

Efficient Learning of Object Placement with Intra-Category Transfer

Adrian Röfer¹, Russell Buchanan^{2,3}, Max Argus¹, Sethu Vijayakumar², and Abhinav Valada¹

Abstract—Efficient learning from demonstration for long-horizon tasks remains an open challenge in robotics. While significant effort has been directed toward learning trajectories, a recent resurgence of object-centric approaches has demonstrated improved sample efficiency, enabling transferable robotic skills. Such approaches model tasks as a sequence of object poses over time. In this work, we propose a scheme for transferring observed object arrangements to novel object instances by learning these arrangements on canonical class frames. We then employ this scheme to enable a simple yet effective approach for training models from as few as five demonstrations to predict arrangements of a wide range of objects including tableware, cutlery, furniture, and desk spaces. We propose a method for optimizing the learned models to enable efficient learning of tasks such as setting a table or tidying up an office with intra-category transfer, even in the presence of distractors. We present extensive experimental results in simulation and on a real robotic system for table setting which, based on human evaluations, scored 73.3% compared to a human baseline. We make the code and trained models publicly available at <https://oplict.cs.uni-freiburg.de>.

Index Terms—Learning from Demonstration, Probabilistic Inference

I. INTRODUCTION

HUMANS excel at teaching each other various skills efficiently. Whether it is setting tables, changing bicycle brakes, or furnishing a room, we can guide another person to proficiency within a handful of training sessions. Deep learning on large datasets has yielded impressive manipulation results on a variety of tasks [1]–[3]. However, the requirement for vast amounts of data by these methods prevents them from achieving human-like efficiency in skill acquisition [4].

One of the reasons for this requirement is the need to distill relevant features from visual observations. Pre-trained vision foundation models that successfully transfer across different object instances have alleviated this problem. They enable one- or few-shot trajectory imitation by using deep vision features for object-centric localization and low-capacity models for encoding motion [5]–[7]. Another method for overcoming the data requirement is the sparsification of tasks.

This paper was recommended for publication by Editor Aleksandra Faust upon evaluation of the Associate Editor and Reviewers’ comments.

This work was funded by the BrainLinks-BrainTools center of the University of Freiburg, the H2020 Project HARMONY, and the Alan Turing Institute. Adrian Röfer acknowledges travel support from the EU H2020 research and innovation programme under grant agreement No 951847.

Manuscript received: November, 5th, 2024; Revised February, 7th, 2025 & August, 12th, 2025; Accepted November, 4th, 2025.

Corresponding Author: Adrian Röfer, email: aroef@cs.uni-freiburg.de.

¹ Department of Computer Science, University of Freiburg, Germany.

² School of Informatics, University of Edinburgh, Scotland.

³ Faculty of Engineering, University of Waterloo, Canada.

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

©2026 IEEE




Fig. 1: Our approach learns object placements sample-efficiently by mapping object instances to a known canonical instance and inferring the placement of the new object in this canonical space. Here, the known setup on the left is matched with the novel one on the right to place the chair.

Subsampling the number of timesteps in motion policies can boost the efficiency of reinforcement-learning [8]. Taken to its extreme, such a sparsification yields a separation of tasks into discrete sequences of motions, with boundary conditions for transitioning between segments [9]–[11]. In this view, instead of learning a full trajectory for placing a monitor on a desk, the agent learns where a monitor should go in relation to keyboard and mouse. The specific motion is then to be inferred at runtime. Such sparse representations have been used successfully in teaching robots simple pick-and-place sequences from a single demonstration [12], [13], while it has also been possible to learn control sequences [14] for the motion segments.

In this work, we aim to support long-horizon task learning by proposing a method for learning *where* objects should be during the *key moments* of sequential placement task. We introduce an approach for learning relative arrangements of objects which is trainable using a minimal number of examples (≤ 5) and transfers novel object instances of the same classes as seen during training. We employ this approach to incremental object placement tasks in which larger everyday scenes of objects have to be arranged. Such natural scenes, similar to Fig. 1, can exhibit complex inter-object dependencies and significant object variability which creates a challenge to reproducing the scenes in a semantically acceptable manner. Our framework operates on distributions of relative poses, for which we introduce a scheme for mapping object observations to a canonical class frame, which enables the desired intra-category transfer. In addition, we introduce a method for minimizing the complexity of models and reducing the impact of spurious correlations formed with unrelated objects. We demonstrate our method in simulations, showing the effectiveness of our method by arranging furniture, tableware, and office space. In simulated experiments, our method was successfully trained with as few as five demonstrations and performed robustly against distractors.

Using the system depicted in Fig. 2, we perform real robot experiments on tableware arrangements. Our system was rated by human evaluators as 73.3% as good as human performance. Our primary contributions are:

- 1) A novel approach for few-shot relative pose learning.
- 2) A framework for mapping observations of objects to a canonical class frame for intra-category transfer.
- 3) A technique for optimizing model complexity and removing distracting observations from the models.
- 4) Real-world robot experiments of autonomously setting tables using both familiar and unfamiliar objects, with the quality of the arrangements assessed by human jurors.
- 5) We publicly release our dataset, code, and models at <https://opliect.cs.uni-freiburg.de>.

II. RELATED WORK

We briefly summarize related work on learning manipulation tasks and learning relative object poses.

Manipulation as a Series of Contacts: Manipulation tasks can be understood as a series of contact state changes instead of dense temporal trajectories. Recent works [12], [13] have demonstrated that it is possible to learn a long-horizon manipulation task from a single demonstration by using changes in contacts as delimiters of robotic actions. Similarly, Mao *et al.* [14] demonstrate that such a delimitation can be used to simplify the combinatorial problem of searching for effective action sequences in long-horizon tasks, with learned parameterizations of actions and success predictors for action primitives and their sequences. In motion planning, this sparse segmentation identifies *kinematic modes* under each of which manipulation is continuous, making it efficiently solvable using numerical optimization [9]. Despite the sparsity, contact-centric representations are also informative as demonstrated by [15] who introduced a representation of tasks as a sequence of contact graphs between objects, which they coined *semantic event chains*, and could recognize objects and tasks from this representation. In [16], they illustrated how their representation can support manipulation, but the focus is on recognition of tasks and objects [17]–[19].

