

Metric Learning for Generalizing Spatial Relations to New Objects

Oier Mees

Nichola Abdo

Mladen Mazuran

Wolfram Burgard

Abstract—Human-centered environments are rich with a wide variety of spatial relations between everyday objects. For autonomous robots to operate effectively in such environments, they should be able to reason about these relations and generalize them to objects with different shapes and sizes. For example, having learned to place a toy inside a basket, a robot should be able to generalize this concept using a spoon and a cup. This requires a robot to have the flexibility to learn arbitrary relations in a lifelong manner, making it challenging for an expert to pre-program it with sufficient knowledge to do so beforehand. In this paper, we address the problem of learning spatial relations by introducing a novel method from the perspective of distance metric learning. Our approach enables a robot to reason about the similarity between pairwise spatial relations, thereby enabling it to use its previous knowledge when presented with a new relation to imitate. We show how this makes it possible to learn arbitrary spatial relations from non-expert users using a small number of examples and in an interactive manner. Our extensive evaluation with real-world data demonstrates the effectiveness of our method in reasoning about a continuous spectrum of spatial relations and generalizing them to new objects.

I. INTRODUCTION

Understanding spatial relations is a crucial faculty of autonomous robots operating in human-centered environments. We expect future service robots to undertake a variety of everyday tasks such as setting the table, tidying up, or assembling furniture. In this context, a robot should be able to reason about the best way to reproduce a spatial relation between two objects, e.g., by placing an item inside a drawer, or aligning two boxes side by side.

However, our everyday environments typically include a rich spectrum of potential spatial relations. For example, each user may have different preferences with respect to object arrangements, which requires robots to be flexible enough to handle arbitrary relations they have not encountered before. Similarly, robots should be able to generalize relations they have learned and achieve them using new objects of different shapes or sizes. For these reasons, it is highly impractical to expect an expert to pre-program a robot with the knowledge it needs to handle all potential situations in the real world, e.g., in the form of symbols. Instead, we aim for a lifelong learning approach that enables non-expert users to teach new spatial relations to robots in an intuitive manner.

One way to do this is to provide a robot with several examples using different objects in order to learn a model for a new relation, e.g., “inside.” On the one hand, this may require generating large amounts of data to learn the new



Fig. 1: We present a novel method based on distance metric learning to reason about the similarity between pairwise spatial relations. Our approach uses demonstrations of a relation given by non-expert teachers (top row) in order to generalize this relation to new objects (bottom row).

relation, which is impractical in setups in which a robot learns from a non-expert teacher. On the other hand, this requires learning a new model for each relation individually, making it hard for the robot to reuse its knowledge from previous relations.

In this paper, we address this problem from the perspective of distance metric learning and focus on learning relations between pairs of objects. We present a novel method that allows the robot to reason about how similar two relations are to each other. By doing so, we formulate the problem of reproducing a relation using two new objects as one of minimizing the distance between the reproduced relation and the teacher demonstrations. More importantly, our approach enables the robot to use a few teacher demonstrations as queries for retrieving *similar* relations it has seen before, thereby leveraging prior knowledge to bootstrap imitating the new relation. Therefore, rather than learning a finite set of individual relation models, our method enables reasoning on a continuous spectrum of relations.

Concretely, we make the following contributions: *i)* we present a novel approach from the perspective of distance metric learning to address the problem of learning pairwise spatial relations and generalizing them to new objects¹, *ii)* we introduce a novel descriptor that encodes pairwise spatial relations based only on the object geometries and their relative pose, *iii)* we demonstrate how our method enables bootstrapping the learning of a new relation by relating it to similar, previously-learned relations, *iv)* we present an interactive learning method that enables non-expert users

All authors are with the Department of Computer Science, University of Freiburg, Germany. {meeso, abdon, mazuran, burgard}@informatik.uni-freiburg.de

¹The Freiburg Spatial Relations Dataset and a demo video of our approach running on the PR-2 robot are available at <http://spatialrelations.cs.uni-freiburg.de>

to teach arbitrary spatial relations from a small number of examples, and v) we present an extensive evaluation of our method based on real-world data we gathered from different user demonstrations.

II. RELATED WORK

In the context of robotics, previous work has focused on leveraging predefined relations in the form of symbolic predicates for solving tasks, as in the case of combined task and motion planning or in the context of relational reinforcement learning [4, 8, 12, 22]. Rather than relying on grounding existing symbols, other works have addressed learning symbols and effects of actions to abstract continuous states for the purpose of high-level planning [1, 2, 10, 14, 20]. As opposed to these works, we reason about the similarity between relations by learning the distance between scenes, allowing us to compute scenes that generalize a relation to new objects.

Related to this is the work by Rosman and Ramamoorthy, which proposes constructing a contact point graph to classify spatial relations [23]. Similarly, Fichtl et al. train random forest classifiers for relations based on histograms that encode the relative position of surface patches [5]. Guadarrama et al. learn models of pre-defined prepositions by training a multi-class logistic regression model using data gathered from crowdsourcing [6]. As opposed to those works, we propose learning a distance metric that captures the similarities between different relations without specifying explicit classes.

Moreover, related to our work is the interactive approach by Kulick et al. for learning relational symbols from a teacher [16]. They use Gaussian Process classifiers to model symbols and therefore enable a robot to query the teacher with examples to increase its confidence in the learned models. Similarly, our method enables a robot to generalize a relation by interacting with a teacher. However, we do this from the perspective of metric learning, allowing the robot to re-use previous demonstrations of other relations.

