

Planning for Multi-Object Manipulation with Graph Neural Network Relational Classifiers

Yixuan Huang¹, Adam Conkey¹, and Tucker Hermans^{1,2}

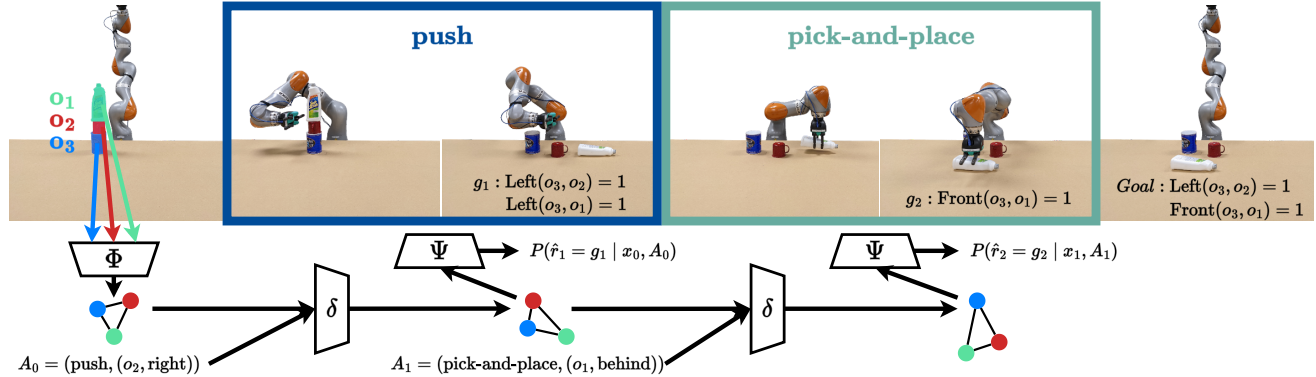


Fig. 1: Planning using RD-GNN: The robot encodes the segmented scene point cloud into a learned latent state space with graph structure using the encoder Φ . The robot then plans a sequence of actions using a learned latent dynamics function δ to predict relations \hat{r} that match both subgoal \mathbf{g}_1 and subgoal \mathbf{g}_2 . The robot first plans to push o_2 to the right achieving the first subgoal \mathbf{g}_1 then pick-and-place object o_1 to be the behind of o_3 achieving the second subgoal \mathbf{g}_2 and thus the goal $\mathbf{g} = \mathbf{g}_1 \wedge \mathbf{g}_2$. Note relations are defined from the camera's perspective.

Abstract—Objects rarely sit in isolation in human environments. As such, we'd like our robots to reason about how multiple objects relate to one another and how those relations may change as the robot interacts with the world. To this end, we propose a novel graph neural network framework for multi-object manipulation to predict how inter-object relations change given robot actions. Our model operates on partial-view point clouds and can reason about multiple objects dynamically interacting during the manipulation. By learning a dynamics model in a learned latent graph embedding space, our model enables multi-step planning to reach target goal relations. We show our model trained purely in simulation transfers well to the real world. Our planner enables the robot to rearrange a variable number of objects with a range of shapes and sizes using both push and pick-and-place skills.

I. INTRODUCTION

Robots operating in human environments must contend with many objects at a time. As such robotic multi-object manipulation and rearrangement has received much attention in the literature [1–9]. The most recent of these works show excellent results in reasoning about novel objects from partial view sensory information [6–12]. However, robots using these approaches operate in a limited capacity manipulating individual objects one at a time. In contrast, a pair of recent works have shown the ability for robots to reason about and manipulate multiple objects at once [13,14]. These works leverage image-based feedback controllers, that lack the level of semantic reasoning and explicit object grounding we desire for multi-step planning to logical goals. How can a robot plan to logical goals while still reasoning about a variable number of dynamically interacting objects?

Our work seeks to answer this question. We propose an approach that can explicitly predict inter-object relations

in multi-object scenes, while also being able to plan to manipulate multiple objects at once. We advocate for the use of logical relations for specifying goals as in [7,9] as they provide a useful language for communication between robot and human. A human can easily construct a goal for tasking a robot by providing a conjunction of desired logical relations between objects in the scene. On the flip side, the robot can use its predictions of logical relations to communicate its belief over the current scene or future states it intends to achieve through manipulation. This contrasts with several recent approaches to rearrangement which provide images as goals to the robot [15]. Generating images for all desired goals requires a much higher burden on the user and in many cases would require the user to actively rearrange the scene, obviating the need for the robot!

To enable reasoning about manipulation effects simultaneously on multiple objects, we propose learning a latent state space in the form of a graph neural network (GNN). Our proposed relational dynamics GNN (RD-GNN) takes as input a segmented, partial view point cloud of the objects in the scene. It encodes this observation into the graph latent space, from which it can predict inter-object relations for both the current scene and future states given a sequence of actions. To enable these future predictions, we learn a dynamics function in the latent graph space. We then use our learned network to perform planning to achieve a desired relational goal. Figure 1 illustrates how our learned model can be used for multi-step planning. Importantly, our planning framework can incorporate multiple distinct robot skills and produce multi-step task plans to achieve the specified goals. Our planner performs diverse multi-object rearrangements including lifting and placing multiple objects at once, building and deconstructing towers, pushing objects to be in contact, and aligning objects spatially.

¹University of Utah. ²NVIDIA. yixuan.huang@utah.edu, adam.conkey@utah.edu, and thermans@cs.utah.edu

We leverage GNNs as their relational inductive bias enables better reasoning about multiple object interactions compared to alternative neural network structures [13]. Our experiments provide further evidence to this effect. We show that RD-GNN outperforms a similarly structured multi-layer perceptron operating directly on object pairs both in terms of planning success rate and predicting post-manipulation relations to enable successful planning. Crucially, using graph neural networks allows the robot to use the same model to reason about a variable number of objects. Further, by directly using partial view point cloud information as input, the robot can reason about objects of novel shape and size without access to explicit object models.