Object Pose Learning: While several works investigate object pose placement from language cues [20]–[22], only a few learn relative object placement directly from category information, without any cues. Image-space approaches are conditioned on language instructions to predict image-space activations for object placement [23]. More recently [24] demonstrated an approach for finding arrangement poses of single objects in cluttered scenes using diffusion processes on point cloud data. In [25], [26], the authors demonstrate approaches that learn language-conditioned rearrangement of objects. These approaches incrementally update a given scene towards the described arrangement until the updates anneal. In [20], [27], object-placement or navigation locations are identified from verbal descriptions. Moreover, recent methods also use pretrained VLMs, or generative models, to generate desired arrangements using language cues with zero-shot transfer. Methods such as [28], [29] either generate desired target arrangements or query pretrained language models for target locations for objects. Although the knowledge that can be

extracted from these models is impressive, they ultimately require both very specific prompts and very good observability of a scene. In contrast, our aim in this work is to solely leverage data obtainable from physical demonstrations without language cues, even without full observability.

The most similar works to our own are [30]–[32], which also do not use language cues. Kapelyukh *et al.* [30] employ graph-neural networks to predict object arrangements from user preferences extracted from demonstrations. They represent the objects in a scene using text embeddings and their positions and train an embedding VAE on the fully connected graph of a scene. They can infer a user’s preference from a scene the user arranges and use their model to rearrange other scenes to the user’s liking. Their follow-up work [31] proposes a model that assigns a score to a given scene, again using a graphical representation and CLIP embeddings [33] of objects. Unlike our approach, these methods require 16 or more training examples, while 5 are sufficient for ours. Their inferences are limited to SE2, while we address full SE3 poses. Most similar to our approach is the *Intention Likelihood* formulation [32]. It also employs relative pose observations and can be trained with as few as 5 demonstrations. Distinctively, it does not learn a distribution over poses but assumes a fixed kernel density and scores poses under it. At inference time, the method places previously seen poses in the scene and scores their likelihood. Our method estimates the density of the pose distributions from the training data and samples these, making our inferences more precise and less sensitive to distractors.

III. PROBLEM DEFINITION

As a multi-step task with complex inter-object dependencies, we study the problem of sequential object placement, which assumes that a sequence of objects $\mathcal{O}_p = \{o_1, \dots, o_n\}$ needs to be placed in a scene in relation to a static set of objects \mathcal{O}_s . Each placed object transitions from \mathcal{O}_p to \mathcal{O}_s . Formally: At time step $t = 1$, the sets are the initial sets \mathcal{O}_p and \mathcal{O}_s , afterwards they change to $\mathcal{O}_{p,t} = \{o_t, \dots, o_n\}$ and $\mathcal{O}_{s,t} = \mathcal{O}_s \cup \{o'_1, \dots, o'_{t-1}\}$ for all $t \leq n$.

Each object o has a world-space pose ${}^W\mathbf{T}_o$ and a class c_o . Each class is associated with a set of feature points, $\mathcal{F}_c = \{(e_1, p_1), \dots, (e_m, p_m)\}$ which consist of an embedding $e \in \mathbb{R}^E$ and a position $p \in \mathbb{R}^3$ in a canonical class frame. The generation of feature points and their utility for manipulation and pose estimation and tracking is widely studied in robotics.

Given this structure, we seek to learn the distribution of object placements in the scene, given its class, and the poses and classes of the other already placed objects: $p_t \left({}^W\mathbf{T}_{o_t} \mid c_{o_t}, c_{o'_1}, \dots, c_{o'_{t-1}}, {}^W\mathbf{T}_{o'_1}, \dots, {}^W\mathbf{T}_{o'_{t-1}} \right)$, with $o'_1, \dots, o'_{t-1} \in \mathcal{O}_{s,t}$.

IV. POSE LEARNING APPROACH

In order to learn poses effectively, we decompose of learning world space poses $p_t \left({}^W\mathbf{T}_{o_t} \mid \dots \right)$ into

$$p_t \left({}^W\mathbf{T}_{o_t} \mid c_{o_t}, c_{o'_1}, \dots, c_{o'_{t-1}}, {}^W\mathbf{T}_{o'_1}, \dots, {}^W\mathbf{T}_{o'_{t-1}} \right) = \prod_{o' \in \mathcal{O}_{s,t}} p_t \left({}^{o'}\mathbf{T}_{o_t} \mid c_{o_t}, c_{o'} \right), \quad (1)$$

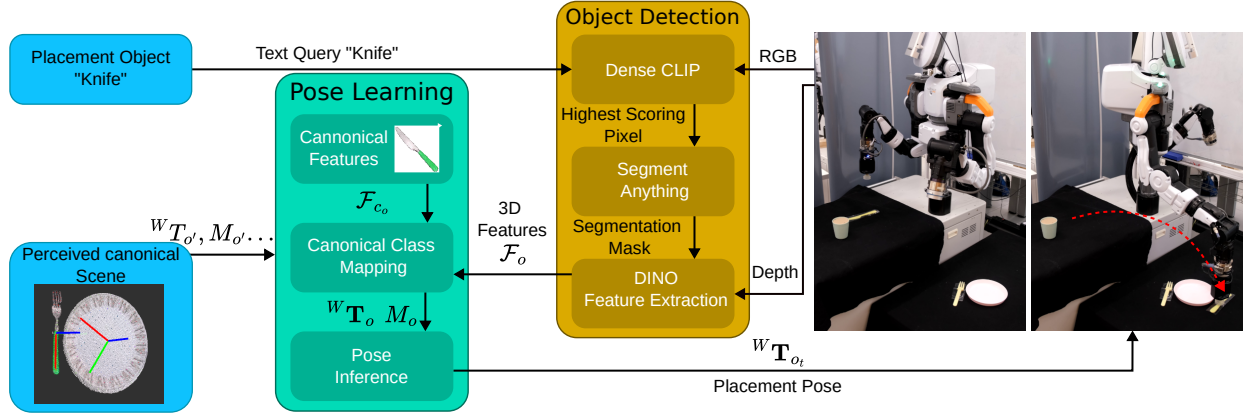


Fig. 2: Full pipeline of our system with real robot experiment. Our proposed pose inference method predicts the ideal object placement pose, which the robot then arranges autonomously. Our approach’s few-shot transfer to other object instances is enabled by our object class mappings which are enabled by several large networks for object detection and feature extraction.