Similar to our work, Zampogiannis et al. model spatial relations based on the geometries of objects given their point cloud models [27]. However, they define a variety of common relations and focus on addressing the problem of extracting the semantics of manipulation actions through temporal analysis of spatial relations between objects. Other methods have also relied on the geometries of objects and scenes to reason about preferred object placements [11] or likely places to find an object [3]. Moreover, Kroemer and Peters used 3D object models to extract contact point distributions for predicting interactions between objects [15].

Finally, our approach leverages distance metric learning for reasoning about the similarity between relations. Metric learning is a popular paradigm in the machine learning and computer vision communities. Learned metrics have been applied to address face recognition [7], image classification [9] and image segmentation [26]. In the context of robotics, metric learning has been used to address problems related to object instance or place recognition [17, 24].

III. NOTATION AND PROBLEM FORMULATION

In this section, we formalize the problem we address in this paper.

A. Object Representation

We consider the problem of learning spatial relations between pairs of objects. We denote an object by o . In this work, we assume to have no semantic knowledge about objects such as their type, e.g., box. Instead, we aim to learn relations based on object geometries and assume to have a 3D model of each object o_k in the form of a point cloud \mathbf{P}_k . We consider only points on the surface of the objects. We model the state using the 3D poses of objects in $SE(3)$ and express the pose \mathbf{T}_k of o_k relative to the world frame as a homogeneous transform consisting of a translation vector $\mathbf{t}_k \in \mathbb{R}^3$ and a rotation matrix $\mathbf{R}_k \in SO(3)$. We denote the pose of o_l relative to o_k by ${}^k\mathbf{T}_l$. Additionally, we assume the world frame to be specified such that the $-z$ -axis aligns with the gravity vector \mathbf{g} .

B. Pairwise Relations

We consider learning pairwise spatial relations between objects, i.e., we consider scenes with two objects only. In this work, we do not address the perception problem and rely on existing techniques to segment the scene and compute the object poses based on their point clouds. Accordingly, we express a scene with o_k and o_l as a tuple $\mathbf{s} := \langle \mathbf{P}_k, \mathbf{P}_l, {}^k\mathbf{T}_l \rangle$. In this work, we assume that one of the objects (o_k) is labeled as the *reference object*, and therefore express the scene using the pose ${}^k\mathbf{T}_l$ of o_l relative to o_k . Given a scene \mathbf{s} , we aim to express the spatial relation between the two objects in it. For this, we rely on a feature function (descriptor) f to express the relation as a K -dimensional feature vector \mathbf{r} , i.e., $f(\mathbf{s}) = \mathbf{r} \in \mathbb{R}^K$. Moreover, our goal is to enable the robot to reason about the how *similar* two scenes are with respect to the pairwise relations in them. We capture the similarity between \mathbf{s}_1 and \mathbf{s}_2 using a distance function dist that computes the distance $\text{dist}(\mathbf{r}_1, \mathbf{r}_2) \geq 0$ between the two scenes with respect to their feature vectors \mathbf{r}_1 and \mathbf{r}_2 . In this work we do not consider relations involving more than two objects, as they can be defined by combining pairwise relations. Finally, we do not explicitly treat object symmetries.

C. The Problem

The problem we address is threefold.

1) *Representing relations*: First, we seek a descriptor f that enables us to capture the underlying spatial relation in a scene based only on the geometries (point clouds) of the objects, their relative poses, and the direction of gravity \mathbf{g} .

2) *Learning the distance between relations*: Given f , we aim to learn a distance metric dist for computing the distance between two scenes. For this, we rely on training data $\mathcal{D} = \{\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(N)}\}$ consisting of N demonstrated scenes. Additionally, we assume to have a symmetric similarity matrix \mathbf{Y} of size $N \times N$ with unit diagonal values. The value $y_{i,j}$ in the i -th row and j -th column of \mathbf{Y} captures the

degree of similarity between scenes $\mathbf{s}^{(i)}$ and $\mathbf{s}^{(j)}$ in \mathcal{D} . In this work, we consider binary similarities $y \in \{0, 1\}$, such that 0 represents dissimilar relations and 1 means that the relations in both scenes are identical. Note that we do not assume \mathbf{Y} to be completely specified, i.e., some entries may be missing. Therefore, we aim for a method that can learn with partially-labeled data with respect to scene similarities. Given \mathcal{D} and \mathbf{Y} , our goal is to learn a *distance metric* dist that captures the distance between scenes \mathbf{s}_1 and \mathbf{s}_2 based on their features $\mathbf{r}_1 = f(\mathbf{s}_1)$ and $\mathbf{r}_2 = f(\mathbf{s}_2)$. We learn this metric such that $\text{dist}(\mathbf{r}_1, \mathbf{r}_2)$ is “small” if \mathbf{s}_1 and \mathbf{s}_2 represent similar relations, and “large” if they represent dissimilar relations.

3) *Generalizing a relation to new objects*: Given f and a distance metric dist , our goal is to learn a new, arbitrary relation from a teacher. We assume the teacher provides a small set of demonstrations $\mathcal{D}' = \{\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(N')}\}$ of the new relation, where $1 \leq N' \ll N$. Given two new objects o_k and o_l , the robot has to “imitate” the demonstrated relation in \mathcal{D}' by computing the pose ${}^k\mathbf{T}_l$ of o_k relative to o_l such that the resulting scene $\mathbf{s}^* = \langle \mathbf{P}_k, \mathbf{P}_l, {}^k\mathbf{T}_l \rangle$ is close to the demonstrations with respect to the corresponding features.

