Learning Mobile Manipulation Actions from Human Demonstrations

Tim Welschehold Christian Dornhege Wolfram Burgard

***Abstract*— Over the past years learning from demonstration has become a popular method to intuitively teach new skills to service robots without explicit programming. However, most teaching approaches in literature use kinesthetic training and do not include mobile platforms. Here, we present a novel approach to learn joint robot base and gripper action models from observing demonstrations carried out by a human teacher. To achieve this we adapt RGBD observations of the human teacher to the capabilities of the robot. We formulate a graph optimization problem that links observations with robot grasping capabilities and kinematic constraints between co-occurring base and gripper poses. In real world experiments we show that the robot is able to learn complex mobile manipulation tasks such as opening and driving through a door.**

I. Introduction

The expected rise of mobile manipulation platforms to service robots acting in households in the near future poses the question of how they should be instructed in their custom environments. Ideally the teaching process should be intuitive and not require expert knowledge neither in programming nor about the robot kinematic and grasping capabilities. Instead of explicit programming or performing kinesthetic teaching the robot should learn actions directly from observing human demonstrations like depicted in Fig. 1. In our previous work we introduced a system that learns manipulation actions from human demonstrations [1]. Here, we focused on manipulation actions, where the base of the robot is at a fixed position. The ability to learn and reproduce actions that require a motion of the platform's base is one extension presented in this paper.

One of the major challenges in this context lies in the fact that the demonstrations provided by the human teacher typically cannot be exactly executed by the robot, due to different kinematic constraints between the base and gripper of the robot compared to human's torso and hand. In addition the grasping capabilities differ due to geometric differences between a human hand and the gripper of the robot. Grasp planners can generate high-quality grasp but lack a connection to the task at hand. I.e. for most tasks an object may need to be grasp in the way demonstrated by the teacher instead of just maximizing the grasp quality. So one needs to trade off the grasp quality against the resemblance to the demonstrations. Further problems arise due to the limited observability of the scene from one static observation angle. This makes the recording prone to measurement errors and even occlusions where the robot might temporarily loose track of the human or the object. In this paper we present an extended approach for teaching robots mobile manipulation actions by solely observing a human teacher. We enhance our previously stated approach that, based on graph optimization, generates consistent hand-object trajectories for manipulation task learning. In this paper the following contributions are made. We adapt the observed human grasp to the robot capabilities as part of the graph optimization. Here, the robot gripper motion is determined to both follow the human demonstration while fulfilling the geometric demands of the task to learn. In addition, we add functionality to learn the motion of the robot base by taking into account observed human torso poses. The motion of the robot base is constrained by the demonstrated human torso poses as well as the kinematic feasibility regarding the co-occurring gripper poses.

All authors are with the Institute of Computer Science, University of Freiburg, Germany. This work has been supported by the Baden-Württemberg Stiftung and the German Research Foundation under research unit BU 865/8-1 (HYBRIS)

Fig. 1: Learning from human demonstrations: The robot observes the human demonstrating a task, adapts this to its own capabilities and executes the task on its own.

We evaluated our approach in both simulation and real-world experiments on a PR2 on challenging tasks demanding automatic positioning as well as task specific motion of the robot base. The experiments demonstrate that our approach is able to learn a combined model of robot gripper and base motions for complex tasks from a low number of human demonstrations observed with an RGBD camera.

II. Related Work

Most learning from demonstration approaches do not deal with human demonstrations of object manipulation tasks and also do not consider mobile manipulation. Often demonstrations are performed by kinesthetic teaching [2], [3]. Pastor *et al.* [4] collect learning data with an exoskeleton arm and define grasp points on the handled objects manually. Like us, Calinon *et al.* [5] track human hand trajectories with markers, but they do not address manipulation tasks. The system presented by Mühlig *et al.* [6] learns from a human teachers using an articulated models of the human body. Although dealing with manipulation actions, grasp poses on the objects are either preprogrammed or assumed as given. None of these approaches is designed to handle a motion of the robot base. Burget *et al.* implement a imitation of human whole-body motions [7]. The focus lies on a stable real-time teleoperation of the robot and not on learning action models. Ratner *et al.* [8] present a web-based infrastructure for recording user demonstrations of mobile manipulation tasks on a simulated robot. This approach deals with mobile manipulation tasks but does not use human task demonstrations. Whole-body motion of humanoids or manipulation of articulated objects are also addressed as a planning problem in [9], [10], the latter also dealing with the task of opening and navigating through a door. In contrast to our work detailed information about the handled objects are needed to set planning constraints.

One of our goals is to allow the robot to position itself according to the task to handle. To this end we use the concept of an inverse reachability map [11]. In contrast to our work they do not address a joint motion of the robot base and gripper. A different idea to this goal is pursued by Stulp *et al.* [12]. They aim at learning to position the robot to perform manipulation tasks through trial-and-error interaction with the environment.

