are observed in each egocentric view, which initialize states in  $\mathcal{G}_{\text{state}}$  and (2) the actions that lead to changes in views, which initialize the transitions in  $\mathcal{G}_{\text{state}}$ . To navigate to a specific scene feature in  $\mathcal{G}_{\text{CSR}}$ , the agent can lookup a corresponding state where this feature was observed, and navigate to it by

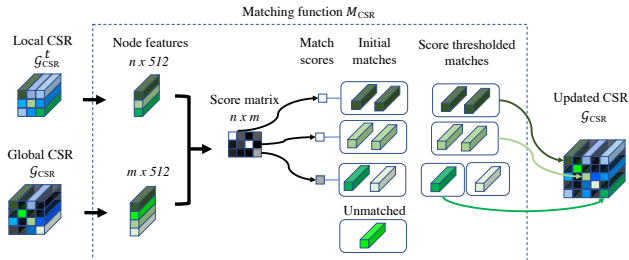


Figure 5. **Matching.** A function  $M_{CSR}$  takes the local and the existing global CSR to update the global CSR.

planning actions from its current state to the target state in  $\mathcal{G}_{state}$ . Since the states in  $\mathcal{G}_{state}$  are linked via actions, the path and actions to the target object can be recovered using a standard breath first search planner over the state graph. Note, this is a simple planning strategy, is independent of CSR creation. In the presence of noise, it is also possible to employ more sophisticated closed-loop planning.

### 3.4. Detecting Changes

In practice, scenes typically change over time. For example, kitchen equipment frequently moves. Given two trajectories, each observing a different static configuration, it is important to determine the objects that changed.

**Motivating between-trajectory object correspondences.** Our insight is to detect changes in the scene by leveraging object correspondences between trajectories. If a kettle is observed on a stove in one trajectory and on a countertop in another, the agent should know that plausibly this kettle has moved. While our scene representation captures nodes and edges in the context of the scene, we would also like to detect objects that shuffle over time. This motivates finding object correspondences between trajectories.

**Between-trajectory object features.** Our idea is to augment the scene nodes of  $\mathcal{G}_{CSR}$  with a object feature vector that is responsible for matching the object instance between trajectories. This feature is learned similarly to the scene node features. However, instead of positives for the contrastive loss coming from different views in the same static scene, they come from different views across varied initializations of the scene. The representation learner cannot rely on consistent scene information to minimize the loss, but should instead focus only on embedding characteristics of the object instances. Using a similar Hungarian matching scheme, as in Sec. 3.2, object features are matched between the two trajectories and thresholded based on CosSim.

**Using low feature similarity to detect change.** Recall, when learning node and edge features, we leveraged that the scene was static. Based on the object matches, detecting violations to the static assumption provides signal for what has moved. Concretely, using the object feature assignments, the CosSim between matched scene nodes and edges (i.e., scene features before object features were intro-

duced) is computed. If scores are lower than a threshold there is signal the visual object has moved (i.e., the underlying spatial relationship has changed). In summary, the object features are used to match object instances between trajectories and then low scene node/edge feature similarity is used to determine what moved.

## 4. Experiments

In this section, we first present our datasets (Sec. 4.1). We show competitive performance of Continuous Scene Representations (CSR) on 2-phase visual room rearrangement [4, 52], *without training for the task* (Sec. 4.2), CSR encodes spatial information that can be recovered via linear probes (Sec. 4.3), CSR can retrieve layout-consistent views of the scene against hard negatives (Sec. 4.4), and CSR is effective for tracking objects in real-world data (Sec. 4.5).

### 4.1. Datasets

For rearrangement experiments (Sec. 4.2), we use the RoomR dataset introduced by Weihs *et al.* [52]. Its test and val sets each contain 1000 room rearrangements.

To train the CSR encoder  $f_{CSR}$ , which embeds nodes and edges, we generate a dataset using AI2-THOR [30]. For each of the 80 train rooms, we render 20 random agent poses in 5 different room initialization with random lighting, texture, and object locations. RGB images for this and all following datasets are  $224 \times 224$ . We mine pairs of objects (i.e., edges) that are viewed from different view points in each room initialization to get a dataset of 900K samples.