Concretely, we seek the best ${}^k\mathbf{T}_l^*$ to a problem of the form:

$$\begin{aligned} & \text{minimize} \quad \mathcal{L}(\mathcal{R}', \mathbf{r}^*) \\ & \text{subject to} \quad {}^k\mathbf{T}_l \in SE(3), \mathbf{s}^* \in \mathcal{S}_{\text{feas}}, \end{aligned} \quad (1)$$

where $\mathcal{R}' = \{\mathbf{r}^{(1)}, \dots, \mathbf{r}^{(N')}\}$ is the set of features for the demonstrated scenes \mathcal{D}' , $\mathbf{r}^* = f(\mathbf{s}^*)$ is the feature vector of the test scene \mathbf{s}^* , and \mathcal{L} is a loss function describing the distance between the demonstrations and the test scene based on dist . Additionally, $\mathcal{S}_{\text{feas}}$ denotes the set of physically feasible scenes. In this work, we focus on computing the desired pose between the two objects and do not consider the problem of manipulating the objects to achieve this pose. Therefore, we consider $\mathcal{S}_{\text{feas}}$ as the set of scenes in which o_k and o_l are not colliding.

IV. PROPOSED FEATURE REPRESENTATION

In this section, we present our model for f and propose a descriptor for modeling pairwise spatial relations, thereby addressing Sec. III-C.1. We model relations based only on the spatial interaction between their point clouds \mathbf{P}_k and \mathbf{P}_l given the direction of the gravity vector \mathbf{g} . Note that in this work, we do not address the correspondence problem between scenes and assume the teacher specifies the reference object. We rely on the directions of the vectors between the surface points of the objects as a signature of the underlying relation between them. Defining these directions purely based on a fixed (world) reference frame is sub-optimal as this results in a descriptor that is affected by translations and rotations of the scene. At the same time, computing a local reference frame using one of the objects (e.g., using PCA) introduces the challenge of ensuring consistency and reproducibility of the axes across different scenes. We address this problem by computing angles between direction vectors between points on both objects and the centroid of the reference object o_k , see Fig. 2. This is analogous to




Fig. 2: Visualization of the descriptor computation for the spatial relation between objects o_k and o_l . We consider the gravity vector \mathbf{g} at the centroid \mathbf{c}_k of the reference object o_k . We compute angles θ and φ based on direction vectors involving all surface points \mathbf{p}_k and \mathbf{p}_l on o_k and o_l , respectively.

methods for computing rotationally-invariant descriptors for 2D images such as RIFT [18]. Accordingly, we propose a descriptor that is based on three histograms as follows:

$$f(\mathbf{s}) := [\mathbf{h}_\theta \quad \mathbf{h}_\varphi \quad \mathbf{h}_d]^\top. \quad (2)$$

With the first histogram \mathbf{h}_θ , we capture the angular relation between the two objects regardless of how the scene is oriented in the global reference frame. We construct \mathbf{h}_θ as a distribution over the angle θ between vectors $(\mathbf{p}_k - \mathbf{c}_k)$ and $(\mathbf{p}_l - \mathbf{p}_k)$ based on all points $\mathbf{p}_k \in \mathbf{P}_k$ and $\mathbf{p}_l \in \mathbf{P}_l$, i.e.,

$$\theta = \arccos \left(\frac{(\mathbf{p}_k - \mathbf{c}_k)^\top (\mathbf{p}_l - \mathbf{p}_k)}{\|\mathbf{p}_k - \mathbf{c}_k\|_2 \|\mathbf{p}_l - \mathbf{p}_k\|_2} \right), \quad (3)$$

where \mathbf{c}_k is the centroid of the reference object o_k , see Fig. 2.

For the same relative pose ${}^k\mathbf{T}_l$, \mathbf{h}_θ provides a unique signature of the scene that is invariant to its translation or rotation in the global reference frame. However, in the context of everyday manipulation tasks, it is typically useful to also reason about spatial constraints with respect to the world frame, e.g., a supporting surface such as a table. For example, this enables the robot to disambiguate scenes in which the two objects are on top of each other from those in which they are next to each other for the same ${}^k\mathbf{T}_l$.

We achieve this disambiguation using the second histogram \mathbf{h}_φ , which is a distribution over the angle φ around the vector $(\mathbf{p}_k - \mathbf{c}_k)$, see Fig. 2. We take this as the angle between two planes. The first plane is defined by the two vectors $(\mathbf{p}_k - \mathbf{c}_k)$ and \mathbf{g} , whereas the second is defined by the two vectors $(\mathbf{p}_k - \mathbf{c}_k)$ and $(\mathbf{p}_l - \mathbf{p}_k)$. We compute φ as the angle between the respective normal vectors \mathbf{n}_1 and \mathbf{n}_2 of those planes, i.e.,

$$\begin{aligned} \mathbf{n}_1 &= \frac{(\mathbf{p}_k - \mathbf{c}_k) \times \mathbf{g}}{\|(\mathbf{p}_k - \mathbf{c}_k) \times \mathbf{g}\|_2}, \mathbf{n}_2 = \frac{(\mathbf{p}_k - \mathbf{c}_k) \times (\mathbf{p}_l - \mathbf{p}_k)}{\|(\mathbf{p}_k - \mathbf{c}_k) \times (\mathbf{p}_l - \mathbf{p}_k)\|_2}, \\ \varphi &= \arccos(\mathbf{n}_1^\top \mathbf{n}_2). \end{aligned} \quad (4)$$

We populate \mathbf{h}_φ by computing φ using all surface points $\mathbf{p}_k \in \mathbf{P}_k$ and $\mathbf{p}_l \in \mathbf{P}_l$. As the direction of \mathbf{g} is fixed, rotating the scene while maintaining ${}^k\mathbf{T}_l$ results in changes in φ , i.e., the discriminative behavior we seek. On the other hand, \mathbf{h}_θ and \mathbf{h}_φ are invariant to translations or rotations around \mathbf{g} .