learning of relative poses ${}^o\mathbf{T}_{o'}$, which we derive as ${}^o\mathbf{T}_{o'} = {}^W\mathbf{T}_{o'}^{-1} \cdot {}^W\mathbf{T}_{o'}$. To enable intra-category transfer, we assume there exists an invertible class mapping $\mathcal{M}(o) = o_c$ which maps an object’s properties, i. e. its pose and feature points, to the canonical categorical instance. These concepts, depicted in Fig. 1, yield the probability of a relative pose given the observed distribution of poses in categorical space as

$$p_t \left({}^o\mathbf{T}_{o'} \mid c_{o'}, c_o \right) = p \left({}^o\mathbf{T}_{o'} \mid \mathcal{M}(o') {}^o\mathbf{T}_o \mathcal{M}(o_t) \right). \quad (2)$$

A. Canonical Class Mappings

Class maps deform an instance of an object to best match a known canonical object. Our approach requires these maps to be affine transformations, but in this work, we only consider linear scaling instances. The simplest considered map is the identity $\mathcal{M}_I = I$, which does not scale an object to its categorical representation. The *uniform linear map*, which was used in [34] applies a single scaling factor s , to scale the observed instance of the object to the class prototype. We denote the map as

$$\mathcal{M}_U(s) = \text{diag}(s, s, s, 1). \quad (3)$$

Given an observed instance o and its feature points $\mathcal{F}_o = \{p_1, \dots, p_m\}$, we can derive the scaling factor s easily, by comparing the distances between point pairs in \mathcal{F}_o and $\mathcal{F}_{c_o} = \{\hat{p}_1, \dots, \hat{p}_m\}$ as

$$s = \frac{1}{|\mathcal{F}_o|^2} \sum_{i=1}^m \sum_{j=1}^m \frac{\|\hat{p}_i - \hat{p}_j\|}{\|p_i - p_j\|}. \quad (4)$$

We find this mapping to perform robustly, as it is simple to derive. Nonetheless, its assumption of a single scaling factor is limiting, as objects can vary quite significantly and non-uniformly in their extents, i. e. slender wine glasses compared to bulbous ones. For such cases, we propose the generalization \mathcal{M}_O of the previous mapping as:

$$\mathcal{M}_O(s_x, s_y, s_z) = \text{diag}(s_x, s_y, s_z, 1), \quad (5)$$

which we refer to as *orthogonal linear map*. Deriving the values for s_x, s_y, s_z is more challenging in this case, as the extents have to be measured in the object’s frame ${}^W\mathbf{T}_o$, while

the estimate of the scaling factors also affects the estimation of this frame. Thus, we jointly optimize fit and pose as

$$\min_{{}^W\mathbf{T}_{o'}, s} \sum_{i=0}^m \|p_i - \mathcal{M}_O(s_x, s_y, s_z)^{-1} {}^W\mathbf{T}_{o'} \hat{p}_i\|, \quad (6)$$

where $\mathcal{M}_O(s_x, s_y, s_z)^{-1}$ is the inverse of the class mapping, thus mapping the class’ feature points to the space of the observed object instance. With this estimate of \mathcal{M}_O , we attempt a better fit of ${}^W\mathbf{T}_o$. This process continues until it converges or until a fixed step limit is reached.

B. Learning Relative Pose Distributions

To learn the relative poses of objects in category space, we use multivariate normal distributions. Our distributions capture the relative pose ${}^b\hat{\mathbf{T}}_a$ of two objects a, b where ${}^b\hat{\mathbf{T}}_a = \mathcal{M}_b {}^b\mathbf{T}_a$, with \mathcal{M}_b being the class mapping for the object b . While ${}^b\hat{\mathbf{T}}_a \in \mathbb{R}^{4 \times 4}$ is a convenient representation for computations, it does not lend itself to learning due to its size and redundancy. Instead, we represent these poses in a lower-dimensional feature space. We represent the encoding into this space as an invertible function f . These distributions capture the conditional probability defined in Eq. (2):

$$p_t \left({}^b\mathbf{T}_a \mid c_a, c_b \right) = p \left(f \left({}^b\hat{\mathbf{T}}_a \right) \right). \quad (7)$$

We use these pair-wise distributions to learn one joint placement distribution for each $o \in \mathcal{O}_p$. The naive approach is the substitution of the pair-wise conditional probability in Eq. (1) for our newly derived pair-wise probability. We refer to this unidirectional model as $p_{U,t}$. A problem with this model is its disregard for the placed object’s geometry, as only the reference object’s class mapping is used. We propose a bidirectional model $p_{B,t}$ which includes the observation of the reference object:

$$p_{B,t} \left({}^o\mathbf{T}_{o'} \mid c_{o'}, c_o \right) = p \left(f \left({}^o\hat{\mathbf{T}}_o \right) \right) p \left(f \left({}^o\hat{\mathbf{T}}_{o'} \right) \right). \quad (8)$$

The full model for placing one object given all other objects in the scene can then be derived by substituting Eq. (1). For inference on either model, we sample uniformly from $o' \in \mathcal{O}_{s,t}$ and draw a sample x from its relative distribution. We use




Fig. 3: Illustration of the observation augmentation procedure used for model pruning. The observation of the cup relative to the sofa in Scene 1 is transferred to Scene 2. While the transferred observation is scored the same from the point of the sofa, from the point of the table, it is scored far lower, informing us that including the table improves our model.

${}^W\mathbf{T}_x = {}^W\mathbf{T}_{o'}(M_{o'}^{-1}f^{-1}(x))$ to compute its corresponding pose and then score the sample under the probability density function of our joint distribution p_t . Naively, we pick the highest-scoring sample. Experimentally we observed that the average sample \hat{x} of the top 10 samples X_{10} almost always scored better than each individual sample, i.e. $\forall x \in X_{10} : p_t(\hat{x}) \geq p_t(x)$. We exploit this in a simple refinement scheme: We draw and initial 10^5 samples, take the top 10 samples according to p_t , and form a new distribution of world-space poses with mean \hat{x} and variance $\text{var}(X_{10})$, represented as position and rotation vector. We sample a new 10^3 from this distribution and repeat the process until we see only minor increases in p_t or until a step limit is reached.