We further test the hypothesis that we can train RD-GNN using only the pre- and post-manipulation relational labels for supervision, in addition to the input point cloud and actions. Our extensive simulated and real world experiments show that using this relational supervision outperforms training to predict changes in object pose, coupled with an analytic approach to predicting relations from the object bounding boxes. Further, we show that training with both the relational and pose estimation losses provides no real benefit over training with relational losses alone.

II. RELATED WORK

Neural networks, including graph neural networks, have been applied to reason about spatial relationships and perform planning based on said reasoning [7–9,16,17]. Paxton et al. [7] present a framework to reason about pairwise relations and plan to find an object placement that is physically feasible and satisfies the goal relations. Liu et al. [8] present a transformer-based framework to manipulate objects into a configuration that satisfies the multi-object semantic goal relations encoded from natural language. This approach reasons about multi-object relations, but like [7] only moves a single object at a time via pick-and-place. It also requires added complexity incorporating language and loses the ability to operate explicitly on logical goals. Yuan et al. [17] present a framework to learn object feature embeddings incorporating single object semantics from RGB images for use in sequential manipulation tasks. Simeonov et al. [16] leverage a GNN as a graph-attention network to select contact points and object transformations from point cloud observations for single-object manipulation. Zhu et al. [9] presents a grounded hierarchical planning framework for long-horizon planning manipulation tasks that leverages a symbolic scene graph to predict high-level plan actions and a geometric scene graph to predict low-level motions. Unlike our work, Zhu et al. [9] do not examine multi-object dynamic interactions. Lou et al. [18] predict spatial relations between objects in clutter using GNNs to aid in finding better grasps, but do not model how relations will change post grasp. Furthermore, Driess et al. [19] learn to predict multi-object interactions using graph nets, with supervised reconstruction for NERF-like embeddings. Unlike our proposed approach they do not predict object relations and learn and plan at a much finer time scale which makes their simulation-only

experiments unlikely to transfer well to the real world. Biza et al. [20] similarly examine learning object-oriented models of the world with pose estimation supervision. They show the ability to embed pose-based goals into a latent space, but do not explicitly reason about relations or manipulating multiple objects at once. In [11] a GNN-based policy learns to perform multi-object rearrangement tasks including stacking and unstacking. However, the policy requires full object pose information and manipulates one object at a time.

Object stacking and unstacking tasks are challenging for robots to perform autonomously [21–25] due to the difficulty in modeling the non-trivial contact dynamics and support relations of the objects being stacked. Some recent works leverage GNNs for object stacking and unstacking tasks [10–12]. Li et al. [10] leverage the GNNs to build relational reinforcement learning framework to help capture multi-object information in object stacking and unstacking. However, this method does not consider generalization to objects of different shape and size, requires expert demonstrations, does not show real-world experiments, and uses only one primitive action. Sharma and Kroemer [12] leverage GNNs to predict the feasibility of an action in object stacking and unstacking. They only consider preconditions, while we focus on leveraging multi-object dynamics to achieve logical goal relations. Furthermore, they require full 3D scene observations from multiple cameras, while we use the partial point cloud from a single camera.

Long-horizon planning has become an important problem for robot manipulation. Task and motion planning (TAMP) [26–32] defines a promising method to solve long horizon problems. TAMP approaches typically assume models of how objects and potentially their relations change. While learning has been used for various aspects of TAMP, no work has shown how to plan with multi-object dynamic interactions from point cloud data. Simeonov et al. [16] propose an approach to object manipulations from point cloud data. They leverage a plan skeleton similar to us to solve long horizon planning problems. However, they do not reason about object relations and only manipulate one object with each action. Liang et al. [32] learn to plan with different skill primitives which sometimes include multi-object dynamic interactions. They perform multi-step skill planning using a heuristic graph search. However, they assume knowledge of object state and do not explicitly reason about object relations for learning.

III. PLANNING TO GOAL RELATIONS WITH GNNs

We assume our robot perceives the world as a point cloud Z with N associated object segments $O_i \subset Z, i = 1, 2, \dots, N$. The robot receives a goal, $\mathbf{g} = r_1 \wedge r_2 \wedge \dots \wedge r_M; r_j \in \mathcal{R}$, defined as M desired object relations. \mathbf{g} represents the goal relation conjunction, r_j represents each goal relation, and \mathcal{R} represents the set of all possible relations. Example relations in \mathcal{R} include planar spatial relations such as “object i is in front of object j ” or 3D relations such as “object i is above object j ” and “object i is in contact with object j .” We assume the robot receives a plan skeleton

$G = (\mathbf{g}_1, \dots, \mathbf{g}_H)$ [16,33] specifying the subgoals for each step in the multi-step plan of length H . We do not find this overly restrictive as several different approaches can generate appropriate plan skeletons [27–29,32].

We provide our robot with a set of L parametric action primitives $\mathcal{A} = \{A_1, \dots, A_L\}$ where A_l defines the discrete skill, which has associated skill parameters θ_l . Example skills include a push skill (A_l) with parameters (θ_l) encoding the end effector pose and push length or a pick-and-place skill defined by the grasp and placement poses.

We define the robot’s planning task as finding a sequence of skills and skill parameters $\tau = (A_0, \dots, A_{H-1})$ that transforms the objects such that they sequentially satisfy each subgoal \mathbf{g}_i in the skeleton G . We propose learning a latent-space dynamics model [34] for use in a planner to solve this task. The latent space model enables the robot to encode its partial view point cloud observations to a representation useful for planning. Crucially we propose a novel graph neural network architecture to encode the latent space with a graph structure. This enables the robot to explicitly reason about a variable number of objects in the scene.