Whereas \mathbf{h}_θ and \mathbf{h}_φ encode the relation with respect to the relative rotation between the two objects, we encode the desired distance between them using the third histogram \mathbf{h}_d . We compute \mathbf{h}_d as a distribution over the Euclidean distance $\|\mathbf{p}_k - \mathbf{p}_l\|_2$ between points $\mathbf{p}_k \in \mathbf{P}_k$ and $\mathbf{p}_l \in \mathbf{P}_l$. Rather than doing so using all $|\mathbf{P}_k| |\mathbf{P}_l|$ pairs of points $\{\mathbf{p}_k, \mathbf{p}_l\}$, we consider the subset of pairs with the smallest 10% distances over all pairs, as this is indicative of how close the two objects are and is less sensitive to differences in object sizes.

We discretize both \mathbf{h}_θ and \mathbf{h}_φ with a bin resolution of 20 deg, and discretize \mathbf{h}_d with a resolution of 6 cm. We normalize all histograms using the number of points used to compute them such that $f(s)$ is independent of the object point cloud densities and object sizes. The resulting descriptor $f(s) = \mathbf{r}$ has 39 dimensions. To speed up computation, besides considering only surface points, we also downsample the point clouds.

V. DISTANCE METRIC LEARNING

In this section, we discuss how we learn a metric dist that models the similarities between relations given the feature representation $\mathbf{r} = f(\mathbf{s})$ above (see Sec. III-C.2). For this, we leverage a popular metric learning technique originally introduced to improve the performance of k -NN classification: large margin nearest neighbor (LMNN) [25].

We follow the terminology of Weinberger and Saul and define the set of *target neighbors* \mathcal{R}_i^+ for an example \mathbf{r}_i as the k nearest neighbors of \mathbf{r}_i that are labeled as similar, i.e. $y_{i,j} = 1$ for $\mathbf{r}_j \in \mathcal{R}_i^+$. These target neighbors define a region around \mathbf{r}_i . We refer to all examples \mathbf{r}_k within this region that are not similar to \mathbf{r}_i (i.e., $y_{i,k} = 0$) as *imposters* \mathcal{R}_i^- . The original LMNN formulation identifies \mathcal{R}_i^+ and \mathcal{R}_i^- by assuming training data that is labeled with pre-specified classes. In our context, we achieve this based on the similarity labels y without requiring class labels to be specified by the teacher.

In the general form, LMNN learns a metric dist_ϕ parametrized by ϕ by minimizing a loss function with two objectives: *i)* for each training relation \mathbf{r}_i , pull target neighbors $\mathbf{r}_j \in \mathcal{R}_i^+$ close, and *ii)* push imposters $\mathbf{r}_k \in \mathcal{R}_i^-$ away such that they are further than target neighbors \mathbf{r}_j by at least a large margin ζ (see [25]), i.e.,

$$\begin{aligned} & \text{pull a similar neighbor } \mathbf{r}_j \text{ close} \\ \text{minimize}_{\phi} & \sum_{\substack{\mathbf{r}_i \in \mathcal{D}, \\ \mathbf{r}_j \in \mathcal{R}_i^+}} \overbrace{\text{dist}_\phi(\mathbf{r}_i, \mathbf{r}_j)^2}^{\text{pull a similar neighbor } \mathbf{r}_j \text{ close}} + \\ & \lambda \sum_{\mathbf{r}_k \in \mathcal{R}_i^-} \underbrace{[\zeta + \text{dist}_\phi(\mathbf{r}_i, \mathbf{r}_j)^2 - \text{dist}_\phi(\mathbf{r}_i, \mathbf{r}_k)^2]_+}_{\text{push a dissimilar neighbor } \mathbf{r}_k \text{ further than } \mathbf{r}_j \text{ by at least } \zeta}, \end{aligned} \quad (5)$$

where $[d]_+ = \max(0, d)$ is the hinge loss and λ is a constant that controls the trade-off between the two objectives.

In this work, we consider three LMNN-based methods for learning a dist_ϕ parametrized by ϕ . The linear LMNN case learns a generalized Euclidean (Mahalanobis) distance by parametrizing dist_ϕ using a linear mapping $\mathbf{L} \in \mathbb{R}^{K \times K}$, i.e., $\phi(\mathbf{r}) = \mathbf{L}\mathbf{r}$, see [25]. χ^2 -LMNN also learns a linear

mapping but uses the χ^2 distance instead of the Euclidean distance, see [13]. Finally, gradient-boosted LMNN (GB-LMNN) models arbitrary non-linear mappings $\phi(\mathbf{r})$ of the input space using gradient-boosted regression trees, see [13].

VI. REPRODUCING A NEW RELATION FROM A FEW DEMONSTRATIONS

In this section, we present our approach for imitating a new relation from a small number of teacher demonstrations (Sec. III-C.3). We assume that the robot is already equipped with a set of relation scenes $\mathcal{D} = \{\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(N)}\}$ and a (partially-filled) $N \times N$ matrix \mathbf{Y} consisting of their similarity labels as in Sec. III-C.2. These are either provided by an expert beforehand, or are accumulated by the robot when learning previous relations over time. Using \mathcal{D} and \mathbf{Y} , we assume the robot has already learned a prior distance metric dist_{ϕ_0} parametrized by ϕ_0 as described in Sec. V. This is done offline and without knowledge of the new relation.

We now consider a teacher providing the robot with a small set of demonstrations \mathcal{D}' of size N' for a new, arbitrary relation. The teacher can use different pairs of objects, such that all scenes in \mathcal{D}' are equally valid ways of achieving this relation, i.e., $y_{i,j} = 1$ for all $\mathbf{s}^{(i)}, \mathbf{s}^{(j)} \in \mathcal{D}'$. Given two objects o_k and o_l and their respective models \mathbf{P}_k and \mathbf{P}_l , our goal is to compute a pose ${}^k\mathbf{T}_l^*$ such that the resulting scene $\mathbf{s}^* = \langle \mathbf{P}_k, \mathbf{P}_l, {}^k\mathbf{T}_l^* \rangle$ corresponds to the intention of the teacher for the new relation, i.e., minimizing \mathcal{L} in Eq. (1).