We additionally create smaller datasets for our various experiments. We create two image datasets of  $\sim 5k$  images in 80/20 train/test splits with spatial relationship labels for pairs of ground truth boxes (e.g., “ $x$  supports  $y$ ”). These are used in linear probes for our edge representations (Sec. 4.3). A test dataset of 4.4K triplets of images is generated, where a query and a positive image have the same static scene layout and a negative image presents the scene with one object moved. We use this set to evaluate the model’s matching ability to retrieve the positive against a hard negative (Sec. 4.4). Finally we use the YCB-Video test set [59] of 12 videos to evaluate the tracking and updating components of our method on real world data (Sec. 4.5).

Across all datasets, train, validation, and test rooms are constant. More details on the Faster-RCNN detector and the exploration policy network can be found in Appx. A, B.

### 4.2. Downstream Rearrangement

The goal is to evaluate CSR for a downstream embodied task, 2-phase visual room rearrangement [52]. Rearrangement serves as a difficult, but practical, end-task with applications to indoor service robots. Furthermore, performing visual room rearrangement requires capable perception, memory, navigation, and planning, thereby testing how the various components of our approach work together.

Experiment	INFERRED				100 · SUCCESS ↑		100 · % FIXEDSTRICT ↑		%E ↓	
	plan	match	box	traj.	Val.	Test	Val.	Test	Val.	Test
CSR (GT-MBT)	✓	✗	✗	✗	8.8	10.0	26.0	27.0	0.74	0.73
CSR (GT-BT)	✓	✓	✗	✗	3.0	2.2	7.9	5.9	0.97	0.98
CSR (GT-T)	✓	✓	✓	✗	1.3	0.7	3.8	2.1	1.07	1.11
VRR [52]	-	-	-	-	0.5	0.2	1.2	0.7	1.15	1.12
VRR+Map [52]	-	-	-	-	0.6	0.3	1.6	1.4	1.15	<b>1.10</b>
CSR (Ours)	✓	✓	✓	✓	<b>1.2</b>	<b>0.4</b>	<b>3.3</b>	<b>1.9</b>	<b>1.13</b>	1.17

Table 1. **Downstream Rearrangement.** Using CSR, we are able to outperform the RL based visual room rearrangement baselines in the hard 2-phase setting (i.e., walkthrough and unshuffle phases happen sequentially). Additionally, the improved performance when including ground truth components, hints that as the various components of our approach improve, our end-to-end solution will also improve.

**Task.** In a walkthrough phase, an agent observes objects in their target configurations, building a CSR and embodied state graph:  $({}^{\text{walk}}\mathcal{G}_{\text{CSR}}, {}^{\text{walk}}\mathcal{G}_{\text{state}})$ . The agent leaves the room and 1-5 objects change (e.g., a mug moves from a countertop to a table). In the unshuffle phase, the agent returns and must restore the room to its original layout. In our case, a second unshuffle CSR and embodied state graph:  $({}^{\text{un}}\mathcal{G}_{\text{CSR}}, {}^{\text{un}}\mathcal{G}_{\text{state}})$  is built. Leveraging the fact the agent enters the room from the same starting location, we fuse the  ${}^{\text{walk}}\mathcal{G}_{\text{state}}$  and  ${}^{\text{un}}\mathcal{G}_{\text{state}}$  into a unified embodied state graph  $({}^{\text{walk,un}}\mathcal{G}_{\text{state}})$ , using the initial state as a point of correspondence between two trajectories. While this simplifying assumption is allowable due to the task construction of [52], it is plausible that state correspondence can be established using features, which is left to future work.

By identifying target objects that moved between creation of  ${}^{\text{walk}}\mathcal{G}_{\text{CSR}}$  and  ${}^{\text{un}}\mathcal{G}_{\text{CSR}}$  (i.e., high object feature similarity but low node feature similarity as discussed in Sec. 3.4), it is possible to query current and target locations for the moved objects in  $({}^{\text{walk,un}}\mathcal{G}_{\text{state}})$ . Using a simple breadth first search planner over  $({}^{\text{walk,un}}\mathcal{G}_{\text{state}})$ , the agent plans to the location of the target object, executes a pick-up action, and navigates to the goal location, where the object is placed.