We learn an encoder to map observed segmented point clouds to the latent space $\Phi : \mathcal{Z} \rightarrow \mathcal{X}$ and an associated decoder to predict inter-object relations from latent states $\Psi : \mathcal{X} \rightarrow \mathcal{R}$. To predict the forward state dynamics for planning we learn a forward dynamics function $\delta : \mathcal{X} \times \mathcal{A} \rightarrow \mathcal{X}$.

We can now formally define our planning objective as maximizing the probability of achieving the goal relations with the following constrained optimization problem:

$$\arg \max_{\tau=(A_0, \dots, A_{H-1})} \prod_{k=1}^H P(\mathbf{r}_k = \mathbf{g}_k | \mathbf{x}_k) = \Psi(\mathbf{x}_k) \quad (1)$$

$$\text{subject to } \mathbf{x}_{k+1} = \delta(\mathbf{x}_k, A_k) \quad \forall k = 0, \dots, H-1 \quad (2)$$

$$\mathbf{x}_0 = \Phi(Z_0) \quad (3)$$

$$A_k \in \mathcal{A} \quad \forall k = 0, \dots, H-1 \quad (4)$$

$$\theta_{min} \preceq \theta_k \preceq \theta_{max} \quad \forall k = 0, \dots, H-1 \quad (5)$$

The constraints in this optimization problem encode the latent space dynamics (Eq. (2)), grounding of the initial latent state from the observed point cloud Eq. (3), and constraints on the action parameters Eqs. (4–5). We thus chain together predicted action effects decoding each state to predict the inter-object relations. Figure 1 visualizes planning with this model.

After solving this optimization problem the robot can execute the planned actions in the physical world. Our proposed network enables the robot to validate if it achieved its goal by computing $\mathbf{r}_k = \Psi(\Phi(Z_k))$. Where Z_k denotes the current point cloud observation.

We now provide a brief introduction to graph neural networks which are crucial building blocks for our proposed model. We then discuss the details of our specific relational dynamics graph neural network architecture visualized in Fig. 2. We follow this with a discussion of training our model before concluding this section with a description of our planning algorithm.

A. Graph Neural Networks

We define a directed graph $G = \{V, E\}$ with nodes $V = \{\mathbf{v}_i\}$ and edges as $E = \{\mathbf{e}_{ij}^{\rightarrow}\}$ where each \mathbf{v}_i and $\mathbf{e}_{ij}^{\rightarrow}$ is a feature vector for node i or the edge from i to j respectively. We seek to encode information associated with this graph into a neural network; following [35] we can reason about our graph network operations in terms of message passing in the graph, where a single graph net layer of *update* and *aggregation* functions performs one round of message passing between neighbors in the graph. By constructing multiple graph layers, information from nodes across the graph can propagate in the form of deeper and deeper features.

Update functions transform individual node or edge features. We use feed-forward multi-layer perceptrons as update functions in this paper. We denote node updates as $\mathbf{v}'_i = f_n(\mathbf{v}_i)$ and edge updates as $\mathbf{e}'_{ij} = f_e(\mathbf{e}_{ij}^{\rightarrow})$. Aggregations take inputs from multiple parts of the graph and reduce them to a fixed feature length, thus enabling consistent output feature dimensions from a variable input size. We denote a message from node i to node j as $\mathbf{m}_{ij}^{\rightarrow} = (\mathbf{v}_i \oplus \mathbf{v}_j \oplus \mathbf{e}_{ij}^{\rightarrow})$ and define our message update functions as $\mathbf{m}'_{ij} = f_m(\mathbf{m}_{ij}^{\rightarrow})$. Here \oplus denotes vector concatenation. To define our aggregation functions we introduce an intermediate variable $\mathbf{y}_i = \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} (\mathbf{m}'_{ji})$ which takes the average of all messages incoming to node \mathbf{v}_i denoted as those coming from nodes in node i ’s neighborhood $\mathcal{N}(i)$. Using this we can define our node aggregation function as $\mathbf{v}'_i = g_v(\mathbf{v}_i \oplus \mathbf{y}_i)$ and the edge aggregations as $\mathbf{e}'_{ij} = g_e(\mathbf{v}_i \oplus \mathbf{v}_j \oplus \mathbf{y}_i \oplus \mathbf{y}_j \oplus \mathbf{e}_{ij}^{\rightarrow})$ where $g_v(\cdot)$ and $g_e(\cdot)$ define MLPs. This edge aggregation thus concatenates and then transforms the features associated with the two neighboring nodes and the messages passing between them. For more details on graph nets including alternative aggregation functions see [35].

B. Learning Relational Dynamics with GNNs

We now turn our attention to our relational dynamics graph neural network, RD-GNN, which takes as input the segmented object point cloud and a candidate action and predicts the current and post-manipulation inter-object relations. We discuss the network in terms of several different components: the encoder, the latent graph dynamics function, and the relational output classifier. An overview of these different components and how they connect to one another is visualized in Figure 2.

Our encoder Φ can be decomposed into two sub-networks: the point cloud encoder and the latent graph encoder. Our point cloud encoder $\phi_{PC}(Z)$ operates on each of the N point cloud segments, $O_i \subset Z$, converting the variable size input point cloud to a fixed-length feature vector \mathbf{v}_i . This feature vector will act as the node feature to our input graph to the GNN. We use *PointConv* [36] as the backbone of our point cloud encoder to output a feature of length 128.