C. Pose Encoding & Model Minimization

The choice of pose-encoding f affects the type of relative relationship which can be captured by $p(o' \mathbf{T}_{o'})$. In this work, we examine a number of choices for f and compare their general utility for inference in Sec. V-A. Instead of choosing a fixed encoding for all relations, we would like to be able to autonomously identify the ideal pose representation for a given set of relative pose observations. To do so we propose selecting the encoding $f_{t,o',o}$ for these observations as

$$f_{t,o',o} = \arg \min_{f \in F} H(\mathcal{N}(\mu(f(o' \mathbf{T}_{o'})); \Sigma(f(o' \mathbf{T}_{o'}))), \quad (9)$$

where H denotes the Shannon information of the distribution fitted to the relative pose observations encoded using f . We use the specialized entropy for multivariate Gaussians $H = \frac{1}{2} \ln((2\pi e)^k \det \Sigma)$. Assuming that two encodings f, g do not produce vastly differently scaled vector spaces, we can use this criterion to evaluate the tightness of the model's fit.

With the described method, we can optimize the choice of observation encoding. As we described the method so far, we form distributions for all relative object observations in a scene. Especially in larger scenes, many of the relative pose observations will not be relevant to the placement of an object. Imagine placing a cup on a coffee table that is set in front of a sofa. While the relative pose of cup and sofa will have a statistical trend, semantically, this correlation is (largely) spurious. When performing inference, maintaining all of these relationships potentially increases the number of samples needing to be drawn and the numerical instability.

Thus, given the fitted distributions $p_{o_1}({}^{o_1}\mathbf{T}_{o_t}), \dots, p_{o_{t-1}}({}^{o_{t-1}}\mathbf{T}_{o_t})$, we seek to identify the minimal set of distributions $O_{s,t}^* \subseteq O_{s,t}$ which represents the data. We propose doing so by using an outlier-discrimination strategy: We can expect all training samples ${}^W\mathbf{T}_{o_t}^k$ to be reasonably probable under the fitted distributions in the originally observed context of scene k . However, we can produce potential outliers by taking observation ${}^{o_1}\mathbf{T}_{o_t}^i$ from scene i and introducing it into scene j as ${}^W\tilde{\mathbf{T}}_{o_t}^j = {}^W\mathbf{T}_{o_1}^j \cdot {}^{o_1}\mathbf{T}_{o_t}^i$. We can now evaluate $y_j = p_{o_2}({}^{o_2}\mathbf{T}_{W^j} \cdot {}^W\tilde{\mathbf{T}}_{o_t}^j)$ and compare it to $y_i = p_{o_2}({}^{o_2}\mathbf{T}_{o_t}^i)$. If p_{o_1} and p_{o_2} identify largely overlapping regions, i.e. they represent redundant information, we would expect $y_i \approx y_j$. If the distributions are not redundant, we would expect $y_i \gg y_j$. Fig. 3 illustrates this process. We use a sampling-based scheme to incrementally build $O_{s,t}^*$. We initialize this set as $O_{s,t}^* = \{o_1'\}$ where o_1' is sampled from $O_{s,t}$ according to $p(o' \in O_{s,t}^*) \propto H(o' \mathbf{T}_{o_t})$. Given this root object, we generate K^2 potential outlier observations. For each object $o_i' \in O_{s,t}$ with $o_i' \neq o_1'$, we compute $y_{i,k,j}$. We define the set of rejected samples of an object o as

$$R(o) = \left\{ (k, j) \mid \frac{y_{o,k,j}}{y_{o,k,k}} < \alpha \right\}, \quad (10)$$

and the rejected samples of a set of objects as $R(O) = \bigcup_o R(o)$. We now incrementally expand $O_{s,t}^*$ by sampling an object $o' \in O_{s,t}/O_{s,t}^*$ with $p(o') \propto |R(o)|/R(O_{s,t}^*)$. For a selected sample, we calculate a score $s_{o'}$ as

$$s_{o'} = \frac{|R(o')/R(O_{s,t}^*)| \cdot \tilde{H}(O_{s,t}^*)}{(\hat{K} - |R(O_{s,t}^*)|)(\tilde{H}(O_{s,t}^*) - \tilde{H}(O_{s,t}^* \cup \{o'\}))}, \quad (11)$$

where $\hat{K} = K^2 - K$ and $\tilde{H}(O) = \frac{1}{|O|} \sum_{o \in O} H(p(o \mathbf{T}_{o_p}))$ is the mean entropy of the distributions of a set. Intuitively, this score trades off the fraction of remaining outliers with the relative rise in mean entropy of the chosen set. We compare $s_{o'}$ to a random sample $\epsilon \in [0, 1]$ and admit o' to $O_{s,t}^*$ if $s_{o'} > \epsilon$. Otherwise, we terminate the assembly of $O_{s,t}^*$. We repeat this process multiple times and pick the best-sampled model according to scoring lowest under $(1 - |R(O)|/\hat{K}) \cdot \tilde{H}(O)$. To limit the computational cost of generating hypothetical scenes, we prefilter a subset $\hat{O}_{s,t} \subseteq O_{s,t}$, where $p(o \in \hat{O}_{s,t}) \propto H(p_o)$ and $|\hat{O}_{s,t}| \propto K$ from which we select objects. The complexity of this process is proportional to $|\hat{O}_{s,t}|K^2$.

V. EXPERIMENTAL EVALUATION

We evaluate our approach both in simulation and on a real robot. First, in simulation, we measure the impact of feature encodings, category maps, relative pose models, impact of training samples, and impact of distractor objects. Second, we deploy the best performing model on a real robot and use it for table setting with seen and unseen cutlery.