On the field of grasp planning there are numerous approaches with focus on detection of stable grasps of objects we can build on. Fischinger *et al.* propose a method to calculate grasps on point clouds using height accumulated features [13]. Ten Pas *et al.* [14] introduce an approach to sample and classify grasps on unknown objects in cluttered scenes using RGBD data. Other approaches focus on human inspired grasps [15]–[17]. Gioioso *et al.* [15] exploit synergetic composition of the human hand mapping the human hand to different robot grippers. In the grasp planning approach presented by Lin and Sun [16] human demonstrated grasp types are used as foundation. The work of Armor *et al.* [17] proposes a direct mapping of human hand to robot joint angles. A generalization to unknown objects is achieved by adapting the demonstrated contact points. The latter two methods use precise hand postures for training. All of them have a focus on grasping itself and not on manipulating objects. Our work builds upon such techniques since we use pre-sampled grasps and aim at choosing a grasp that balances grasp quality and the resemblance to the demonstrated manipulation task.

III. Problem Statement

For our problem setting a demonstration of a manipulation action consists of three 6-dof trajectories: one for the teacher's hand, one for the teacher's torso and one for the manipulated object. A trajectory $\boldsymbol{X}$ is a sequence of n poses $\boldsymbol{X} = \langle \boldsymbol{x}_1, \dots, \boldsymbol{x}_n \rangle$. These three independent trajectories are not required to be the same length and may contain gaps due to occlusions resulting in missing observations. A pose $\boldsymbol{x}$ is build of a translational part $pos(\boldsymbol{x})$ given by a 3d vector $(x, y, z)^T$ and a rotational part $rot(\boldsymbol{x})$ given by a quaternion $(q_x, q_y, q_z, q_w)^T$. Each pose $\boldsymbol{x}$ has a time stamp $time(\boldsymbol{x})$ allowing us to associate matching poses. We assume that all demonstrations can be segmented into reaching, manipulation and retreat parts. Manipulation is identified by a joint hand and object motion. Since our robot is incapable of in-hand manipulation we expect a fixed grasp. For more elaborate grippers this assumption might be relaxed.

We neither require expert demonstrations, in which the human hand trajectory precisely leads to a suitable grasp for the robot nor a given transform between the human hand and the robot gripper. Our goal is to adapt human demonstrations in a way that it allows the robot to learn a joint motion model for both the robot gripper and base enabling it to reproduce the intended action. We achieve this by formulating a hypergraph optimization problem.

IV. Graph Structure

In this section we show how to generate data suited for robot motion learning given a set of human teacher demonstrations. The observations are recorded with a RGBD camera. Note that these observations are subject to noise from the detection algorithms and might contain gaps where either the hand, the torso or the object are obstructed from the camera view. The recorded trajectories are preprocessed and segmented as described in our previous work [1]. We first compute matching same-length trajectories $\boldsymbol{X}_{hand}$, $\boldsymbol{X}_{torso}$ and $\boldsymbol{X}_{object}$. Whenever for a pose with time t in one of the trajectories there is no matching one with the same time in the others, we create a new pose with the missing time t in the other trajectories by linear interpolation for the translational part and slerp for the rotational part. Afterwards for each $\boldsymbol{x}_i^h \in \boldsymbol{X}_{hand}$ there is a respective $\boldsymbol{x}_i^o \in \boldsymbol{X}_{object}$ and $\boldsymbol{x}_i^t \in \boldsymbol{X}_{torso}$ with $time(\boldsymbol{x}_i^h) = time(\boldsymbol{x}_i^o) = time(\boldsymbol{x}_i^t)$.

Next we build a graph in which the vertices correspond to the recorded trajectories $\boldsymbol{X}_{hand}, \boldsymbol{X}_{torso}, \boldsymbol{X}_{object}$. Additionally we add nodes representing trajectories for the robot gripper $\boldsymbol{X}_{gripper}$ and the robot base $\boldsymbol{X}_{base}$. These are initialized at the same poses as the hand, respectively torso poses. We refer to the estimates of the vertices, which are initialized with the recorded poses, as $\hat{\boldsymbol{x}}_i^h$, $\hat{\boldsymbol{x}}_i^t$, $\hat{\boldsymbol{x}}_i^o$, $\hat{\boldsymbol{x}}_i^g$, $\hat{\boldsymbol{x}}_i^b$ for the hand, torso, object, gripper and base trajectories. These vertices are connected via different types of edges addressing assumptions and requirements we pose on the motions. Each k-edge relates its vertices' estimates $\hat{\boldsymbol{x}}_1, \dots, \hat{\boldsymbol{x}}_k$ with a measurement $z_{1,\dots,k}$ given for the edge by an error function $\mathbf{e} \stackrel{\text{def}}{=} \mathbf{e}(\hat{\boldsymbol{x}}_1, \dots, \hat{\boldsymbol{x}}_k, z_{1,\dots,k})$. The sum of squared error $\mathbf{e}^T \boldsymbol{\Omega} \mathbf{e}$ is minimized weighted with information matrices $\boldsymbol{\Omega}$ for all edges using g^{2}o [18]. Thus, for the optimization, besides the structure of the graph, we only need to define the error functions and the information matrices.