Given a metric dist_{ϕ^*} , there are different ways to model \mathcal{L} to express the distance between the features $\mathbf{r}^* = f(\mathbf{s}^*)$ of the test scene and the features \mathcal{R}' of the demonstrations \mathcal{D}' . In general, as $|\mathcal{R}'| \geq 1$, Eq. (1) is a multi-objective optimization problem seeking to minimize the distance between \mathbf{r}^* and all $\mathbf{r}' \in \mathcal{R}'$. In such settings, it is typically challenging to satisfy all objectives. Minimizing the (mean) distance to \mathcal{D}' can thus lead to sub-optimal solutions that “average” the demonstrated scenes. Instead, we consider each demonstration to represent a mode of the target relation and seek the best solution with respect to any of them as follows:

$$\mathcal{L}(\mathcal{R}', \mathbf{r}^*) := \min_{\mathbf{r}' \in \mathcal{R}'} \text{dist}_{\phi^*}(\mathbf{r}', \mathbf{r}^*). \quad (6)$$

A. Interactive Local Metric Learning

One way to obtain dist_{ϕ^*} in Eq. (6) is to use the prior metric dist_{ϕ_0} . However, we learned this metric using a set of previous relations \mathcal{D} and their similarities. Therefore, it is not directly clear if dist_{ϕ_0} is able to generalize to novel relations shown by the teacher.

We answer this question using an interactive approach. For each demonstration in \mathcal{D}' , we use dist_{ϕ_0} to retrieve the Q nearest neighbor examples corresponding to scenes from the database \mathcal{D} . We query the teacher with these examples and ask her to indicate whether they align with her intention ($y = 1$) for the new relation or not ($y = 0$). In our experiments, we achieved this by means of a graphical user interface visualizing $Q = 8$ nearest neighbors per query. Let $\mathcal{D}_{NN} \subset \mathcal{D}$ be the set of all nearest neighbor scenes for \mathcal{D}' . We measure the confidence in the ability of dist_{ϕ_0} to

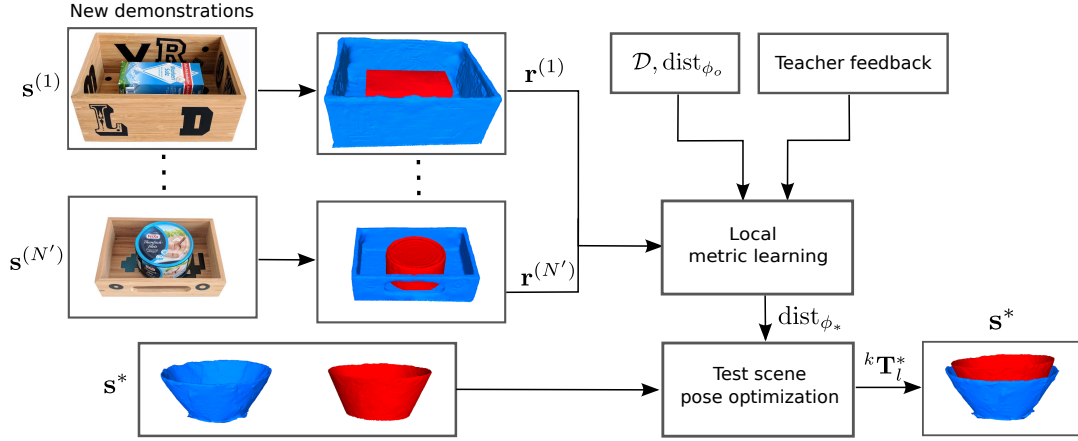


Fig. 3: Overview of our interactive approach for learning to reproduce a new relation. Given a small number of demonstrations $s^{(1)}, \dots, s^{(N')}$ by a teacher and two objects in the test scene s^* , we aim to compute a pose ${}^kT_l^*$ transforming s^* in order to imitate the demonstrated relation and generalize the intention of the teacher. Our approach enables the robot to leverage its prior knowledge of scenes \mathcal{D} representing other relations and the distances between them based on a metric dist_{ϕ_0} .

generalize to the new relation as the ratio ϵ_{NN} of scenes in \mathcal{D}_{NN} for which the teacher indicated a similarity to the new relation (i.e., $y = 1$).

ϵ_{NN} values larger than a threshold ϵ_* indicate that we are able to relate the new relation to ones the robot has seen in the past. Therefore, we use dist_{ϕ_0} to compute Eq. (6), i.e., $\phi_* = \phi_0$. We empirically set ϵ_* to 77% in our experiments. On the other hand, $\epsilon_{NN} < \epsilon_*$ indicates that the new relation is far in feature space from (target neighbor) relations in \mathcal{D} . We address this by learning a new *local* metric dist_{ϕ_*} using the set of scenes $\mathcal{D}^* = \mathcal{D}' \cup \mathcal{D}_{NN}$ and labels \mathbf{Y}^* of size $N^* \times N^*$, where $N^* = |\mathcal{D}^*|$. This is a smaller problem (compared to learning the prior metric dist_{ϕ_0}) in which \mathbf{Y}^* is completely specified. We set the similarity $y_{i,j}$ to 1 for all $s^{(i)}, s^{(j)} \in \mathcal{D}'$. For rows i and columns j corresponding to scenes $s^{(i)} \in \mathcal{D}'$ and $s^{(j)} \in \mathcal{D}_{NN}$ (or vice versa), we set $y_{i,j}$ to the similarity labels obtained from querying the teacher. We use the transitivity property to set the similarity between $s^{(j)}, s^{(k)} \in \mathcal{D}_{NN}$. For example, if the teacher labeled $y_{i,j} = 0$ and $y_{i,k} = 0$, we set $y_{j,k} = 1$.