A. Simulation Experiments

For the simulated experiments, we use the five scenes shown in Fig. 4. We train our models with different feature and category maps for predicting the placement poses of the




Fig. 4: We evaluate five different training scenarios with 16 variations for training (only 5 shown here) and 5 for evaluation each. The scale of objects changes non-uniformly and the placements are varied. They are hand-crafted to ensure that they are semantically meaningful. For each object category, we generate 12 key points, which remain the same across all instances. From the top left: *Dinner places*, *Bread-cutting*, *Desks*, *Living room*, and *TV setup*.

objects with predetermined placement order. We evaluate model performance in n -step inference, where n truncates the length of the task from the back, starting at the final placement step t . This is best understood in reference to the placement sequence of object \mathcal{O}_p : The n -step inference of object o_t reduces \mathcal{O}_p to $\hat{\mathcal{O}}_p = \{o_{t-n}, \dots, o_t\}$, with the initial observation set $\mathcal{O}_{s,t-n}$. For example, 0-step inference is the placement of the current object, where previously placed objects have ground truth pose. As n increases, the error in the scenes accumulates for each placed object.

We evaluate the performance of different combinations of category maps \mathcal{M} , pose representations f , models p_U, p_B , and the impact of model minimization by fitting our models to 5 training samples and evaluating on 5 test scenes. Our choices for pose encodings are $f_{quat}, f_{AA}, f_{euler}$ which all encode poses as Cartesian position and rotations as a quaternion, rotation vector, and Euler angles, respectively. The $f_{\mathfrak{se3}}$ encoding uses the log-map to represent a pose as an element of $\mathfrak{se3} \subset \mathbb{R}^6$. The f_{mix} encoding chooses between the aforementioned ones according to the heuristic described in Sec. IV-C. We start by comparing the impact of combinations of \mathcal{M} and f and the performance of our two competing models p_U, p_B . Based on the results we select one model combination to ablate for long-horizon performance and to study the impact of our model minimization approach in the face of distractors.

Results: We present the results in Tab. I for 0-step inference. We find that the orthonormal class map \mathcal{M}_O and f_{AA} feature encoding achieve the most accurate result combining position and rotation. We note that the difference in positional error between the \mathcal{M}_O and \mathcal{M}_U is minor, while not using either scaling map incurs large errors. We note that the encoding f_{quat} performs the worst out of all our pose encodings. Similarly, the automatically chosen pose representation f_{mix} does not yield any benefit in 0-step inference. We compare the performance of p_U, p_B using the \mathcal{M}_O category map and f_{AA} pose representation and find p_B to perform better on position ($5.3\% < 5.7\%$), and angle ($16.7^\circ < 21.0^\circ$). To evaluate our model’s performance, we compare with SceneScore

TABLE I: Simulation results with a comparison of class maps and different feature encodings in 0-step inference. We bold-face the best performance in each row. We normalized the positional errors by the extent of the scenes before averaging them to account for different scene scales. The column *min.* indicates the incurred error when a model is minimized.

Class map	Pose encoding	$\Delta\%$		Δ°	
		base	min.	base	min.
\mathcal{M}_I	f_{quat}	14.6	14.2	17.6	17.1
	f_{mix}	14.2	14.3	18.1	18.0
	f_{euler}	13.9	13.9	18.0	20.0
	f_{AA}	13.6	14.8	18.3	21.3
	$f_{\mathfrak{se3}}$	13.4	14.6	18.3	21.3
\mathcal{M}_O	f_{quat}	5.8	5.9	17.6	18.2
\mathcal{M}_U	f_{quat}	5.8	5.7	17.0	18.4
	f_{mix}	5.6	5.6	20.4	20.7
\mathcal{M}_O	$f_{\mathfrak{se3}}$	5.6	6.7	17.7	20.0
	f_{mix}	5.5	5.7	20.4	19.9
\mathcal{M}_U	f_{AA}	5.5	5.7	18.8	21.1
	f_{euler}	5.5	5.5	18.0	19.7
	$f_{\mathfrak{se3}}$	5.4	5.8	17.1	17.8
\mathcal{M}_O	f_{euler}	5.4	5.5	18.7	19.4
	f_{AA}	5.4	5.5	18.6	19.1

(SS) [31] and the *Intention Likelihood* (IL) model [32], a VLM baseline [35], as well as standard models. We obtain the implementation details of IL from [36] and use the same hyperparameters as [36]. For SS, we use the official code release and reduce our problem to 2D inferences on the XY plane. We train SS with the original training routine and hyperparameters on $k = 16$ training samples, as this is the lowest sample count reported by the authors. We present the VLM with overhead images of a current scene and a cropped image of the object to be placed. We prompt the model to state the pixel location where the new object should be placed. Please see the supplementary* for further details. In addition to these works, we also deploy an MLP on our scenes by stacking pose observations encoded as f_{AA} and object scales

*Supplementary: <https://arxiv.org/abs/2411.03408>

TABLE II: Comparison of bi-directional model p_B using the f_{mix} and f_{AA} pose representations against the three best baselines from Fig. 5 in the face of a growing number of distractors. Columns $t < n$ show the mean error over a n -step inference. The lowest inference errors are highlighted per d -block and column. As the number of distractors increases, the minimization manages to eliminate distractor objects. The correlated rise of inference error and number of distractors indicates that neither method filters out all distractors. We note that Intention Likelihood (IL) is particularly sensitive to distractors, but also SceneScore’s (SS¹⁶) performance is impacted. In the right-most four columns we report the impact of the number of training samples. Given a larger number of samples, the benefit of our minimization scheme dissipates, while it demonstrates a particular effectiveness given fewer training samples.