**Model.** We use an exploration agent to explore the room during the two phases of the task. We use a simple exploration policy network with three convolutional layers and a GRU, which takes RGB images and outputs a distribution over navigation actions  $\mathcal{A}_{\text{nav}} = \{\text{MOVEFORWARD}, \text{MOVELEFT}, \text{MOVERIGHT}, \text{MOVEBACK}, \text{ROTATERIGHT}, \text{ROTATELEFT}, \text{LOOKUP}, \text{LOOKDOWN}\}$ . The network is trained using DD-PPO [44, 54] in the AllenAct framework [53], with a reward that encourages visiting more states and objects. Note: this objective has no notion of rearrangement. To generate the object region proposals used to generate box mask inputs for the CSR encoder, we use a FasterRCNN [38] detector. This detector is trained on LVIS [24] and AI2-THOR data. Leveraging the policy network and detector during walkthrough and unshuffle, we get CSR representations for each trajectory, with corresponding embodied state graphs. We then detect changes in the scene and

execute a plan to restore objects. For more details on the exploration training, reward, detector training, loss, and hyperparameters, see Appx. A, B, and C.

**Metrics.** We adopt rearrangement metrics [52]:

- **SUCCESS:** The percentage of episodes in which *all* the objects are correctly returned to their original states.
- **%FIXEDSTRICT:** The percentage of objects returned to their proper place over all rooms. However, if a mistake is made in a room (e.g., an object is moved that should not be), the value is 0 for that rearrangement even if other objects are returned correctly.
- **% E:** A ratio between the energy of a room before and after rearrangement. The energy function decreases monotonically as the agent brings objects closer to their target destinations. A value of 0 corresponds to a perfect rearrangement, 1 to no change, and greater than 1 to a final configuration further from the goal than the initialization.

**Ablation.** We ablate CSR in the context of the rearrangement task. We substitute ground truth (GT) for key components of our model to ablate the impact of their performances on the rearrangement task. We study the effects of using GT **Matching**, **Boxes**, and **Trajectories**. For GT matching, we substitute the matching function  $M_{\text{CSR}}$  for GT object instance labels, which allow for perfect matching, both within a trajectory and between walkthrough and unshuffle trajectories. For GT boxes, we substitute the FasterRCNN detector for GT boxes. For GT trajectories, we substitute the trained exploration policy for a heuristic policy that visits all objects that get shuffled in both walkthrough and unshuffle phases. Concretely, we compare against the following ablated models:

- **CSR (GT-MBT):** GT object **Matches**, **Boxes**, and heuristic **Trajectories**.
- **CSR (GT-BT):** Matching is estimated within a trajectory and between the two trajectories, keeping box and trajectory GT constant from CSR (GT-MBT).
- **CSR (GT-T):** Boxes are estimated using the Faster-RCNN detector, keeping the heuristic exploration trajectories constant from CSR (GT-BT).

**Baselines.** We compare against the following prior work for the challenging 2-phase rearrangement setting:

- *VRR*: Corresponds to the (RN18, PPO+IL) from [52]. The model uses a frozen ImageNet pretrained ResNet-18 and a mixture of imitation learning and reinforcement learning to train *directly* on the rearrangement task.
- *VRR+Map*: Similar to VRR, but with the addition of Active Neural SLAM [9] for mapping. Corresponds to (RN18+ANM,PPO+IL) from [52].

**Is CSR competitive v.s. sophisticated task-specific baselines?** The results in Tab. 1 support that using a simple exploration policy and detector in concert with CSR performs favorably on the SUCCESS and %FIXEDSTRICT metrics compared to the sophisticated baselines trained directly on the task. This shows the effectiveness of CSR since *no part of our approach was trained to rearrange rooms.*

**How do improved submodules improve the solution?** We replace each CSR submodule with GT as a proxy for an improved submodule. See Tab. 1. Looking at CSR (GT-T), we see that performance is comparable to CSR (Ours), suggesting that the trained exploration agent generalizes sufficiently. Comparing CSR (GT-BT), we see larger gains suggesting that with better detection, performance would improve by  $\sim 3$  times. Finally when comparing to CSR (GT-MBT), we find that there is room for improvement in matching. These results suggest a high ceiling to our pipeline.

