Learning Relevant Features for Manipulation Skills using Meta-Level Priors

Oliver Kroemer and Gaurav S. Sukhatme

University of Southern California

Abstract—**Robots can generalize manipulation skills between different scenarios by adapting to the features of the objects being manipulated. Selecting the set of relevant features for generalizing skills has usually been performed manually by a human. Alternatively, a robot can learn to select relevant features autonomously. However, feature selection usually requires a large amount of training data, which would require many demonstrations. In order to learn the relevant features more efficiently, we propose using a meta-level prior to transfer the relevance of features from previously learned skills. The experiments show that the meta-level prior more than doubles the average precision and recall of the feature selection when compared to a standard uniform prior. The proposed approach was used to learn a variety of manipulation skills, including pushing, cutting, and pouring.**

I. Introduction

In the future, we want robots to perform a wide range of manipulation tasks in everyday environments. These robots will need to be capable of learning new skills from demonstrations and generalizing these skills between different objects and scenarios. Variations in objects and their state can be represented by object features. Object features could, for example, include the 3D position and size of a container's opening. Using these features, the robot can learn to generalize a pouring skill between cups and bowls at different positions in the robot's workspace. Similarly, a robot could adapt a cutting skill using the position and length features of a knife's edge. A robot can thus represent task scenarios using geometric features of affordance-bearing parts.

Many of the extracted object features will however not be relevant for generalizing the manipulation skill. One of the key challenges for learning manipulation skills in unstructured environments is therefore to determine which object features are relevant for adapting and generalizing the skills. One approach would be to have the human demonstrator manually label the relevant features. However, this approach would require a considerable amount of additional effort from the human demonstrator and limit the robot's autonomy. Instead of relying on prior human knowledge, the robot should utilize information about relevant features from skills that it has previously learned. Using this information, the robot can autonomously create a prior regarding which features are relevant for generalizing novel manipulation skill.

Predicting the relevance of an object feature is not a trivial task. As there is no one-to-one mapping between features of different tasks, the robot will need to generalize relevance between different features. This approach may seem impossible if we consider a feature as simply being a number. However,

for manipulation tasks, the features are generally grounded in the geometry and properties of the objects being manipulated. Hence, we can define meta features that capture characteristics of the features and how they relate to the skill being learned. Using these meta features, the robot can learn a meta prior for computing a prior over the relevance of features in novel manipulation tasks [30]. This prior can then be incorporated in the feature selection process in order to determine the relevant features from demonstrations more efficiently.

In this paper, we investigate learning meta-level priors for generalizing the relevance of features between different manipulation skills. An overview of the proposed approach is shown in Fig. 1. In order to learn a new versatile manipulation skill, the robot is first provided with multiple demonstrations of the task using kinaesthetic teaching. The robot is also provided with 3D point cloud models of the manipulated objects. The robot uses these demonstrations and models to extract a set of object features and their corresponding skill parameters for training the skill. The ultimate goal of the proposed framework is to select a sparse set of relevant object features and to learn the corresponding parameters for adapting the skill execution accordingly. The part-based feature generation process is explained in Section III, and the parameterized skill representation is described in Section IV-A.

For each object feature, the robot also extracts a set of meta features, which describe the relationship between the feature and the skill being learned. These meta features are used to compute a prior over the relevance of the feature. For example, as each object feature is associated with an object part, one meta feature may describe the distance between the robot's hand and the feature's part at the start of the demonstration. Another meta feature may represent this distance at the end of the demonstration. Using these two meta feature, the robot may learn that a feature is more likely to be relevant if the robot's hand moves closer to the feature's part during the task. The mapping from meta features to the prior over the feature's relevance is based on a meta prior. The meta-level prior is learned using the meta features and corresponding feature relevances from previous skills, as illustrated in the top row of Fig. 1. The meta features and meta-level prior are explained in Sections III-B and IV-B respectively.