Distractors	Minimized	Model	f	Position Error in $\Delta\%$			Angular Error in Δ°			$\Delta\%, t < 15$			
				$t < 5$	$t < 10$	$t < 15$	$t < 5$	$t < 10$	$t < 15$	$k = 3$	$k = 4$	$k = 5$	$k = 16$
$d = 0$		SS ¹⁶		18.2	18.5	18.5	38.0	39.5	40.1	–	–	–	18.5
		MLP	f_{AA}	10.0	10.2	10.2	27.7	30.6	31.4	13.5	11.3	10.2	7.3
		IL		9.6	9.6	9.6	23.5	25.1	25.6	9.1	9.4	9.6	8.7
	×	p_B	f_{mix}	5.5	5.6	5.7	17.1	18.9	19.3	6.6	5.9	5.7	4.7
	✓	p_B	f_{mix}	5.4	5.5	5.6	16.4	18.4	19.0	6.6	5.9	5.6	4.7
	×	p_B	f_{AA}	5.3	5.5	5.5	16.7	18.6	19.1	6.6	5.9	5.5	4.7
	✓	p_B	f_{AA}	5.3	5.4	5.4	16.7	18.8	19.3	6.7	5.8	5.4	4.6
$d = 1$		IL		35.1	34.9	34.9	57.4	58.5	58.9	30.1	36.6	34.9	56.1
		SS ¹⁶		27.9	28.6	28.7	48.4	50.3	50.6	–	–	–	28.7
		MLP	f_{AA}	13.7	13.8	13.8	40.1	44.2	45.2	16.7	14.5	13.8	9.5
	×	p_B	f_{AA}	9.6	9.8	9.9	24.1	25.9	26.3	6.5	8.5	9.9	4.8
	×	p_B	f_{mix}	8.9	9.1	9.1	20.6	22.5	23.0	7.1	7.8	9.1	4.7
	✓	p_B	f_{AA}	8.9	9.1	9.1	22.2	24.2	24.7	6.6	7.4	9.1	4.7
	✓	p_B	f_{mix}	8.6	8.8	8.8	18.4	20.3	20.9	7.2	6.8	8.8	4.7
$d = 5$		IL		60.2	59.5	59.3	99.4	99.6	99.7	56.1	59.8	59.3	80.8
		SS ¹⁶		36.1	36.2	36.1	92.4	94.0	94.0	–	–	–	36.1
	×	p_B	f_{AA}	35.0	35.6	35.7	77.6	79.9	80.5	22.1	28.1	35.7	4.9
		MLP	f_{AA}	19.9	20.3	20.3	57.8	61.0	61.7	30.9	24.9	20.3	12.5
	×	p_B	f_{mix}	17.8	18.1	18.2	42.2	44.4	45.1	13.8	19.9	18.2	4.7
	✓	p_B	f_{AA}	12.8	13.0	13.0	26.8	28.9	29.3	9.9	8.9	13.0	4.8
	✓	p_B	f_{mix}	12.8	13.0	13.0	25.3	27.3	27.9	11.4	6.9	13.0	4.7




Fig. 5: Results of our baseline comparison. For each method, we report the best-performing configuration. Our approach produces the lowest inference errors, followed by an MLP without object scaling information, Intention Likelihood, and Scene Score with 16 training samples SS¹⁶. The VLM performs slightly lower than SS¹⁶ and seems distracted by rotation inference.

and requiring it to infer the next object’s pose. As MLPs require more training data, we augment the original k samples by applying random translations and rotations to the scene. We report metrics in Fig. 5. We include our Top-3 competitors (MLP, IL, SS¹⁶) in our comparisons.

We study both the long-horizon performance of our models and their sensitivity to distractor objects. In Tab. II, we show the average inference errors over longer inference sequences by computing the mean error incurred up to a step length of t . Additionally, we contrast the performances under a rising number of distractors. We do so by generating a set of up to 5 distractor objects d per scene instance and including them as observations in the model’s fitting process. The distractors are distributed uniformly across the volume of the scene with an additional 50% margin and possess uniformly sampled orientation and scale. We note that inference errors remain steady over the length of an inference within one condition. As the number of distractors increases, our proposed minimization

approach demonstrates its benefit. Yet, the rise in inference error overall does indicate that some distractor objects remain under consideration from the model. In Tab. II, we also show cumulative position error with respect to the number of training samples k . While the most similar baseline IL, performs well in the noise-free condition with ample samples ($d = 0, k = 5$), its constant kernel density assumption becomes problematic in the noisier settings. Moreover, as the success of its inference is dependent on its training samples, it also suffers under scarcer training data. While the deep models (MLP, SS¹⁶) are more robust to noise, the MLP’s inference is impacted by decreasing training samples. Given the results in Fig. 5, we do not ablate k for SS¹⁶. Interesting is that our model’s performance improves with lower k in $d > 0$ setting. This can be attributed to the regularization of the model size dependent on the number of training samples as described in Sec. IV-C. It seems the chosen hyperparameter of our minimization is more apt for less training data. We present a qualitative impression of the best and worst inferences according to the model’s own scores in Fig. 6. Despite these good results, we want to acknowledge that our model can only learn unimodal setups, while our baselines MLP, IL, and SS¹⁶ do not share this limitation.

B. Real Robotic Experiments: Table Setting

To evaluate our method on a real robot, we consider a table setting task. We implement a perception pipeline for detection and pose estimation of several intra-category instances of common tableware such as plates, bowls, cutlery, etc. Our pipeline infers object pose ${}^W\mathbf{T}_{o_t}$ and category map \mathcal{M}_{o_t} from RGB-D data as shown in Fig 2. We leverage several publicly available deep learning models for perception. We associate each class c with a natural language label i. e. *fork*




Fig. 6: Renderings of scenes with highest max -step inference score, instead of lowest metric error. **GT**: Ground truth scene; **Best**: Best inference on this scene instance. **Worst**: Inference with the lowest joint score across all inference steps on the scene. The samples seem very close with no noticeable difference between them as a result of the refinement sampling. The scenes seem semantically plausible, indicating that model score is indicative of success. Yet, in the bread-cutting scene and the desk scene, thin objects such as a knife and keyboard get placed slightly below the supporting object’s surface.