### 4.3. Representation Linear Probes

We probe if CSR encodes interpretable relationships.

**Tasks.** We linearly probe CSR features for two relationship classification tasks. First, given a pair of box regions of interest, predict if one object supports the other. We call this the [SUPPORT] task. This is a three-way classification problem, where object 1 either supports object 2, vice versa, or there is no support relationship. The second task is predicting, given a pair of boxes, if they are supported by the same surface. We call this the [SIBLING] task. In this binary task, the direction of the edge does not matter.

**Model.** We fit a linear layer on a frozen CSR feature extractor  $f_{CSR}$ . We train two linear heads for [SUPPORT] and [SIBLING] respectively.

**Metrics.** We consider the mean accuracy across both tasks.

**Baselines.** We compare to the following baselines:

- *Random Features*: Randomly initialized and frozen ResNet-18 receiving 5 channels as input (3 for RGB and 2 for object boxes).
- *Supervised* [SUPPORT]: a ResNet-18 backbone trained end-to-end on the [SUPPORT] task. The representation is frozen and a separate linear head is trained for the [SIBLING] task.
- *Supervised* [SIBLING]: Similar to the *Supervised* [SUPPORT] model but trained end-to-end on the SIBLING task.

Experiment	Linear Probe mAcc	Retrieval Acc.
Random Chance	41.7	50.0
Random Features	51.8	52.4
Supervised [SUPPORT]	84.3	69.3
Supervised [SIBLING]	70.2	54.9
Fine-tuned [SUPPORT]	77.7	64.6
Fine-tuned [SIBLING]	78.3	57.6
CSR features	<b>85.3</b>	<b>84.3</b>

Table 2. **Linear Probes and Retrieval.** CSR features outperform supervised and fine-tuned baseline features.

- *Fine-tuned* [SUPPORT], *Fine-tuned* [SIBLING]: Same as *Supervised* SUPPORT, SIBLING; however, encoders are initialized with CSR contrastively pre-trained weights instead of random initialization.

### Benefits to continuous instead of discrete relations.

While the baselines outperforms CSR for the specific tasks they are trained for, when averaging performance over both tasks, CSR outperforms the baselines, as seen in Tab. 2 (left). These results support that training on specific relations can be detrimental in transfer setting and that our features can be used to recover specific spatial relations.

### 4.4. Retrieving Consistent Relationships

Our CSR learning objective assumes scenes are static, enforcing the similarity of relations between objects seen from different viewpoints. We hypothesize that our features should allow for detecting violations to this assumption.

**Task.** Consider triplets of images. The first is a query image containing a pair of objects. The second is the *positive* image taken from a different vantage point that also captures the same pair. The third is the *negative* image taken from the same pose as the second image; however, one of the objects in the pair is moved. The task then is to match the query and the positive image with CosSim of features as the matching criteria. Fig. 6 shows sample data.

**Model.** We use our  $f_{CSR}$  directly to embed three relationships. We compute the CosSim between the query features against the features from the positive and the negative images. To solve the classification task, we choose the image that has the higher CosSim score.

**Metrics.** We measure accuracy of the predictions.

**Baselines.** We adopt the same baselines as in Sec. 4.3. The Supervised [SUPPORT] baseline is the main point of comparison. This model determines if two objects participate in a support relationship (i.e., one on top of the other). Because changes in the dataset involve objects shifting to different supporting surfaces, these events should correspond to drastic changes in feature space.

**Results.** Based on the results in Tab. 2 (right), we find that CSR outperforms the baselines in terms of matching the query and positive images. See Fig. 6, which qualitatively illustrates the difficulty of the task. The results suggest that CSR learns some underlying spatial relationships.