We attach unary edges to all vertices from measured (non-interpolated) poses. The poses $\boldsymbol{x}_i$ from the trajectories serve as the measurements for the estimates $\hat{\boldsymbol{x}}_i$. Defining the error function as

$$\mathbf{e}(\hat{\boldsymbol{x}}_i, \boldsymbol{x}_i) = \boldsymbol{x}_i^{-1} \cdot \hat{\boldsymbol{x}}_i \tag{1}$$

keeps the estimates near the actual observations. The information matrix is set to represent the accuracy of the corresponding tracking algorithm. To model that trajectories are smooth in the sense that consecutive poses are near to each other we add binary edges E_s along the individual

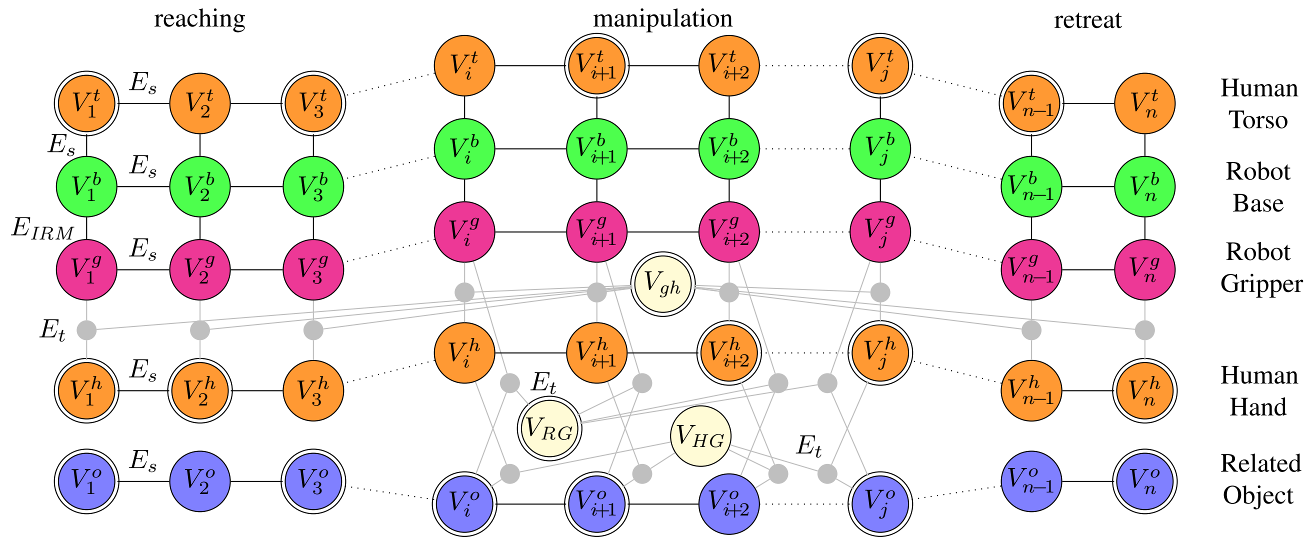

Fig. 2: A hyper-graph is constructed to compute the best estimate from the observations. Vertices for poses from the hand V^h and the torso V^t are shown in orange, vertices for object poses V^o are shown in blue and vertices for the robot gripper V^g and base V^b are shown in magenta and green, respectively. Vertices with unary edges that link matched poses to a fixed frame are shown with double circles. Binary edges representing the smoothness of trajectories are shown as black lines between vertices. Three lines connected by a dot display ternary edges that estimate a single transformation. These transformations represent the human grasp V_{HG}, the robot grasp V_{RG} and the transform between gripper and hand V_{gh}.

trajectories. For the hand, torso and object trajectories they keep the relative transformation between poses similar to that of the input trajectories' poses $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$ by

$$\mathbf{e}(\hat{\boldsymbol{x}}_i, \hat{\boldsymbol{x}}_j, \boldsymbol{x}_i, \boldsymbol{x}_j) = \left(\boldsymbol{x}_i^{-1} \cdot \boldsymbol{x}_j\right)^{-1} \cdot \left(\hat{\boldsymbol{x}}_i^{-1} \cdot \hat{\boldsymbol{x}}_j\right). \quad (2)$$

For the robot gripper and base poses we simply want consecutive poses to change as little as possible. We define the error function as

$$\mathbf{e}(\hat{\boldsymbol{x}}_i, \hat{\boldsymbol{x}}_j) = \hat{\boldsymbol{x}}_i^{-1} \cdot \hat{\boldsymbol{x}}_j. \quad (3)$$

Corresponding information matrices reflect the smoothness we assume for demonstrations scaled by the distance between the poses. Further used nodes and edges are described in the following. The structure of the graph with all modeled relations is presented in Fig. 2.