Given the output of our point cloud encoder, we define our input graph $\mathbf{x}^I = (V^I, E^I)$ with nodes $V^I = \{\phi_{PC}(O_i) \oplus k\}$ where k denotes a one-hot encoding providing a unique

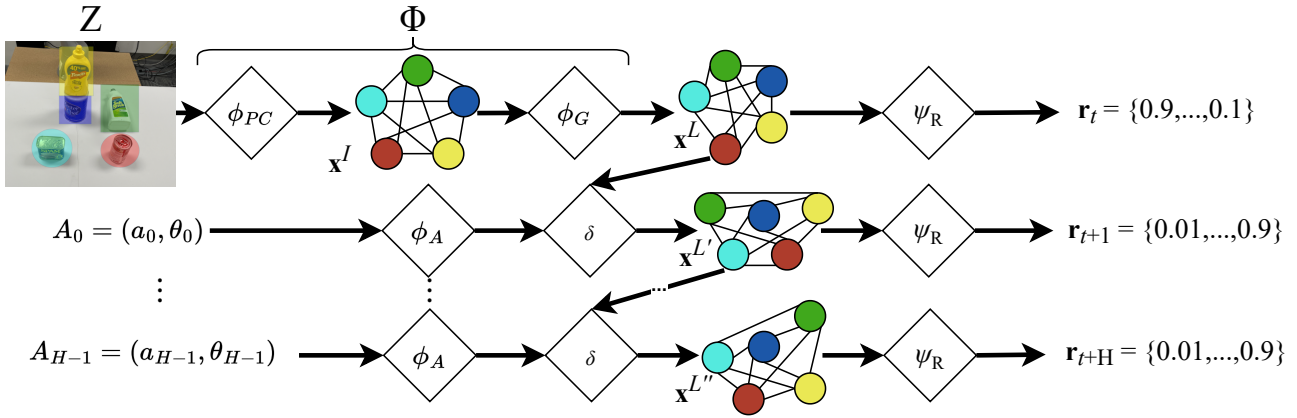


Fig. 2: Overview of the components of RD-GNN. A point cloud encoder ϕ_{PC} transforms a segmented point cloud into N feature vectors acting as the node features for the fully connected input graph \mathbf{x}^I . The graph encoder ϕ_G transforms \mathbf{x}^I into a latent graph state embedding \mathbf{x}^L . This latent graph in turn acts as input to a relational classifier ψ_R and dynamics function δ along with actions encoded via ϕ_A .

identity label for each node. To improve the generalization ability we randomly generate the object IDs during training [37] over a range larger than the highest number of objects expected to be seen at deployment. We chose 16 in the paper since it was all we needed for our experiments, although it could be set higher for other applications.

We define edges to and from all node pairs in the graph creating a fully-connected, directed input graph. We set all input edge feature $\mathbf{e}_{ij} \in E^I$ to be empty. This topology enables message passing between all nodes, but provides no explicit edge features as input for learning.

We use our graph encoder to transform our input graph, \mathbf{x}^I to a latent graph $\mathbf{x}^L = \phi_G(\mathbf{x}^I)$. Here $\phi_G(\cdot)$ represents a layer of graph message passing and aggregation as defined in the previous section. We use our latent graph embedding as input to two sub-networks: our relational classifier, $\mathbf{r} = \psi_R(\mathbf{x}^L)$ and our latent graph dynamics function $\mathbf{x}^{L'} = \delta(\mathbf{x}^L, A)$.

We construct our relational classifier as an MLP that operates on a pair of nodes and their associated edges from \mathbf{x}^L , taking the form of an edge aggregation network $\mathbf{r}_{ij} = \psi_R(\mathbf{v}_i, \mathbf{v}_j, \mathbf{y}_i, \mathbf{y}_j, \mathbf{e}_{ij})$. We predict relations for all object pairs by running this classifier for each pair of nodes in the graph as a form of graph convolution. While some relations may be mutually exclusive, in general the spatial relations are independent of one another, necessitating individual binary classifiers and not a softmax-based multi-class classifier. Note we never specify mutually exclusive goal relations.

We additionally examine learning to predict the object pose (defined as its centroid and bounding box orientation in simulation) for all objects in the scene. To this end we learn a pose regressor $\mathbf{x}_i = \psi_P(\mathbf{x}^L)$ which we train using a node aggregation network with an output MLP with 3 outputs encoding position and 6 encoding orientation as in [38].

The final piece to define is our latent graph dynamics function $\mathbf{x}^{L'} = \delta(\mathbf{x}^L, A)$. Recall that A defines the action (skill) including its skill parameters being evaluated through the dynamics. We encode any discrete skill variables (e.g. object identity) using a one-hot-encoding for use as input into the network. We pass this action through an action encoder $A' = \phi_A(A)$ which we implement as an MLP. We build separate node $\delta_v(\cdot)$ and edge $\delta_e(\cdot)$ dynamics

functions which respectively take as input the node or edge features of the latent graph concatenated with the encoded action. As output they predict the change in graph features $\Delta \mathbf{v}_i^{L'}, \Delta \mathbf{e}_{ij}^{L'}$. Given these definitions we define our graph dynamics functions as $\mathbf{v}_i^{L'} = \mathbf{v}_i^L + \delta_v(\mathbf{v}_i^L \oplus \phi_A(A))$ and $\mathbf{e}_{ij}^{L'} = \mathbf{e}_{ij}^L + \delta_e(\mathbf{e}_{ij}^L \oplus \phi_A(A))$. We incorporate multiple skills by learning a separate dynamics functions for each skill, using the same shared latent space.

C. Multi-step Loss Functions and Model Training

We train our model end-to-end using a combination of loss function terms. The first term defines the classification (cross-entropy) loss of predicting the current relations from the currently observed point cloud $\mathcal{L}_{REL} = \sum_{t=1}^H \text{CE}(\mathbf{r}_t, \hat{\mathbf{r}}_t = \psi_R(\Phi(Z_t)))$. Our second loss component provides regularization that the predicted, post-manipulation latent graph should match the latent graph encoded from the point cloud observation after executing the associated action [34] $\mathcal{L}_{DYN} = \sum_{t=1}^H \sum_{i=0}^{H-t-1} \|\mathbf{x}_{t+i}^L - \mathbf{x}_{t+i}^{L'}\|_2^2$. Where $\mathbf{x}_{t+i}^L = \Phi(Z_{t+i})$ defines the latent graph encoded directly from an observed point cloud and $\mathbf{x}_{t+i}^{L'} = \delta(\Phi(Z_{t+i}), A_{t+i})$ defines the predicted latent graph state through the learned dynamics function. We recursively use our δ functions to predict the latent state for all time steps in the executed plan.