Finally, we highlight two main advantages of leveraging the previous relations \mathcal{D} and prior metric dist_{ϕ_0} . Firstly, we enable the robot to consider whether its previous knowledge is sufficient to reproduce the new relation or not. Secondly, even for new relations that are significantly different from previously-known ones, we are able to augment the teacher's demonstrations \mathcal{D}' with additional training data \mathcal{D}_{NN} consisting of target neighbors and imposters retrieved from \mathcal{D} without requiring the teacher to demonstrate them.

B. Sample-Based Pose Optimization

Given the metric dist_{ϕ_*} to model Eq. (6), we present our approach for solving Eq. (1) to compute ${}^kT_l^*$ for reproducing the new relation using o_k and o_l . In this work, we simplify this problem by assuming that the reference object o_k is stationary and therefore only reason about desirable poses of o_l relative to it. Due to the discretization in computing

our descriptor f , we cannot rely on gradient-based methods as our loss function is piecewise constant.

We address this using a sample-based approach. We discretize the space of poses by searching over a grid of translations ${}^k\mathbf{t}_l$ of o_l relative to o_k . For each translation, we sample rotations ${}^k\mathbf{R}_l$ uniformly. We use the resulting ${}^k\mathbf{T}_l$ to transform \mathbf{P}_l and compute \mathcal{L} based on the corresponding feature value \mathbf{r}^* of the scene. Whenever we find a new local minima during optimization, we check for collisions between the two objects and reject infeasible solutions. Finally, we take ${}^k\mathbf{T}_l^*$ as the feasible pose minimizing \mathcal{L} over all sampled poses. We implemented this process efficiently by parallelizing the grid search over translations. Fig. 3 shows an overview of our method for reproducing a new relation.

VII. EXPERIMENTAL EVALUATION

In this section, we present the experimental evaluation of our approach. Through our experiments, we demonstrate the following: *i*) our proposed descriptor is able to capture different spatial relations and to generalize to the shapes and sizes of the objects, *ii*) using distance metric learning, we are able to capture the similarities between scenes even for relations not encountered before, *iii*) our interactive learning method enables non-expert users to teach new relations based on small number of examples, and *iv*) we outperform several baselines that do not learn a metric based on the similarities between scenes.

A. Baselines

In our experiments, we used three variants of LMNN-based metrics: vanilla (linear) LMNN, χ^2 -LMNN, and GB-LMNN, which we learned as in Sec. V. We compared those learned metrics to a variety of standard distance metrics. This includes the Euclidean, χ^2 , Bhattacharyya, and the correlation distances, as well as the Kullback-Leibler divergence (KL) and the Jensen-Shannon divergence (JS).

B. Dataset

We recorded 3D models of 26 household objects and used SimTrack to detect them and compute their poses in a scene using a Kinect camera [21]. Using this setup, we recorded a set of demonstrations \mathcal{D} consisting of 546 scenes, see Fig. 4 for examples. For the purpose of evaluation, we manually labeled the similarities \mathbf{Y} between all scenes.




Fig. 4: Examples of the training scenes we recorded for our evaluation. We recorded a set of 546 scenes and manually provided ground truth labels for their similarities.

C. Nearest Neighbor Classification

In this experiment, we evaluated the ability of distance metric learning to relate scenes based on the similarities of their relations. We formulated this as a k -NN classification problem, with $k = 5$, and evaluated using 15 random splits. For each split we used 75% of the data for the training set and 25% for the test set. We considered a success if at least 3 out of 5 of the retrieved nearest neighbors were similar to the test example.

Method	Accuracy(%)
Euclidean	82.32 ± 2.56
KL	82.61 ± 3.10
Correlation	82.66 ± 2.43
χ^2	82.81 ± 3.20
Bhattacharyya	83.26 ± 3.15
JS	83.30 ± 3.14
χ^2 -LMNN	86.46 ± 2.84
LMNN	86.52 ± 1.98
GB-LMNN	87.60 ± 1.94

TABLE I: Performance of different methods for retrieving at least 3 out of 5 target neighbors of scenes, averaged over 15 random splits.

The results are shown in Tab. I. LMNN-based metrics outperform the baselines, i.e., the learned metrics can better capture the distances between scenes. We achieved the highest success rate of 87.6% using GB-LMNN. Note that by directly computing the Euclidean distance in the original feature space, we are able to achieve a success rate of 82.32%. This demonstrates that our proposed feature descriptor is suitable for encoding arbitrary spatial relations. Fig. 5 shows a qualitative example of the nearest neighbors of a test scene using both the Euclidean distance and LMNN.




Fig. 5: Example of the 5 nearest neighbors of a query scene using Euclidean distance (top) and LMNN (bottom). LMNN better captures the distances between relations.




Fig. 6: t-SNE visualization of six relations mapped using LMNN. We trained this metric using three of the relations only, yet it is able to capture the semantic similarity between those relations and the test relations.

D. Distance to New Relations

In this qualitative experiment, we investigated the ability of a learned metric to capture the similarities between relations that were not used for training. We trained LMNN with data from three relations (upper row of Fig. 4), which can be semantically described as “on top”, “inside”, and “next to”. We used the learned metric to map all six relations in Fig. 4 to the new space and visualized the data using t-SNE, a popular non-linear embedding technique for visualizing high dimensional data [19]. We show this in Fig. 6. This qualitatively illustrates the separation between the three relations used for training the metric. Moreover, the metric is able to capture the semantic similarity between the relations used for training and the new ones, which we denote by “inclined”, “on top corner” and “inclined inside” (bottom row of Fig. 4).