A. Consistent Hand-Object Trajectories

We use the structure presented in our previous work [1] to consistently correct measurement errors and occlusions in the human hand $\boldsymbol{X}_{hand}$ and object $\boldsymbol{X}_{object}$ motion in the manipulation segment. Under the assumption that during manipulation the relative transform between hand and object is fixed, i.e., we have a fixed (human) grasp, we insert an additional vertex V_{HG} into the graph that estimates this grasp transformation. The grasp gives us an estimate for the pose of the hand in the object's frame $\hat{\boldsymbol{x}}_h^o$. For each pair of poses $(\boldsymbol{x}_i^h, \boldsymbol{x}_i^o)$ in the manipulation segment we add a ternary edge E_t that links their estimates $\hat{\boldsymbol{x}}_i^h$, $\hat{\boldsymbol{x}}_i^o$ to $\hat{\boldsymbol{x}}_h^o$. The corresponding error function is

$$\mathbf{e}(\hat{\boldsymbol{x}}_i^h, \hat{\boldsymbol{x}}_i^o, \hat{\boldsymbol{x}}_h^o) = \hat{\boldsymbol{x}}_h^{o^{-1}} \cdot \left(\hat{\boldsymbol{x}}_i^{o^{-1}} \cdot \hat{\boldsymbol{x}}_i^h\right) \quad (4)$$

During the optimization this minimizes the difference between the grasp estimate $\hat{\boldsymbol{x}}_h^o$ and the relative transformation between all pairs of pose estimates for the hand and object.

By introducing V_{HG} and using ternary edges the human grasp is estimated as part of the optimization process.

B. Human Hand to Robot Gripper

On top of the described correction of the measured hand and object poses we generate a feasible robot gripper trajectory. It should follow the demonstrated hand motion as close as possible while ensuring a fixed and geometrically valid robot grasp on the manipulated object. To achieve this we use a similar approach as described above. We insert a vertex V_{RG}, with pose estimate $\hat{\boldsymbol{x}}_g^o$, that is connected to the pairs of gripper $\hat{\boldsymbol{x}}_i^g$ and object $\hat{\boldsymbol{x}}_i^o$ pose estimates via ternary edges. Using an error function similar to equation (4) we minimize the difference between all relative transforms from gripper to object pose estimates and the estimate of the robot grasp.

$$\mathbf{e}(\hat{\boldsymbol{x}}_i^g, \hat{\boldsymbol{x}}_i^o, \hat{\boldsymbol{x}}_g^o) = \hat{\boldsymbol{x}}_g^{o^{-1}} \cdot \left(\hat{\boldsymbol{x}}_i^{o^{-1}} \cdot \hat{\boldsymbol{x}}_i^g\right) \quad (5)$$

For this to also result in a good grasp for the robot we add a unary edge on V_{RG} addressing the grasp quality. We generate a set of possible grasp poses G with corresponding quality on the handled object beforehand [13]. We define the associated error function as

$$\mathbf{e}(\hat{\boldsymbol{x}}_g^o, G) = \hat{\boldsymbol{x}}_g^{o^{-1}} \cdot g \quad (6)$$

where g is the grasp closest to $\hat{\boldsymbol{x}}_g^o$ with quality above a certain threshold. To ensure similarity to the demonstration we add another vertex V_{gh} representing the transform between human hand and robot gripper poses, again following the same pattern as before. All pairs of gripper and hand poses are connected to V_{gh} with the goal to keep this transform constant over the action. Again an error function like equations (4, 5) is used.

$$\mathbf{e}(\hat{\boldsymbol{x}}_i^g, \hat{\boldsymbol{x}}_i^h, \hat{\boldsymbol{x}}_g^h) = \hat{\boldsymbol{x}}_g^{h^{-1}} \cdot \left(\hat{\boldsymbol{x}}_i^{h^{-1}} \cdot \hat{\boldsymbol{x}}_i^g\right) \quad (7)$$

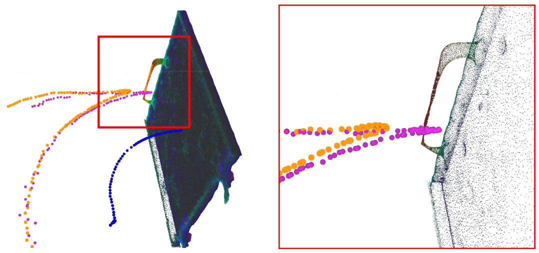

Fig. 3: Adaption of the grasp. The poses for the gripper (magenta dots) are shifted towards the handle of the door leading to a successful robot grasp. By just imitating the human hand motion (orange dots) the grasp would fail.