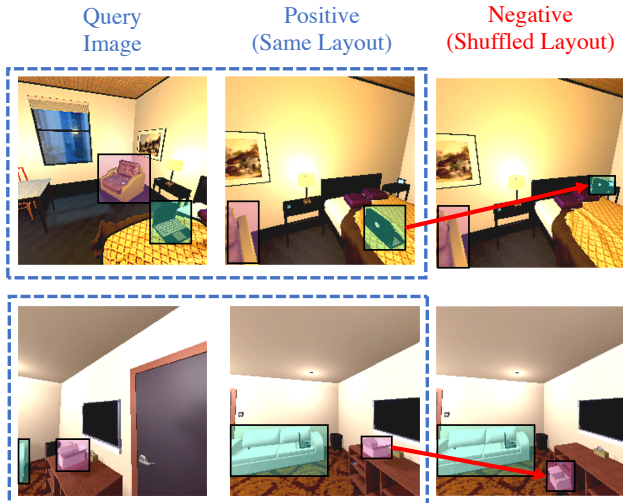


Figure 6. **Consistent relationship sample data.** Features for edges in the same room layout should be more similar than features in different layouts even with relatively small pixel differences in the positive and negative images.

#### 4.5. Real World Tracking

To demonstrate the potential of the representation on real-world data, we additionally evaluate CSR features, trained in simulation on object tracking in YCB-Video [59].

**Task.** The goal is to track objects over a trajectory. We formulate the problem as clustering, where, with each new observation, a detected object can either be assigned to an existing cluster (representing a previously observed object) or to a new cluster. The target GT clustering is one cluster for each object instance, with a count per cluster corresponding to the number of times the object was detected.

**Model.** CSR feature extraction and matching within a trajectory. Note, we use GT boxes for detections in lieu of a detector. Each node represents a cluster with number of elements equal to the number of assignments to that cluster.

**Metrics.** For each video, we measure the Adjusted Rand Index (ARI) [27], a common clustering accuracy metric, between the CSR and GT assignments. A value of 1 corresponds to perfect clustering and 0 to random clustering.

##### Ablations/Baselines:

- *No-Update*: Node features are fixed at initialization without the feature updating discussed in Sec. 3.2.
- *Opt.*: Oracle with feature updating and matching threshold set optimally.

**Results.** In Fig. 7 we see near perfect matching (compared to the baseline *Opt.*) and the usefulness of updating features (compared to *No-Update* baseline). These preliminary results suggest that CSR trained in simulation can be applied to real-world data distributions and that updating the representation on-the-fly is critical for tracking.

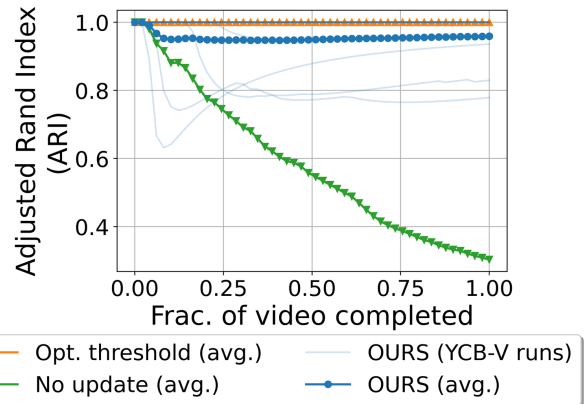


Figure 7. **Matching on YCB-Video.** Tracking performance is close to optimal and outperforms an ablated version without feature updating.

## 5. Limitations

In this work, we mainly focus on capturing the pose of objects and ignore the state of the objects (e.g., open v.s. closed microwave). The state of objects plays an important role in task planning. Future work will consider incorporating the state of the object. Also, the interaction of the agents with the objects is not used to update the scene representations. Hence, future work will address agent-object interactions. Additionally developing learned matching functions, which go beyond cosine similarity, is a promising direction to develop better matching performance.