Our final loss term operates on the post-manipulation relations predicted via the latent graph dynamics, $\mathcal{L}_{REL'} = \sum_{t=1}^H \sum_{i=0}^{H-t-1} \text{CE}(\mathbf{r}_{t+i}, \hat{\mathbf{r}}'_{t+i})$. We make the distinction between $\hat{\mathbf{r}}_{t+i}$ and $\hat{\mathbf{r}}'_{t+i}$ explicit, where $\hat{\mathbf{r}}_{t+i} = \psi_R(\mathbf{x}_{t+i}^L)$ and $\hat{\mathbf{r}}'_{t+i} = \psi_R(\mathbf{x}_{t+i}^{L'})$. This operates in an analogous multi-step fashion to the latent dynamics regularization.

For pose estimation we define an L2 loss on the current and predicted object poses in an analogous manner replacing ψ_R with ψ_P . We examine the effect of this loss in our experiments.

D. Planning Algorithm Implementation

Any number of methods can solve the optimization problem defined in Eq. (1). For this paper we use the cross entropy method (CEM) [39,40]. CEM is a derivative-free,




Fig. 3: For the same initial scene (left) we show different valid states found by our planner and model for two different goal settings. For the first goal relation, the robot can either pick the green object or the red object to place atop the yellow object. For the second goal relation, the robot can either push the green object or pick-and-place the green object to deconstruct the towers. Here we label the objects using color instead of object id to make the figure easier to interpret.

sampling-based procedure that begins with an initial Gaussian distribution over the decision variables $\mathcal{N}(\tau \mid \mu_0, \Sigma_0)$. We generate a set of samples $S = \{\tau_i \sim \mathcal{N}(\tau \mid \mu_0, \Sigma_0)\}$, evaluate each one under the specified cost function, and select the top- k low-cost samples $S_{\text{top}k}$. We recompute the mean and covariance for the current iteration as $\mu_1 = \text{mean}(S_{\text{top}k})$ and $\Sigma_1 = \text{cov}(S_{\text{top}k})$, and then proceed to the next iteration. We repeat this process for K iterations and select the final mean μ_M as the result. Our optimization problem requires a mixed discrete-continuous optimization due to having to find a sequence of skills (discrete) and the parameters to those skills that are themselves mixed discrete-continuous. Because our robot has a small number of skills, we independently search for each discrete skill and select the one with the lowest cost (highest success probability). To improve the numerical stability we minimize the log of Eq. (1) instead of Eq. (1). Given the subgoal skeleton we can greedily solve the continuous and discrete action search for each step of the plan. We propagate the predicted latent state, \mathbf{x}_{t+1} resulting from running action A_t as initial state for the next plan step.

IV. EXPERIMENTS & RESULTS

We now describe the training data collection and experimental validation for our approach to learning and planning with RD-GNN. In our experiments we examine the following relations: left, right, behind, in-front, above, below and in-contact. We get the in-contact relation directly from simulation and we define other relations following Paxton et al. [7]. We train and evaluate our model on multi-object rearrangement tasks using pushing and pick-and-place skills. We conduct experiments in simulation and on a physical robot manipulating both blocks and YCB objects [41].

Dataset collection in simulation: We conduct large scale data collection using the Isaac Gym simulator [42]. We collect a dataset by generating scenes with a variable number of cuboid objects of random size with arbitrary pose. Scenes contain objects in either one or two vertical stacks. We then execute a random push or pick-and-place action on one of the objects in the scene. We record the partial view point cloud before and after the manipulation, the executed action, and the ground truth relations between all object pairs in the scene. We collected a total of 39,600 push and pick-and-place attempts. Fig. 3 shows an example scene with various

pushing and pick-and-place actions and outcomes from the simulator.

Baseline Approaches: We implement several baselines for comparison to our proposed model *RD-GNN*. *PointConv Relations (PCR)*: defines a *PointConv* based network that takes in a pair of segmented objects and predicts their relations without any GNN, similar to the relational classifiers in [7]. *Pairwise MLP Relational Dynamics (MLP)*: predicts relations and dynamics for pairs of objects using an MLP instead of a GNN to construct the latent space and dynamics. *Direct Pose Dynamics GNN (DPD-GNN)*: uses a GNN to predict the pose for each object conditioned on a chosen action. We use an analytic relational classifier to predict relations from the predicted poses and their associated bounding boxes. *Pose Estimation GNN (PE-GNN)*: We replace the relational output heads on our model with pose estimation regressors. We again use analytic relational classifiers for evaluation. *Combined Relational Dynamics and Pose Estimation (RD-PE-GNN)*: This combines our model with the pose estimation regressor for both the current and next time step. *Relational Dynamics without Latent Regularization (RD-GNN-w/o-LR)*: We train a version of our model without using the \mathcal{L}_{DYN} loss.