To keep the similarity as good as possible we add a unary edge on V_{gh} with error function

$$\mathbf{e}(\hat{\boldsymbol{x}}_g^h) = \hat{\boldsymbol{x}}_g^{h^{-1}}, \tag{8}$$

which pulls the estimate for the relative transform between hand and gripper $\hat{\boldsymbol{x}}_g^h$ towards identity. The weighting between resemblance to the demonstrations and grasp quality is done with the respective information matrices on the unary edges of V_{RG} and V_{gh}. An example of the effect of this setup is shown in Fig. 3.

C. Including Base Poses in the Graph

One drawback of our previous approach was the necessity to manual position of the robot base and the restriction that only actions that do not require base motions could be performed. In this new extended approach we also learn how to position and move the base while executing a task. To this end we introduce the nodes for the robot base V_i^b in the graph. These are initialized with the same estimates $\hat{\boldsymbol{x}}_i^b$ as the measured human torso pose estimates $\hat{\boldsymbol{x}}_i^t$ with the restriction that they need to be upright. These estimates are connected with binary edgeswith error function as equation 3. To ensure that the base poses form kinematically legal pairs with the corresponding gripper poses we use the information of an inverse reachability map (IRM) [11]. For a given gripper pose the IRM provides all feasible poses that the base may take to reach the desired gripper pose. The resolution of the used IRM was $5\,\mathrm{cm}$ and $45°$, respectively. We include this constraint in our graph by adding edges E_{IRM} between base V_i^b and gripper V_i^g vertices with error function

$$\mathbf{e}(\hat{\boldsymbol{x}}_i^g, \hat{\boldsymbol{x}}_i^b) = (\hat{\boldsymbol{x}}_i^{g^{-1}} \cdot \hat{\boldsymbol{x}}_i^b)^{-1} \cdot \boldsymbol{x}_{IRM} \tag{9}$$

where $\boldsymbol{x}_{IRM}$ is the closest pose in the IRM to $(\hat{\boldsymbol{x}}_i^{g^{-1}} \cdot \hat{\boldsymbol{x}}_i^b)$. This concept is easily extended to take information about obstacles in the environment into account by excluding poses from the IRM that would result in a collision. The information matrices place high weight on the inverse reachability constraints an low ones on the absolute positioning.

D. Merging Multiple Demonstrations

To handle multiple demonstrations of one task, all demonstrations are merged into one graph. The individual demonstrations are connected via the nodes for the transforms between the object and the human hand respectively the robot gripper. Explicitly the nodes V_{HG} and V_{RG} are shared among all demonstrations. Further we introduce binary edges connecting the last gripper pose in the reaching segment of each demonstration, i.e. the grasp pose, to the grasp poses in all other demonstrations. These edges use an error function like equation (3) punishing deviations between grasp poses of different demonstrations. For the transform between hand and robot gripper an individual node V_{gh} is added for each demonstration.

Recapitulating, the input for our graph optimization are recorded trajectories for the human hand, torso and object poses. In the optimization process we generate corrected poses for the aforementioned and the robot gripper and base trajectories. Additionally we estimate the human grasp and a corresponding robot grasp on the object.

V. Generate Robot Motions

After optimizing the trajectories towards the robots capabilities, we use these to learn time-driven models using mixtures of Gaussians [5]. We learn a combined model in Cartesian space for both the gripper and base motion. As we are only interested in the applicability of the trajectories, we use the learned models to generate trajectories. These are executed on the robot in the corresponding object frame. During trajectory generation we perform obstacle avoidance to prevent collisions between the robot base and the map [19].

VI. Experiments

We present four experiments to evaluate different aspects of our approach to learning from human demonstrations.

A. Adapting Human Demonstrations to Robot Requirements

In the first experiment we evaluate our proposed graph structure and the optimization. We evaluate our method on four different tasks. The first consists of grasping a door handle, pushing it down and then opening and moving through the door and releasing the handle. The other three consist of operating small furniture pieces, i.e., opening a shelf swivel door, a drawer and a sliding door. We record the poses of the human hand and torso with attached markers (see Fig. 1). The door handle is also tracked using a marker. For the tracking of the other objects we use SimTrack [20]. For the drawer and sliding door task we did not use recorded torso positions but instead estimated them based on the hand positions. For these two tasks sets of 40 grasp samples each were manually defined instead of using [13]. Even then our system is capable of optimizing the data for all tasks. Table I gives a summary of the numerical evaluation of all tasks.