Fig. 7: Eight examples of table settings by our robot. Our system is trained only with examples using real cutlery objects. Using our approach, the trained models transfer zero-shot to novel object instances, even though these vary significantly in their size. Using the training-free perception pipeline we describe, we are able to correctly associate object types despite stark visual differences.

and use CLIP [33] and the MaskCLIP technique [37] to extract the most relevant region in the RGB image. We feed the center of the highest scoring region to Segment Anything [38], which predicts a segmentation of the relevant object. Similar to Goodwin *et al.* [34], we use DINOv2 [39] to generate dense features for the pixels under the mask. Using back projection from the depth image, we obtain a 3D feature cloud \mathcal{F}_o . While [34] presents a method for category-level 6D pose estimation using such features, the stability of this approach is not sufficient for our task. Instead, we use the moments of our gathered points \mathcal{F}_o to initialize a 6D pose. We apply a three-dimensional PCA transform from our categorical feature set \mathcal{F}_c to \mathcal{F}_o and compare the resulting features at the extremes of the longest moment of our point clouds. If the points match their opposite end better than the current one, we rotate the 6D pose around the shortest moment. The category map \mathcal{M}_U is then identified from the ratios between the moments. In real robot experiments, due to noise in the depth image, we were unable to estimate an \mathcal{M}_O mapping reliably and used \mathcal{M}_U . We used a Kawada Robotics Nextage dual-armed robot to perform our table setting experiments. Our perception pipeline requires an average of 1050 ms to extract an object and match its pose and up to 7.5 GB of VRAM on a NVIDIA RTX 4060 Ti with 16 GB of VRAM. As we only invoke the pipeline once per placement, we consider this latency acceptable.

Experimental Setup: For this experiment, we set up two tables as shown in Fig 2. On one table, the objects are arranged randomly. On the other table is a single static object to initiate the relative pose prediction. The robot is given a list of objects and the order they must be placed in. However, object detection, grasping, predicting placement pose, and placement are all performed autonomously. First, the robot examines the table of placement objects, detects the query object, and grasps it. The robot then examines the table of static objects, locating all

the current static objects, and predicts where to place the new object, and then places it. The process continues until there are no more placement objects. We train models for these tasks using four demonstrations, each with only the real tableware. As some objects, such as plates and cups, are rotationally invariant, we minimize the variance in rotations for these objects among the demonstrations.

Results: In Fig. 7, we show eight examples of tables set by the robot. This includes novel object instances and combinations of different cutlery sets. The robot adapted object placement for different-sized objects. For example, in Fig. 7 g) in spite of the smaller children’s plate, the robot still places the fork and knife close to the plate. We also randomized the initial placement of the tray, and the robot inferred that the whole scene should be rotated. We found the pose prediction to be very robust. Most failures in placement were due to errors in the perception system, such as sensor noise in the depth measurement.

In an effort to better understand our method’s performance, we surveyed 111 people, asking them to rate thirty images of table settings done by our system on a 10-point scale. We included 10 images of table settings done by humans to establish a baseline. We compared two scenes, one without a tray (similar to Fig. 7 a-d) and one with a tray (Fig. 7 e-h). Scaling the average robot scores by average human scores, the robot was rated 73.3% as good as humans for setting tables without the tray. With the tray, this reduced to 62.2%. This can be attributed due to the respondents rating the human tray setting higher than the table setting. We suspect the presence of the tray led respondents to rate the human settings more highly as the cutlery was better aligned with the tray.

VI. CONCLUSION

In this work, we presented an approach for efficient learning of relative object placement poses with intra-categorical transfer.

We achieved this by introducing a mapping from observed objects to canonical class features, enabling transfer to unseen instances of different poses and scales. In simulated evaluations, we identify that a bi-directional model with an encoding of poses as positions and rotation vectors performs the best quantitatively, and we find that our model minimization approach is successful at removing distractors. We demonstrated that our approach can be deployed on real robotic systems to set tables with different, unseen object instances of varied scales. In human evaluations, our method was rated as good as 73.3% compared to a human table setting baseline. We view this approach as a successful step towards efficient learning of object placements from demonstrations.

Going forward, we would like to consider including multi-modal models and exploring further pose encoding options, such as spherical coordinates, in our approach. The formulation of our approach lends itself to the inclusion of multi-modal models, but determining the number of modes can be challenging [40]. Further, we would like to consider a different type of class mappings. Our chosen linear projections work but do not consider the affordances of objects, implying that these are always located similarly. This would make it difficult to reliably place a key in a keyhole on a door. Learning relative poses of feature points might be more useful, but may also require more data, i. e. as done in [5].

ACKNOWLEDGMENT

We would like to thank our colleague Imen Mahdi for offering her expertise and help in creating the VLM-baseline. In addition, we thank [kenney.nl](#) for their free 3D assets.