Predicting and Detecting Inter-Object Relations: We first examine the efficacy of our model in correctly predicting which relations will be present after executing a specified action. Then we examine the ability of our model to detect inter-object relations for objects in the observed scene post manipulation. We show here the prediction accuracy between the predicted relations and the ground truth relations post-manipulation. On our simulation test data across 300 skill executions, the average prediction F1 score across all relational classifiers predictions for RD-GNN is 0.906, MLP is 0.678, RD-PE-GNN is 0.879, DPD-GNN is 0.319, PE-GNN is 0.133, and RD-GNN-w/o-LR is 0.693. We find that the RD-GNN and RD-PE-GNN perform comparably and outperform other baselines in terms of prediction F1 score. We next show the F1 score between the detected and ground truth relations post manipulation. For our simulation test data, the average F1 score of the post manipulation relational classifiers for RD-GNN is 0.974, MLP is 0.971, RD-PE-GNN is 0.977, PointConv is 0.985, PE-GNN is 0.899, and RD-GNN-w/o-LR is 0.977. We find that our approach performs comparable to the PCR baselines, MLP, RD-PE-GNN, and RD-GNN-w/o-LR. Furthermore, our approach outperforms PE-GNN

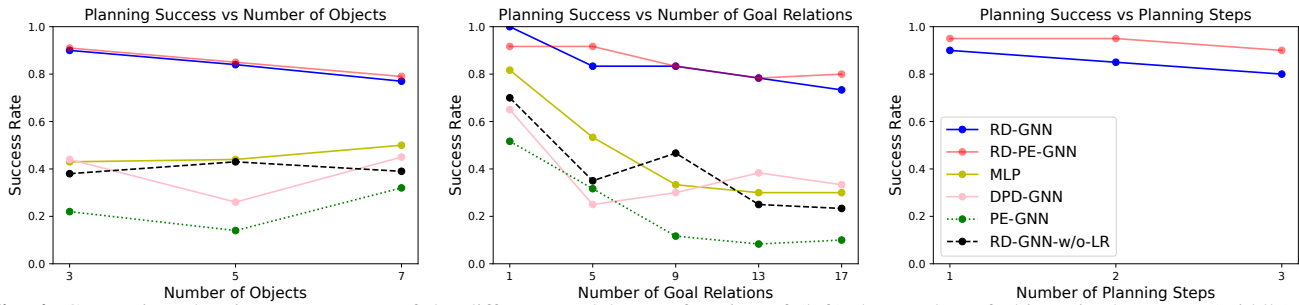


Fig. 4: Comparing planning success rate of the different models as a function of (left) the number of objects in the scene, (middle) the number of relations specified in the goal, and (right) the number of steps. The legend applies to all three plots. We see that RD-GNN and RD-PE-GNN achieve comparable performance while significantly outperforming the baseline models. The success rate drops for all models as we specify more relations in the goal. Even when fully constrained the top performing models achieve high success rates.

in terms of detection F1 score. Note the PCR method alone cannot be used for planning dynamic interactions, which is the main focus of this work.

Planning to Desired Goal Relations: We now examine the ability of our model to plan to desired goal relations with a single action step. We ran 20 planning trials containing varying numbers of objects and goal relations using each model in simulation. We only used pushing tasks for large-scale statistics. Fig. 4 shows that our model, RD-GNN, as well as our model with pose estimation, RD-PE-GNN, dominate all competitors. Fig. 3 shows a variety of successfully executed single-step plans using RD-GNN. Notably we generate diverse plans for the same goal and initial setting.

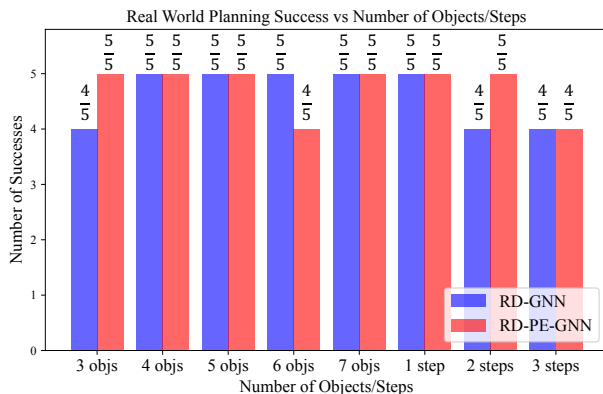


Fig. 5: Number of successes on real-world YCB object manipulation tasks. We compare results for a varying number of objects and varying plan horizon length as denoted by the horizontal labels.

We now turn our attention to multi-step planning focusing on only RD-GNN and RD-PE-GNN as the best single step performers. We show planning success rates for plans ranging in length from 1 to 3 planning steps in simulation in Fig. 4 (right). We use four objects for the multi-step test in simulation and the real world. We see that for both the RD-GNN and RD-PE-GNN approaches, the success rate drops with plan length. Both models achieve high success rates.

We show planning success rate for real-world experiments in Fig. 5. We ran test trials using YCB objects with 5 trials for each setting of varying number of scene objects or plan steps. For all the real-world experiments, we use 5 relations in the goal. Our results verify that our method transfers to real-world settings without any fine tuning and generalizes to real-world objects when trained only on cuboids in simu-

lation. Fig 1 illustrates real-world multi-step plan execution.

V. CONCLUSION

We presented an approach to multi-object manipulation based on learning relational classifiers via graph neural networks. We showed through extensive simulated and physical robot experiments that the relational inductive bias of the graph net provided improved planning success. We can attribute this to better accuracy when predicting both inter-object relations and latent space dynamics. Further we showed inter-object relations provide a better source of supervision for training our model for planning than using a pose estimation loss.

While we find our results quite exciting, several limitations exist in the approach as currently formulated. On a theoretical level we have no proof that the relations we use in training our representation provide a sufficient basis for predicting all inter-object interactions of interest. Currently we only have empirical results to show they seem to work well. As an additional training issue, we only trained with block-shaped objects, while this proved sufficient for demonstrating the benefits of relational graph dynamics, we do not capture detailed shape information that robots must reason about for more complicated tasks and interactions. At a low-level we have not closely integrated the motion planner we use for skill generation with the higher-level CEM planning. This causes the robot to sometimes reject samples that fail to generate motion plans, decreasing efficiency and coverage of our planner. For the multi-step planning, we do not use replanning which requires high accuracy of the latent dynamics model. We think this mechanism will not generalize to very long horizons like 20 steps. So we plan to do replanning for future works.