The evaluation is particularly interesting for a complex task that requires the base to navigate through a narrow passage while having a constrained gripper motion. Therefore we discuss the open and drive through door task in detail (first column in Table I). The task is demonstrated 10 times by a human teacher. In total, after inserting and interpolating missing poses (see Section IV-A), we have 1529 pose tuples. Each tuple consists of an object, a hand and a torso pose. A visualization of a few iterations steps of the optimization

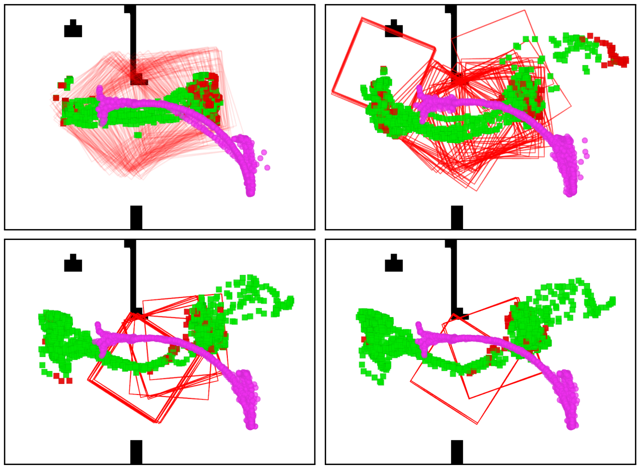

Fig. 4: Graph optimization procedure. The images show the graph poses after 0, 1, 5 and 20 iterations. The black parts in the images outline the map of the environment. The green and red dots describe the base pose of the robot where red signals kinematic infeasibility. The red squares indicate the footprint of the robot base. Note that only footprints that are in collision with the map are displayed here. At iteration zero base poses correspond to the demonstrated torso poses. As the robots base is wider than the teachers body most of the demonstrated poses are in collision with the map if executed by the robot. By reducing the inverse reachability map by poses that collide with the map during optimization most collisions with the map can be avoided for the learning data. The magenta dots describe the robot gripper trajectories. The displayed data covers 10 human demonstrations.

is shown in Fig. 4. First we analyze the distance between the closest sampled valid grasp pose and the gripper poses at the end of the grasping motions. As expected we see that the initial gripper poses are not well suited to perform the grasp with an offset of $2.8\,\mathrm{cm}$ to the next sampled grasp pose. After the optimization this distance drops to $0.5\,\mathrm{cm}$. We see the same behavior in the rotational part. As a result of this the mean distance between the demonstrated hand poses and the gripper poses increases from $0\,\mathrm{cm}$ to $2.2\,\mathrm{cm}$, which is necessary to correct the grasp. As for the transform between object and hand respectively gripper we are only interested in it being constant. Thus, we calculate the corresponding standard deviations. For these connections we see a drop in the standard deviation for both the translational and rotational parts to low values indicating that a stable transform was found. We also observe that almost all collisions with the map as well as kinematic problems were resolved. Note that individual low quality poses are not an issue since this is data is only used to learn a feasible action model. The results for the other tasks show similar behavior.

B. Automatic Positioning

In this experiment we show that our approach is able to reproduce actions independent of the starting pose. Specifically if reachability of all intermediate poses and the goal poses is not given from the starting robot base configuration. We test this in real world scenarios on the tasks of grasping the swivel and the room doors handles. Each task is reproduced on the PR2 with five different starting poses. We vary both the base and gripper configuration. We use Simtrack as before to detect the swivel door and a marker for the room door handle. Fig. 5 shows all five generated paths with our learned model for both the robot gripper and base for grasping the swivel door handle. All trials of grasping the handles were successful.

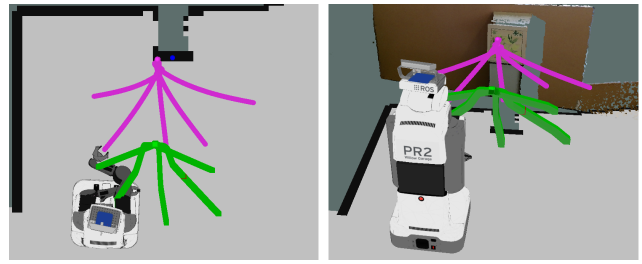

Fig. 5: Generated trajectories for different initial poses for the task of grasping the handle of the swivel door. The images show a top and a 3rd person view of the scene. The right image also displays the environment including the swivel door as a point cloud.