## 6. Conclusion

We present Continuous Scene Representations (CSR) to model objects and their relations as continuous valued feature vectors. We describe an algorithm for updating the CSR as an agent moves and connect the representation to the agent itself. We discuss use of our pipeline to tackle visual room rearrangement. Our experiments show (1) A simple planner using CSR outperforms sophisticated baselines on visual room rearrangement, even though our model is not trained directly for this task. (2) The relational feature from CSR can be used to recover semantically meaningful spatial relations. (3) The representation can detect subtle visual changes in the scene. We hope this work will encourage investigation of techniques in scene representation learning for other downstream embodied tasks, going beyond learning task specific representations.

**Acknowledgements.** We would like to thank Winson Han for help with figures, Klemen Kotar for detector weights, and Luca Weihs for help with AllenAct. SS is supported by NSF CMMI-2037101 and NSF IIS-2132519. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the sponsors.



## References

- [1] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton Van Den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *CVPR*, 2018. 2
- [2] Iro Armeni, Zhi-Yang He, JunYoung Gwak, Amir R Zamir, Martin Fischer, Jitendra Malik, and Silvio Savarese. 3d scene graph: A structure for unified semantics, 3d space, and camera. In *ICCV*, 2019. 1, 2
- [3] Sid Yingze Bao, Mohit Bagra, Yu-Wei Chao, and Silvio Savarese. Semantic structure from motion with points, regions, and objects. In *CVPR*, 2012. 2
- [4] Dhruv Batra, Angel X Chang, Sonia Chernova, Andrew J Davison, Jia Deng, Vladlen Koltun, Sergey Levine, Jitendra Malik, Igor Mordatch, Roozbeh Mottaghi, et al. Rearrangement: A challenge for embodied ai. *arXiv*, 2020. 2, 5
- [5] Dhruv Batra, Aaron Gokaslan, Aniruddha Kembhavi, Oleksandr Maksymets, Roozbeh Mottaghi, Manolis Savva, Alexander Toshev, and Erik Wijmans. Objectnav revisited: On evaluation of embodied agents navigating to objects. *arXiv*, 2020. 2
- [6] Peter W. Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, and Koray Kavukcuoglu. Interaction networks for learning about objects, relations and physics. In *NeurIPS*, 2016. 2
- [7] Sean L Bowman, Nikolay Atanasov, Kostas Daniilidis, and George J Pappas. Probabilistic data association for semantic slam. In *ICRA*, 2017. 2
- [8] Devendra Singh Chaplot, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural slam. In *ICLR*, 2020. 2
- [9] Devendra Singh Chaplot, Ruslan Salakhutdinov, Abhinav Gupta, and Saurabh Gupta. Neural topological slam for visual navigation. In *CVPR*, 2020. 2, 7
- [10] Kevin Chen, Juan Pablo de Vicente, Gabriel Sepulveda, Fei Xia, Alvaro Soto, Marynel Vazquez, and Silvio Savarese. A behavioral approach to visual navigation with graph localization networks. In *RSS*, 2019. 2
- [11] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv*, 2020. 4, 12