REFERENCES

- [1] S. Bahl, R. Mendonca, L. Chen, U. Jain, and D. Pathak, "Affordances from human videos as a versatile representation for robotics," in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2023.
- [2] C. Chi, et al., "Diffusion policy: Visuomotor policy learning via action diffusion," *arXiv:2303.04137*, 2024.
- [3] E. Chisari, N. Heppert, M. Argus, T. Welschhold, T. Brox, and A. Valada, "Learning robotic manipulation policies from point clouds with conditional flow matching," *Conference on Robot Learning*, 2024.
- [4] E. Chisari, T. Welschhold, J. Boedecker, W. Burgard, and A. Valada, "Correct me if i am wrong: Interactive learning for robotic manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, 2022.
- [5] J. Gao, Z. Tao, N. Jaquier, and T. Asfour, "K-vil: Keypoints-based visual imitation learning," *IEEE Tran. on Rob. and Aut.*, 2023.
- [6] N. Heppert, M. Argus, T. Welschhold, T. Brox, and A. Valada, "Ditto: Demonstration imitation by trajectory transformation," in *Int. Conf. on Intelligent Robots and Systems*, 2024.
- [7] J. O. von Hartz, E. Chisari, T. Welschhold, W. Burgard, J. Boedecker, and A. Valada, "The treachery of images: Bayesian scene keypoints for deep policy learning in robotic manipulation," *IEEE Robot. and Autom. Lett.*, 2023.
- [8] I. Nematollahi, E. Rosete-Beas, A. Röfer, T. Welschhold, A. Valada, and W. Burgard, "Robot skill adaptation via soft actor-critic gaussian mixture models," in *IEEE Int. Conf. on Robotics and Automation*. IEEE, 2022, pp. 8651–8657.
- [9] M. Toussaint, J. Harris, J.-S. Ha, D. Driess, and W. Hönig, "Sequence-of-constraints mpc: Reactive timing-optimal control of sequential manipulation," in *Int. Conf. on Intelligent Robots and Systems*. IEEE, 2022, pp. 13 753–13 760.
- [10] Y. Wang, N. Figueroa, S. Li, A. Shah, and J. Shah, "Temporal logic imitation: Learning plan-satisficing motion policies from demonstrations," in *Conf. on Robot Learning*. PMLR, 2023, pp. 94–105.
- [11] E. Merlo, M. Lagomarsino, E. Lamon, and A. Ajoudani, "Exploiting information theory for intuitive robot programming of manual activities," *IEEE Transactions on Robotics*, pp. 1–17, 2025.
- [12] D. Guo, "Learning multi-step manipulation tasks from a single human demonstration," *arXiv:2312.15346*, 2023.
- [13] Y. Zhu, A. Lim, P. Stone, and Y. Zhu, "Vision-based manipulation from single human video with open-world object graphs," *arXiv:2405.20321*, 2024.
- [14] J. Mao, T. Lozano-Pérez, J. B. Tenenbaum, and L. P. Kaelbling, "Learning reusable manipulation strategies," in *Conf. on Robot Learning*, 2023.
- [15] E. E. Aksoy, A. Abramov, J. Dörr, K. Ning, B. Dellen, and F. Wörgötter, "Learning the semantics of object–action relations by observation," *Int. Journal of Robotics Research*, 2011.
- [16] E. E. Aksoy, B. Dellen, M. Tamosiunaite, and F. Wörgötter, "Execution of a dual-object (pushing) action with semantic event chains," in *IEEE-RAS Int. Conf. on Humanoid Robots*, 2011, pp. 576–583.
- [17] A. Toumpa and A. G. Cohn, "Object-agnostic affordance categorization via unsupervised learning of graph embeddings," *Journal of Artificial Intelligence Research*, vol. 77, pp. 1–38, 2023.
- [18] F. Ziaetabar, R. Safabakhsh, S. Momtazi, M. Tamosiunaite, and F. Wörgötter, "A hierarchical graph-based approach for recognition and description generation of bimanual actions in videos," *arXiv:2310.00670*, 2023.
- [19] G. Sochacki, A. Abdulali, N. K. Hosseini, and F. Iida, "Recognition of human chef's intentions for incremental learning of cookbook by robotic salad chef," *IEEE Access*, vol. 11, pp. 57 006–57 020, 2023.
- [20] D. Kim, N. Oh, D. Hwang, and D. Park, "Lingo-space: Language-conditioned incremental grounding for space," *arXiv:2402.01183*, 2024.
- [21] G. Zhai, et al., "Sg-bot: Object rearrangement via coarse-to-fine robotic imagination on scene graphs," *arXiv:2309.12188*, 2023.
- [22] N. Gkanatsios, A. Jain, Z. Xian, Y. Zhang, C. Atkeson, and K. Fragkiadaki, "Energy-based models are zero-shot planners for compositional scene rearrangement," *Robotics: Science and Systems*, 2023.
- [23] O. Mees, A. Emek, J. Vertens, and W. Burgard, "Learning object placements for relational instructions by hallucinating scene representations," in *IEEE Int. Conf. on Robotics and Automation*, 2020, pp. 94–100.
- [24] A. Simeonov, et al., "Shelving, stacking, hanging: Relational pose diffusion for multi-modal rearrangement," in *Conf. on Robot Learning*, 2023, pp. 2030–2069.
- [25] W. Liu, T. Hermans, S. Chernova, and C. Paxton, "Strucdiffusion: Object-centric diffusion for semantic rearrangement of novel objects," *arXiv:2211.04604*, 2022.
- [26] Z. Yang, et al., "Compositional Diffusion-Based Continuous Constraint Solvers," in *Conf. on Robot Learning*, 2023.
- [27] R. Kartmann and T. Asfour, "Interactive and incremental learning of spatial object relations from human demonstrations," *Frontiers in Robotics and AI*, vol. 10, p. 1151303, 2023.
- [28] I. Kapelyukh, Y. Ren, I. Alzugaray, and E. Johns, "Dream2real: Zero-shot 3d object rearrangement with vision-language models," in *First Workshop on Vision-Language Models for Navigation and Manipulation at ICRA 2024*, 2024.
- [29] J. Wu, et al., "Tidybot: Personalized robot assistance with large language models," *Auton. Robot.*, vol. 47, no. 8, pp. 1087–1102, 2023.
- [30] I. Kapelyukh and E. Johns, "My house, my rules: Learning tidying preferences with graph neural networks," in *Conf. on Robot Learning*, 2022, pp. 740–749.
- [31] I. Kapelyukh and E. Johns, "Scenescore: Learning a cost function for object arrangement," in *Workshop on Learning Effective Abstractions for Planning*, 2023, *Conf. on Robot Learning*.
- [32] T. Welschhold, N. Abdo, C. Dornhege, and W. Burgard, "Combined task and action learning from human demonstrations for mobile manipulation applications," in *Int. Conf. on Intelligent Robots and Systems*. IEEE, 2019, pp. 4317–4324.
- [33] A. Radford, et al., "Learning transferable visual models from natural language supervision," in *Int. Conf. on Machine Learning*, 2021.
- [34] W. Goodwin, I. Havoutis, and I. Posner, "You only look at one: Category-level object representations for pose estimation from a single example," in *Conf. on Robot Learning*, 2023, pp. 1435–1445.
- [35] A. Hurst, et al., "Gpt-4o system card," *arXiv:2410.21276*, 2024.
- [36] N. Abdo, *Learning to solve manipulation tasks from non-expert users in human-centered environments*. Universität Freiburg, 2017.
- [37] C. Zhou, C. C. Loy, and B. Dai, "Extract free dense labels from clip," in *Europ. Conf. on Computer Vision*, 2022, pp. 696–712.
- [38] A. Kirillov, et al., "Segment anything," *arXiv:2304.02643*, 2023.
- [39] M. Oquab, et al., "Dinov2: Learning robust visual features without supervision," *arXiv:2304.07193*, 2023.
- [40] N. B. Figueroa Fernandez and A. Billard, "A physically-consistent bayesian non-parametric mixture model for dynamical system learning," in *Conf. on Robot Learning*, 2018.