C. Task Relevant Base Motion

In the previous experiment we demonstrated that our approach is applicable to automatically position the robot for simple tasks. In this experiment we show that our approach is also able to generate a model for tasks that require a complex combined motion of the robot base and gripper. We demonstrate this on the task of opening and driving through a door. Fig. 6 shows generated trajectories for one of the trials for the robot base and gripper while opening and driving through the door. We conducted five trials in the real world setting and 20 trials in a simulated environment. In both settings the PR2 is used. All trials were executed successfully, which shows that our system is able to learn subtle gripper motions as needed to press down the door handle at the same time with the rougher overall path for both gripper and base to open the door. The opening of the swivel door was also successfully carried out five times in real world experiments.

D. Transition Between Actions

In a last experiment we demonstrate that the capability of our system to position the robot according to the desired task can be used to append actions that require different base positions. The transition between the actions is handled implicitly by the motion generation. To demonstrate this we append the task of opening the swivel door after opening the room door in the simulations of the previous experiment. We choose two different positions of the swivel door in the simulated world. We run 10 trials in each setup to append the tasks. The positions of the door handles in the simulated environments are given. All 20 runs to append the actions were executed successfully.

VII. Conclusion

We presented an approach that learns mobile manipulation actions from human demonstrations. Our approach does not rely on accurate tracking or robot tailored demonstrations. We use graph optimization to correct measurement errors and

	Room Door (10, 1529)		Swivel Door (4, 419)		Drawer (6, 656)		Sliding Door (10, 1482)	
	Before Opt.	After Opt.	Before Opt.	After Opt.	Before Opt.	After Opt.	Before Opt.	After Opt.
Euclidean distance gripper-grasp	2.82 cm	0.55 cm	2.36 cm	0.49 cm	6.33 cm	0.37 cm	3.23 cm	0.60 cm
Angular distance gripper-grasp	18.3°	8.0°	5.3°	0.7°	5.4°	1.6°	6.5°	0.5°
Euclidean distance gripper-hand	–	2.2 cm	–	2.68 cm	–	5.54 cm	–	3.13 cm
Angular distance gripper-hand	–	13.5°	–	3.0°	–	2.8°	–	5.8°
Std dev on gripper-object trans.	1.7 cm	0.53 cm	2.35 cm	0.21 cm	2.66 cm	0.18 cm	0.51 cm	0.12 cm
Std dev on gripper-object rot.	20.5°	2.4°	19.3°	1.6°	0.88°	0.21°	3.4°	0.34°
Std dev on hand-object trans.	1.7 cm	0.5 cm	2.35 cm	0.16 cm	2.66 cm	0.28 cm	0.51 cm	0.16 cm
Std dev on hand-object rot.	20.5°	4.6°	19.3°	0.9°	0.88°	0.3°	3.4°	0.6°
Map collision free poses	89.2 %	99.74 %	–	–	–	–	–	–
Kinematically achievable	69.8 %	96.86 %	85.9 %	99.52 %	87.3 %	100 %	63.2 %	99.93 %

TABLE I: Results for the optimization for all four trained tasks. The numbers after the task name denote the number of demonstrations for the task and the total number of recorded poses, after interpolating missing ones. The shown distance between gripper and grasp poses is a mean over the endpoints of the reaching segments of the demonstrations. For the distance between gripper and hand as well as the collisions and the kinematic feasibility all pose tuples are considered. Kinematic feasibility expresses the lookup in the inverse reachability map. For the relation between object and robot gripper respectively human hand a mean over all poses in the manipulation segments is calculated. Since gripper poses are initialized with the measured hand poses no meaningful distance before optimization can be given. For the three furniture operating tasks no collisions with the map are considered.

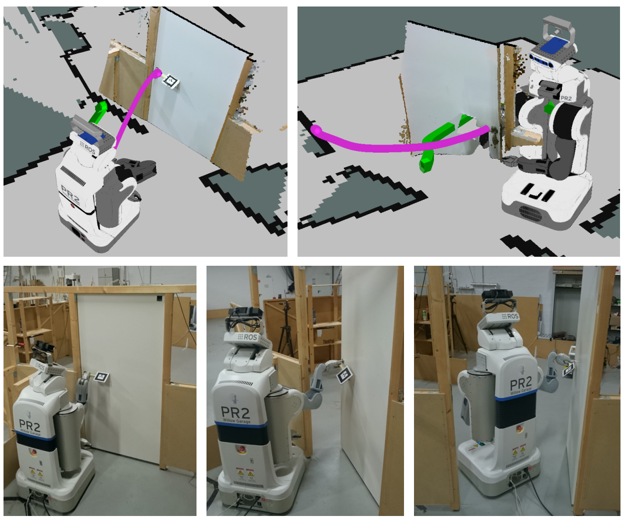

Fig. 6: In the top row the figure illustrates exemplary generated trajectories to grasp the handle (left) and then open the door (right). The bottom row shows the PR2 during task execution.