E. Generalizing a Relation to New Objects

In this experiment, we evaluated our approach for reproducing a demonstrated relation using two new objects, see Sec. III-C.3. We recorded 30 demonstrations for each of five new relations. We then selected two new objects that were not used in the demonstrations and evaluated our




Fig. 7: Results for generalizing a relation to two new objects. GB-LMNN and LMNN outperform the other metrics in identifying the correct ways of reproducing a relation with two new objects. Our approach enables leveraging previous demonstrations when querying the teacher to learn a new relation. In each case, we evaluated the generalization given five new demonstrations, which we then added to the set of prior demonstrations \mathcal{D} . The x-axis shows the cumulative number of demonstrations over time.

method’s ability to generalize each of the five relations to the new objects. In each case, we provided our method with $|\mathcal{D}'| = 5$ examples. Using those examples, we retrieved nearest neighbor queries \mathcal{D}_{NN} from \mathcal{D} as described in Sec. VI. As we aimed for a quantitative evaluation in this experiment, we implemented an “oracle” that provides the binary labels for the queries automatically, and used this to learn a local metric as in Sec. VI-A.

For evaluation, we provided our method with a set of 75 poses ${}^k\mathbf{T}_l$ between the new objects. Only 15 of them are correct ways of imitating the relation in question, whereas the rest represent other relations. For each pose, we computed \mathbf{r}^* of the test scene and the corresponding \mathcal{L} (Eq. (6)). We sorted the poses according to \mathcal{L} . Ideally, the correct poses should be in the top 15 positions. We evaluated this using the mean average precision.

After each such test, we added the five demonstrations \mathcal{D}' to the database \mathcal{D} and extended \mathbf{Y} with the new labels from the nearest neighbor queries. We used this to re-learn the prior metric dist_{ϕ_0} . We did this six times, each with five new demonstrations, until the 30 demonstrations have been used. We repeated the whole experiment 50 times using different random orders of providing five demonstrations. We report the results in Fig. 7 averaged over all runs and five relations.

The metrics we learned with GB-LMNN and LMNN outperform the other metrics in their ability to identify the correct ways of reproducing a relation with two new objects. Although in each case we compute Eq. (6) based on five new demonstrations only, our approach enables those metrics to re-use demonstrations added to \mathcal{D} from the previous tests to learn the local metric. Accordingly, GB-LMNN achieves a mean average precision of 0.82 after having seen at least five demonstrations in the past. This demonstrates the ability of our approach to use previous demonstrations when generalizing a new relation in a lifelong learning manner.



Fig. 8: Left: three demonstrations of a relation by a teacher in our survey. Right: the generalized relation with two new objects using our approach.

F. Interactive Learning of a New Relation

We conducted a small survey to evaluate our approach for learning a new relation interactively. We recorded 250 scenes consisting of 50 relations demonstrated by nine different teachers. Each teacher provided five demonstrations per relation using different objects they chose. We used only three of those demonstrations to learn each relation. In each case, we queried the teacher with nearest neighbor examples from \mathcal{D} and used the result to generalize the three demonstrations \mathcal{D}' and reproduce one of the scenes we left out from the training as in Sec. VI. We computed the best pose for reproducing the relation using the sample-based approach in Sec. VI-B.

We showed the reproduced relations to their corresponding teachers in a 3D visualization environment and asked them to score the quality of the result with 0, 0.5, or 1, where 0 represents unsuccessful and 1 represents successful. The teachers scored results from different baselines shown in random order.

We show the mean scores in Tab. II. The reproduced scenes using both LMNN and GB-LMNN metrics were judged to be the best by the teachers, achieving an average score of 0.72 and 0.71 respectively. Fig. 8 illustrates one such generalization from our experiments. The results confirm that our approach enables non-expert users to teach arbitrary spatial relations to a robot from a small number of examples.

Method	Score
Euclidean	0.49
Bhattacharyya	0.54
KL	0.56
JS	0.56
χ^2	0.60
χ^2 -LMNN	0.63
LMNN	0.72
GB-LMNN	0.71

TABLE II: Mean scores for reproducing 50 relations from nine different teachers. For each relation, the teacher scored the result with 0 (unsuccessful), 0.5, or 1 (successful).

As a final qualitative evaluation, we carried out a small survey in which we asked six participants to judge the quality of scenes generated by our sample-based approach in Sec. VI-B. For this, we selected 30 scenes computed by LMNN, and which were scored with either 0.5 or 1 in the experiment above. We asked the six participants to judge whether the relations in those scenes were demonstrated by a human or generated by a computer (whereas in fact they were all computed using our approach). Despite the fact that we do not consider physics checks (e.g., scene stability) or optimize

computed scenes to make them more realistic, 63.05% of the scenes were thought to have been produced by a human.

VIII. CONCLUSION

In this paper, we presented a novel approach to the problem of learning pairwise spatial relations and generalizing them to different objects. Our method is based on distance metric learning and enables a robot to reason about the similarity between scenes with respect to the relations they represent. To encode a relation, we introduced a novel descriptor based on the geometries of the objects. By learning a distance metric using this representation, our method is able to reproduce a new relation from a small number of teacher demonstrations by reasoning about its similarity to previously-encountered ones. In this way, our approach allows for a lifelong learning scenario by continuously leveraging its prior knowledge about relations to bootstrap imitating new ones. We evaluated our approach extensively using real-world data we gathered from non-expert teachers. Our results demonstrate the effectiveness of our approach in reasoning about the similarity between relations and its ability to reproduce arbitrary relations to new objects by learning interactively from a teacher. In the future, we plan to extend the approach to partial observations of point cloud data.