Overall, our approach provides the first example of predicting multi-object manipulation sequences using learned relational dynamics from partial view point clouds. We leverage these predictions for planning and executing multi-object dynamic rearrangements on a physical robot.

ACKNOWLEDGMENTS

The authors thank Mohit Sharma, Chris Paxton, and Mohanraj Devendran Shanthy for useful discussion. This work was partially supported by NSF Award #2024778, by DARPA under grant N66001-19-2-4035, and by a Sloan Research Fellowship.

REFERENCES

- [1] A. Cosgun, T. Hermans, V. Emeli, and M. Stilman, "Push Planning for Object Placement on Cluttered Table Surfaces," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 11 2011. [Online]. Available: <http://www.cs.utah.edu/~thermans/papers/cosgun-iros2011.pdf>
- [2] L. Chang, J. R. Smith, and D. Fox, "Interactive singulation of objects from a pile," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2012, pp. 3875–3882. [Online]. Available: <https://rse-lab.cs.washington.edu/postscripts/object-singulation-icra-12.pdf>
- [3] M. Gupta and G. S. Sukhatme, "Using manipulation primitives for brick sorting in clutter," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2012, pp. 3883–3889. [Online]. Available: <https://ieeexplore.ieee.org/document/6224787>
- [4] S. Panda, A. A. Hafez, and C. Jawahar, "Learning support order for manipulation in clutter," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2013, pp. 809–815. [Online]. Available: <https://ieeexplore.ieee.org/document/6696444>
- [5] M. R. Dogar, M. C. Koval, A. Tallavajhula, and S. S. Srinivasa, "Object search by manipulation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 4973–4980. [Online]. Available: <https://ieeexplore.ieee.org/document/6631288>
- [6] A. Murali, A. Mousavian, C. Eppner, C. Paxton, and D. Fox, "6-dof grasping for target-driven object manipulation in clutter," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 6232–6238. [Online]. Available: <https://arxiv.org/abs/1912.03628>
- [7] C. Paxton, C. Xie, T. Hermans, and D. Fox, "Predicting Stable Configurations for Semantic Placement of Novel Objects," in *Conference on Robot Learning (CoRL)*, 11 2021. [Online]. Available: <https://arxiv.org/abs/2108.12062>
- [8] W. Liu, C. Paxton, T. Hermans, and D. Fox, "StructFormer: Learning Spatial Structure for Language-Guided Semantic Rearrangement of Novel Objects," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022. [Online]. Available: <https://sites.google.com/view/structformer>
- [9] Y. Zhu, J. Tremblay, S. Birchfield, and Y. Zhu, "Hierarchical planning for long-horizon manipulation with geometric and symbolic scene graphs," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021. [Online]. Available: <https://arxiv.org/abs/2012.07277>
- [10] R. Li, A. Jabri, T. Darrell, and P. Agrawal, "Towards practical multi-object manipulation using relational reinforcement learning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4051–4058. [Online]. Available: <https://arxiv.org/abs/1912.11032>
- [11] Y. Lin, A. S. Wang, E. Undersander, and A. Rai, "Efficient and interpretable robot manipulation with graph neural networks," *IEEE Robotics and Automation Letters*, 2022. [Online]. Available: <https://arxiv.org/abs/2102.13177>
- [12] M. Sharma and O. Kroemer, "Relational learning for skill preconditions," in *Conference on Robot Learning (CoRL)*, 2020. [Online]. Available: <https://arxiv.org/abs/2012.01693>
- [13] M. Wilson and T. Hermans, "Learning to Manipulate Object Collections Using Grounded State Representations," in *Conference on Robot Learning (CoRL)*, 11 2019. [Online]. Available: <https://arxiv.org/abs/1909.07876>
- [14] H. Suh and R. Tedrake, "The surprising effectiveness of linear models for visual foresight in object pile manipulation," in *Workshop on Algorithmic Foundations of Robotics (WAFR)*, 2020. [Online]. Available: <https://arxiv.org/abs/2002.09093>
- [15] A. H. Qureshi, A. Mousavian, C. Paxton, M. Yip, and D. Fox, "NeRP: Neural Rearrangement Planning for Unknown Objects," in *Proceedings of Robotics: Science and Systems*, Virtual, July 2021. [Online]. Available: <https://arxiv.org/abs/2106.01352>
- [16] A. Simeonov, Y. Du, B. Kim, F. R. Hogan, J. Tenenbaum, P. Agrawal, and A. Rodriguez, "A long horizon planning framework for manipulating rigid pointcloud objects," in *Conference on Robot Learning (CoRL)*, 2020. [Online]. Available: <https://arxiv.org/abs/2011.08177>
- [17] W. Yuan, C. Paxton, K. Desingh, and D. Fox, "Sornet: Spatial object-centric representations for sequential manipulation," in *Conference on Robot Learning (CoRL)*. PMLR, 2022, pp. 148–157. [Online]. Available: <https://openreview.net/forum?id=mOLu2rODIJF>
- [18] X. Lou, Y. Yang, and C. Choi, "Learning object relations with graph neural networks for target-driven grasping in dense clutter," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022. [Online]. Available: <https://arxiv.org/abs/2203.00875>
- [19] D. Driess, Z. Huang, Y. Li, R. Tedrake, and M. Toussaint, "Learning multi-object dynamics with compositional neural radiance fields," in *Conference on Robot Learning (CoRL)*, 2022. [Online]. Available: <https://openreview.net/forum?id=qUvTmyGpnm7>
- [20] O. Biza, T. Kipf, D. Klee, R. Platt, J.-W. van de Meent, and L. L. Wong, "Factored world models for zero-shot generalization in robotic manipulation," *arXiv preprint arXiv:2202.05333*, 2022. [Online]. Available: <https://arxiv.org/abs/2202.05333>
- [21] Y. Bisk, K. Shih, Y. Choi, and D. Marcu, "Learning interpretable spatial operations in a rich 3d blocks world," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018. [Online]. Available: <https://arxiv.org/abs/1712.03463>
- [22] C. Paxton, X. Barnoy, K. Katyal, R. Arora, and G. D. Hager, "Visual robot task planning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019, pp. 8832–8838. [Online]. Available: <https://arxiv.org/abs/1804.00062>
- [23] H. Zhang, X. Lan, S. Bai, L. Wan, C. Yang, and N. Zheng, "A multi-task convolutional neural network for autonomous robotic grasping in object stacking scenes," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 6435–6442. [Online]. Available: <https://arxiv.org/abs/1809.07081>
- [24] Y. Jiang, M. Lim, C. Zheng, and A. Saxena, "Learning to place new objects in a scene," *The International Journal of Robotics Research*, vol. 31, no. 9, pp. 1021–1043, 2012. [Online]. Available: <https://arxiv.org/abs/1202.1694>
- [25] Y. Jiang and A. Saxena, "Hallucinating humans for learning robotic placement of objects," in *Experimental Robotics: The 13th International Symposium on Experimental Robotics*. Springer, 2013, pp. 921–937.
- [26] B. Kim, Z. Wang, L. P. Kaelbling, and T. Lozano-Pérez, "Learning to guide task and motion planning using score-space representation," *The International Journal of Robotics Research*, vol. 38, no. 7, pp. 793–812, 2019. [Online]. Available: <https://arxiv.org/abs/1807.09962>
- [27] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling, "Sample-based methods for factored task and motion planning," in *Robotics: Science and Systems*, 2017. [Online]. Available: <https://dspace.mit.edu/bitstream/handle/1721.1/137701/garrett-rss17.pdf?sequence=2&isAllowed=y>
- [28] B. Kim and L. Shimanuki, "Learning value functions with relational state representations for guiding task-and-motion planning," in *Conference on Robot Learning (CoRL)*, 2019. [Online]. Available: <http://people.csail.mit.edu/beomjoon/publications/kim-corl19.pdf>
- [29] D. Driess, J.-S. Ha, and M. Toussaint, "Deep visual reasoning: Learning to predict action sequences for task and motion planning from an initial scene image," in *Proceedings of Robotics: Science and Systems*, 2020. [Online]. Available: <https://arxiv.org/abs/2006.05398>
- [30] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling, "Pddlstream: Integrating symbolic planners and blackbox samplers via optimistic adaptive planning," in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 30, 2020, pp. 440–448. [Online]. Available: <https://arxiv.org/abs/1802.08705>
- [31] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez, "Integrated task and motion planning," *Annual review of control, robotics, and autonomous systems*, vol. 4, pp. 265–293, 2021. [Online]. Available: <https://arxiv.org/abs/2010.01083>
- [32] J. Liang, M. Sharma, A. LaGrassa, S. Vats, S. Saxena, and O. Kroemer, "Search-Based Task Planning with Learned Skill Effect Models for Lifelong Robotic Manipulation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022. [Online]. Available: <https://arxiv.org/abs/2109.08771>
- [33] T. Lozano-Pérez and L. P. Kaelbling, "A constraint-based method for solving sequential manipulation planning problems," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2014, pp. 3684–3691. [Online]. Available: <https://lis.csail.mit.edu/pubs/tpkc-iros14.pdf>
- [34] M. Watter, J. Springenberg, J. Boedecker, and M. Riedmiller, "Embed to control: A locally linear latent dynamics model for control from raw images," *Advances in Neural Information Processing Systems*, vol. 28, 2015. [Online]. Available: <https://arxiv.org/abs/1506.07365>
- [35] P. Battaglia, J. B. C. Hamrick, V. Bapst, A. Sanchez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner,