- [12] Yixin Chen, Siyuan Huang, Tao Yuan, Siyuan Qi, Yixin Zhu, and Song-Chun Zhu. Holistic++ scene understanding: Single-view 3d holistic scene parsing and human pose estimation with human-object interaction and physical common-sense. In *ICCV*, 2019. 2
- [13] Yuren Cong, Wentong Liao, Hanno Ackermann, Bodo Rosenhahn, and Michael Ying Yang. Spatial-temporal transformer for dynamic scene graph generation. In *ICCV*, 2021. 2
- [14] Bo Dai, Yuqi Zhang, and Dahua Lin. Detecting visual relationships with deep relational networks. In *CVPR*, 2017. 2
- [15] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *CVPR*, 2018. 2
- [16] Andrew J Davison, Ian D Reid, Nicholas D Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *TPAMI*, 2007. 2
- [17] Heming Du, Xin Yu, and Liang Zheng. Learning object relation graph and tentative policy for visual navigation. In *ECCV*, 2020. 2
- [18] Kiana Ehsani, Winson Han, Alvaro Herrasti, Eli VanderBilt, Luca Weihs, Eric Kolve, Aniruddha Kembhavi, and Roozbeh Mottaghi. Manipulathor: A framework for visual object manipulation. In *CVPR*, 2021. 2
- [19] Jakob Engel, Thomas Schöps, and Daniel Cremers. Lsdslam: Large-scale direct monocular slam. In *ECCV*, 2014. 2
- [20] Matthew Fisher, Manolis Savva, and Pat Hanrahan. Characterizing structural relationships in scenes using graph kernels. In *ACM SIGGRAPH*, 2011. 2
- [21] Samir Yitzhak Gadre, Kiana Ehsani, and Shuran Song. Act the part: Learning interaction strategies for articulated object part discovery. *ICCV*, 2021. 2
- [22] Chuang Gan, Jeremy Schwartz, Seth Alter, Damian Mrowca, Martin Schrimpf, James Traer, Julian De Freitas, Jonas Kubilius, Abhishek Bhandwalder, Nick Haber, Megumi Sano, Kuno Kim, Elias Wang, Michael Lingelbach, Aidan Curtis, Kevin Tyler Feiglis, Daniel Bear, Dan Gutfreund, David Daniel Cox, Antonio Torralba, James J. DiCarlo, Joshua B. Tenenbaum, Josh McDermott, and Daniel LK Yamins. ThreeDWorld: A platform for interactive multi-modal physical simulation. In *NeurIPS*, 2021. 2
- [23] Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. Iqa: Visual question answering in interactive environments. In *CVPR*, 2018. 2
- [24] Agrim Gupta, Piotr Dollár, and Ross B. Girshick. Lvis: A dataset for large vocabulary instance segmentation. *CVPR*, 2019. 6, 12
- [25] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *CVPR*, 2017. 2
- [26] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *CVPR*, 2020. 4, 12
- [27] L. Hubert and P. Arabie. Comparing partitions. *Journal of Classification*, 1985. 8
- [28] Jingwei Ji, Ranjay Krishna, Li Fei-Fei, and Juan Carlos Niebles. Action genome: Actions as compositions of spatio-temporal scene graphs. In *CVPR*, 2020. 1, 2
- [29] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David A Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *CVPR*, 2015. 1
- [30] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. AI2-THOR: An Interactive 3D Environment for Visual AI. *arXiv*, 2017. 2, 5
- [31] Xinghang Li, Di Guo, Huaping Liu, and Fuchun Sun. Embodied semantic scene graph generation. In *CoRL*, 2021. 2
- [32] Yikang Li, Wanli Ouyang, Bolei Zhou, Kun Wang, and Xiaogang Wang. Scene graph generation from objects, phrases and region captions. In *ICCV*, 2017. 1, 2