ACKNOWLEDGMENTS

This work has partly been supported by the German Research Foundation under research unit FOR 1513 (HY-BRIS) and grant number EXC 1086. We thank Gian Diego Tipaldi, Marc Toussaint, and Oliver Kroemer for the valuable discussions during the early stages of this work.

REFERENCES

- [1] N. Abdo, H. Kretschmar, L. Spinello, and C. Stachniss. Learning manipulation actions from a few demonstrations. In *Int. Conf. on Robotics & Automation (ICRA)*, 2013.
- [2] S. R. Ahmadzadeh, A. Paikan, F. Mastrogiovanni, L. Natale, P. Kormushev, and D. G. Caldwell. Learning symbolic representations of actions from human demonstrations. In *Robotics and Automation (ICRA)*, 2015 *IEEE International Conference on*, pages 3801–3808. IEEE, 2015.
- [3] A. Aydemir and P. Jensfelt. Exploiting and modeling local 3d structure for predicting object locations. In *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2012.
- [4] M. Beetz, L. Mösenlechner, and M. Tenorth. Cram-A cognitive robot abstract machine for everyday manipulation in human environments. In *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2010.
- [5] S. Fichtl, A. McManus, W. Mustafa, D. Kraft, N. Krüger, and F. Guerin. Learning spatial relationships from 3d vision using histograms. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 501–508. IEEE, 2014.
- [6] S. Guadarrama, L. Riano, D. Golland, D. Go, Y. Jia, D. Klein, P. Abbeel, T. Darrell, et al. Grounding spatial relations for human-robot interaction. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1640–1647. IEEE, 2013.
- [7] M. Guillaumin, J. Verbeek, and C. Schmid. Is that you? metric learning approaches for face identification. In *2009 IEEE 12th International Conference on Computer Vision*, pages 498–505. IEEE, 2009.
- [8] S. Höfer and O. Brock. Coupled learning of action parameters and forward models for manipulation. In *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2016.
- [9] S. C. Hoi, W. Liu, and S.-F. Chang. Semi-supervised distance metric learning for collaborative image retrieval. In *Computer Vision and Pattern Recognition*, 2008. *CVPR 2008. IEEE Conference on*, pages 1–7. IEEE, 2008.
- [10] N. Jetchev, T. Lang, and M. Toussaint. Learning grounded relational symbols from continuous data for abstract reasoning. *ICRA Workshop on Autonomous Learning*, 2013.
- [11] Y. Jiang, M. Lim, C. Zheng, and A. Saxena. Learning to place new objects in a scene. *Int. J. of Robotics Research (IJRR)*, 2012.
- [12] L. P. Kaelbling and T. Lozano-Pérez. Integrated task and motion planning in belief space. *Int. J. of Robotics Research (IJRR)*, 32(9-10): 1194–1227, 2013.
- [13] D. Kedem, S. Tyree, F. Sha, G. R. Lanckriet, and K. Q. Weinberger. Non-linear metric learning. In *Advances in Neural Information Processing Systems*, pages 2573–2581, 2012.
- [14] G. Konidaris, L. P. Kaelbling, and T. Lozano-Perez. Constructing symbolic representations for high-level planning. In *National Conf. on Artificial Intelligence (AAAI)*, 2014.
- [15] O. Kroemer and J. Peters. Predicting object interactions from contact distributions. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3361–3367. IEEE, 2014.
- [16] J. Kulick, M. Toussaint, T. Lang, and M. Lopes. Active learning for teaching a robot grounded relational symbols. In *IJCAI*, 2013.
- [17] K. Lai, L. Bo, X. Ren, and D. Fox. Sparse distance learning for object recognition combining rgb and depth information. In *Robotics and Automation (ICRA)*, 2011 *IEEE International Conference on*, pages 4007–4013. IEEE, 2011.
- [18] S. Lazebnik, C. Schmid, and J. Ponce. A sparse texture representation using local affine regions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1265–1278, 2005.
- [19] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.
- [20] H. M. Pasula, L. S. Zettlemoyer, and L. P. Kaelbling. Learning symbolic models of stochastic domains. *Journal of Artificial Intelligence Research*, 29:309–352, 2007.
- [21] K. Pauwels and D. Kragic. Simtrack: A simulation-based framework for scalable real-time object pose detection and tracking. In *Intelligent Robots and Systems (IROS)*, 2015 *IEEE/RSJ International Conference on*, pages 1300–1307. IEEE, 2015.
- [22] C. Paxton, F. Jonathan, M. Kobilarov, and G. D. Hager. Do what i want, not what i did: Imitation of skills by planning sequences of actions. In *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2016.
- [23] B. Rosman and S. Ramamoorthy. Learning spatial relationships between objects. *The International Journal of Robotics Research*, 30(11):1328–1342, 2011.
- [24] M. Shahid, T. Naseer, and W. Burgard. Dtlc: Deeply trained loop closure detections for lifelong visual slam. In *Visual Place Recognition: What is it Good For? Workshop at the Robotics Science and Systems (RSS)*, June. 2016.
- [25] K. Q. Weinberger and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research*, 10(Feb):207–244, 2009.
- [26] S. Xiang, F. Nie, and C. Zhang. Learning a mahalanobis distance metric for data clustering and classification. *Pattern Recognition*, 41(12):3600–3612, 2008.
- [27] K. Zampogiannis, Y. Yang, C. Fermüller, and Y. Aloimonos. Learning the spatial semantics of manipulation actions through preposition grounding. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1389–1396. IEEE, 2015.