- C. Gulcehre, F. Song, A. Ballard, J. Gilmer, G. E. Dahl, A. Vaswani, K. Allen, C. Nash, V. J. Langston, C. Dyer, N. Heess, D. Wierstra, P. Kohli, M. Botvinick, O. Vinyals, Y. Li, and R. Pascanu, "Relational inductive biases, deep learning, and graph networks," *arXiv*, 2018. [Online]. Available: <https://arxiv.org/pdf/1806.01261.pdf>
- [36] W. Wu, Z. Qi, and L. Fuxin, "PointConv: Deep Convolutional Networks on 3D Point Clouds," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 9621–9630. [Online]. Available: <https://arxiv.org/abs/1811.07246>
- [37] H. Cui, Z. Lu, P. Li, and C. Yang, "On positional and structural node features for graph neural networks on non-attributed graphs," in *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 2022, pp. 3898–3902. [Online]. Available: <https://arxiv.org/abs/2107.01495>
- [38] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, "On the continuity of rotation representations in neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 5738–5746. [Online]. Available: <https://arxiv.org/abs/1812.07035>
- [39] R. Y. Rubinfeld and D. P. Kroese, *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation, and machine learning*. Springer, 2004, vol. 133.
- [40] M. Kobilarov, "Cross-entropy randomized motion planning," in *Proceedings of Robotics: Science and Systems*, vol. 7, 2012, pp. 153–160. [Online]. Available: <http://www.roboticsproceedings.org/rss07/p22.pdf>
- [41] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The YCB object and model set: Towards common benchmarks for manipulation research," in *International Conference on Advanced Robotics (ICAR)*, 2015. [Online]. Available: https://www.eng.yale.edu/grablab/pubs/calli_icar2015.pdf
- [42] V. Makovychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa *et al.*, "Isaac gym: High performance gpu-based physics simulation for robot learning," in *Advances in Neural Information Processing Systems*, 2021. [Online]. Available: <https://sites.google.com/view/isaacgym-vidia>