- [33] Chenchen Liu, Yang Jin, Kehan Xu, Guoqiang Gong, and Yadong Mu. Beyond short-term snippet: Video relation detection with spatio-temporal global context. In *CVPR*, 2020. 2
- [34] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *ECCV*, 2016. 1, 2
- [35] Raul Mur-Artal and Juan D Tardós. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE Trans. on Robotics*, 2017. 2
- [36] Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural scene graphs for dynamic scenes. In *CVPR*, 2021. 2
- [37] Claudia Pérez-D’Arpino, Can Liu, Patrick Goebel, Roberto Martín-Martín, and Silvio Savarese. Robot navigation in constrained pedestrian environments using reinforcement learning. In *ICRA*, 2021. 2
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 6
- [39] Antoni Rosinol, Arjun Gupta, Marcus Abate, Jingnan Shi, and Luca Carlone. 3d dynamic scene graphs: Actionable spatial perception with places, objects, and humans. In *RSS*, 2020. 2
- [40] 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*, 2021. 2
- [41] Martin Rünz and Lourdes Agapito. Co-fusion: Real-time segmentation, tracking and fusion of multiple objects. In *ICRA*, 2017. 2
- [42] Renato F Salas-Moreno, Richard A Newcombe, Hauke Strasdat, Paul HJ Kelly, and Andrew J Davison. Slam++: Simultaneous localisation and mapping at the level of objects. In *CVPR*, 2013. 2
- [43] Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. In *ICLR*, 2018. 2
- [44] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv*, 2017. 6, 12
- [45] Xindi Shang, Tongwei Ren, Jingfan Guo, Hanwang Zhang, and Tat-Seng Chua. Video visual relation detection. In *ACM Multimedia*, 2017. 2
- [46] Bokui Shen, Fei Xia, Chengshu Li, Roberto Martín-Martín, Linxi Fan, Guanzhi Wang, Shyamal Buch, Claudia D’Arpino, Sanjana Srivastava, Lyne P Tchapmi, et al. igibson, a simulation environment for interactive tasks in large realistic scenes. In *IROS*, 2021. 2
- [47] Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. Alfred: A benchmark for interpreting grounded instructions for everyday tasks. In *CVPR*, 2020. 2
- [48] Michael Strecke and Jorg Stuckler. Em-fusion: Dynamic object-level slam with probabilistic data association. In *ICCV*, 2019. 2
- [49] Yao-Hung Hubert Tsai, Santosh Divvala, Louis-Philippe Morency, Ruslan Salakhutdinov, and Ali Farhadi. Video relationship reasoning using gated spatio-temporal energy graph. In *CVPR*, 2019. 2
- [50] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv*, 2018. 4, 12
- [51] Johanna Wald, Helisa Dharmo, Nassir Navab, and Federico Tombari. Learning 3d semantic scene graphs from 3d indoor reconstructions. In *CVPR*, 2020. 2
- [52] Luca Weihs, Matt Deitke, Aniruddha Kembhavi, and Roozbeh Mottaghi. Visual room rearrangement. In *CVPR*, 2021. 2, 5, 6, 7, 13
- [53] Luca Weihs, Jordi Salvador, Klemen Kotar, Unnat Jain, Kuo-Hao Zeng, Roozbeh Mottaghi, and Aniruddha Kembhavi. Allenact: A framework for embodied ai research. *arXiv*, 2020. 6, 12
- [54] Erik Wijmans, Abhishek Kadian, Ari S. Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, and Dhruv Batra. Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames. In *ICLR*, 2020. 2, 6, 12
- [55] Yu-Shiang Wong, Changjian Li, Matthias Niessner, and Niloy J. Mitra. Rigidfusion: Rgb-d scene reconstruction with rigidly-moving objects. *Computer Graphics Forum*, 2021. 2
- [56] Mitchell Wortsman, Kiana Ehsani, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Learning to learn how to learn: Self-adaptive visual navigation using meta-learning. In *CVPR*, 2019. 2
- [57] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 12
- [58] Yi Wu, Yuxin Wu, Aviv Tamar, Stuart Russell, Georgia Gkioxari, and Yuandong Tian. Bayesian relational memory for semantic visual navigation. In *ICCV*, 2019. 2
- [59] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *RSS*, 2018. 2, 5, 8
- [60] Binbin Xu, Wenbin Li, Dimos Tzoumanikas, Michael Bloesch, Andrew Davison, and Stefan Leutenegger. Mid-fusion: Octree-based object-level multi-instance dynamic slam. In *ICRA*, 2019. 2
- [61] Danfei Xu, Yuke Zhu, Christopher B Choy, and Li Fei-Fei. Scene graph generation by iterative message passing. In *CVPR*, 2017. 1, 2
- [62] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph r-cnn for scene graph generation. In *ECCV*, 2018. 1, 2
- [63] Wei Yang, Xiaolong Wang, Ali Farhadi, Abhinav Gupta, and Roozbeh Mottaghi. Visual semantic navigation using scene priors. In *ICLR*, 2019. 2
- [64] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *CVPR*, 2018. 1, 2
- [65] Yibiao Zhao and Song-Chun Zhu. Scene parsing by integrating function, geometry and appearance models. In *CVPR*, 2013. 1

- [66] Yang Zhou, Zachary While, and Evangelos Kalogerakis. Scenegraphnet: Neural message passing for 3d indoor scene augmentation. In *ICCV*, 2019. [2](#)
- [67] Yifeng Zhu, Jonathan Tremblay, Stan Birchfield, and Yuke Zhu. Hierarchical planning for long-horizon manipulation with geometric and symbolic scene graphs. In *ICRA*, 2021. [2](#)