adapt the demonstrations to the robot capabilities. In contrast to our previous work, in this paper we also learn feasible robot base motions accounting for kinematic constraints between robot base and gripper. We jointly integrate this with adapting the gripper motion, while we consider grasping capabilities of the robot as well as occlusions in demonstrations. Our experiments demonstrate that our approach is able to reproduce a task that requires subtle gripper motion, while the robot base needs to navigate through a narrow passage. It also enables the robot to position itself according to the demands of an action. We believe this to be specifically beneficial in the context of task planning as it facilitates the transition between individual actions when solving complex tasks, which we plan to address in future work.

References

[1] T. Welschehold, C. Dornhege, and W. Burgard, "Learning manipulation actions from human demonstrations," in *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2016.

[2] S. Calinon, T. Alizadeh, and D. G. Caldwell, "On improving the extrapolation capability of task-parameterized movement models," in *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2013.

[3] O. Kroemer, C. Daniel, G. Neumann, H. van Hoof, and J. Peters, "Towards learning hierarchical skills for multi-phase manipulation tasks," in *Int. Conf. on Robotics & Automation (ICRA)*, 2015.

[4] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal, "Learning and generalization of motor skills by learning from demonstration," in *Int. Conf. on Robotics & Automation (ICRA)*, 2009.

[5] S. Calinon, Z. Li, T. Alizadeh, N. G. Tsagarakis, and D. G. Caldwell, "Statistical dynamical systems for skills acquisition in humanoids," in *Int. Conf. on Humanoid Robots (Humanoids)*, 2012.

[6] M. Mühlig, M. Gienger, and J. Steil, "Interactive imitation learning of object movement skills," *Autonomous Robots*, vol. 32, no. 2, pp. 97–114, 2012.

[7] J. Koenemann, F. Burget, and M. Bennewitz, "Real-time imitation of human whole-body motions by humanoids," in *Int. Conf. on Robotics & Automation (ICRA)*, 2014.

[8] E. Ratner, B. Cohen, M. Phillips, and M. Likhachev, "A web-based infrastructure for recording user demonstrations of mobile manipulation tasks," in *Int. Conf. on Robotics & Automation (ICRA)*, 2015.

[9] F. Burget, A. Hornung, and M. Bennewitz, "Whole-body motion planning for manipulation of articulated objects," in *Int. Conf. on Robotics & Automation (ICRA)*, 2013.

[10] S. Gray, S. Chitta, V. Kumar, and M. Likhachev, "A single planner for a composite task of approaching, opening and navigating through non-spring and spring-loaded doors," in *Int. Conf. on Robotics & Automation (ICRA)*, 2013.

[11] N. Vahrenkamp, T. Asfour, and R. Dillmann, "Robot placement based on reachability inversion," in *Int. Conf. on Robotics & Automation (ICRA)*, 2013.

[12] F. Stulp, A. Fedrizzi, and M. Beetz, "Learning and performing place-based mobile manipulation," in *2009 IEEE 8th International Conference on Development and Learning*, June 2009, pp. 1–7.

[13] D. Fischinger, A. Weiss, and M. Vincze, "Learning grasps with topographic features," *Int. Journal of Robotics Research*, vol. 34, no. 9, pp. 1167–1194, 2015.

[14] A. ten Pas and R. Platt, "Using geometry to detect grasps in 3d point clouds," in *Int. Symposium of Robotics Research (ISRR)*, 2015.

[15] G. Gioioso, G. Salvietti, M. Malvezzi, and D. Prattichizzo, "An object-based approach to map human hand synergies onto robotic hands with dissimilar kinematics," in *Robotics: Science and Systems (RSS)*, 2013.

[16] Y. Lin and Y. Sun, "Robot grasp planning based on demonstrated grasp strategies," *Int. Journal of Robotics Research*, vol. 34, no. 1, pp. 26–42, 2015.

[17] H. B. Amor, O. Kroemer, U. Hillenbrand, G. Neumann, and J. Peters, "Generalization of human grasping for multi-fingered robot hands," in *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2012.

[18] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "g2o: A general framework for graph optimization," in *Int. Conf. on Robotics & Automation (ICRA)*, 2011.

[19] S. M. Khansari-Zadeh and A. Billard, "A dynamical system approach to realtime obstacle avoidance," *Auton. Robots*, vol. 32, no. 4, pp. 433–454, May 2012, ePFL.

[20] K. Pauwels and D. Kragic, "Simtrack: A simulation-based framework for scalable real-time object pose detection and tracking," in *Int. Conf. on Intelligent Robots and Systems (IROS)*, 2015.