Given the training features and skill parameters, as well as the feature relevance prior, the next step is to autonomously select a sparse set of features for generalizing the skill between scenarios. The feature selection is performed using stochastic search variable selection (SSVS) [12]. We extend this method to include the meta-level prior into the feature selection
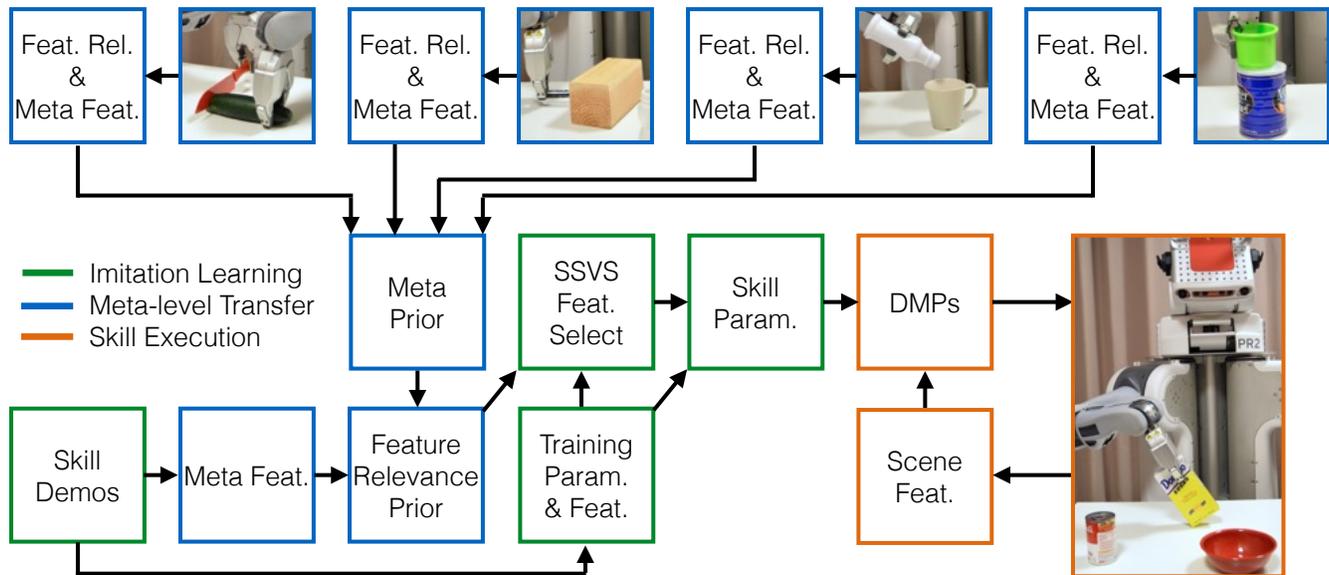
Figure 1. The figure shows an overview of the proposed approach. The green blocks correspond to learning the skill parameters from demonstrations. The orange blocks correspond to executing the skill in novel situations using the learned parameters. The blue blocks correspond to learning the meta-level prior, which can be used to compute a feature relevance prior for the new skill based on the meta features extracted from the skill demonstrations.

process. The set of relevant features are inferred using a Gibbs sampling approach. The feature selection process is explained in Sections IV-B and IV-C. The robot subsequently learns a set of skill parameters for the selected features.

When presented with a new scene, the robot first computes the set of object features for the current scenario. It then adapts the skill execution based on the learned parameters in order to perform the manipulation task. The proposed framework was evaluated on placing, pushing, tilting, cutting, pouring, and wiping tasks using a PR2 robot. The experiments are described and discussed in Section V. Using the meta-level prior, the robot was able to more than double the average precision and recall of the feature selection.

II. RELATED WORK

Recent work in both imitation learning and policy search reinforcement learning have allowed robots to learn and execute complicated motor skills. Example skills include playing ball games [21, 8, 40, 27], opening doors [19, 22], scrubbing surfaces [33, 6], and manipulating objects [38, 32, 4]. These learned skills adapt to different features of the task, e.g., the position of a ball or the goal location. However, the features used to adapt the skill executions are usually predefined and do not include irrelevant features. These methods therefore focus on determining how to adapt to the features, and not learning which features to adapt to.

Some frameworks learn relevant features in order to select actions more accurately in otherwise ambiguous situations [41, 31]. These approaches often perform an implicit pose estimation of an object or part, which the robot can then use to adapt its actions. However, the features are usually learned for specific sets of objects and not to generalize between objects with different shapes.

Motor skills can also be adapted to different situations by selecting suitable task frames. These task frames are often defined relative to objects for manipulation tasks and, hence, there are usually multiple potential task frames to choose from. The robot can learn to select a task frame from multiple demonstrations based on the the variance of the trajectories in each potential task frame [37, 36]. Task frames can be generalized between different sets of objects by associating the coordinate frame to object parts with similar shapes [50, 2, 24]. In these cases, learning the task frame is approached as a part detection and pose estimation problem. Although task frames are fundamental to generalizing manipulations, many skill adaptations rely on additional information regarding the size and shape of the objects being manipulated [53].

Another important challenge for skill learning is determining which objects are involved in a manipulation [20]. The set of relevant objects can often be extracted from demonstrations using visual cues, such as motionese [29]. The object selection problem is however distinct from the feature selection problem addressed in this paper, and not all of the features associated with an object will be relevant for generalizing a skill.

Several works have investigated transfer and multi-task learning in the field of robotics [49, 34, 9]. These approaches often focus on transferring trajectories or controllers between different tasks. The different tasks are also often quite similar and share the same feature space, e.g., reaching for different locations in the task frame may be considered as different tasks. In our work, the tasks have distinct sets of features and the robot is learning a meta-level prior for transferring the relevance of features between skills. Meta features have previously been used to transfer knowledge about the relevance of features between tasks in applications such as predicting movie ratings, text classification, and object recognition [30, 52, 26].

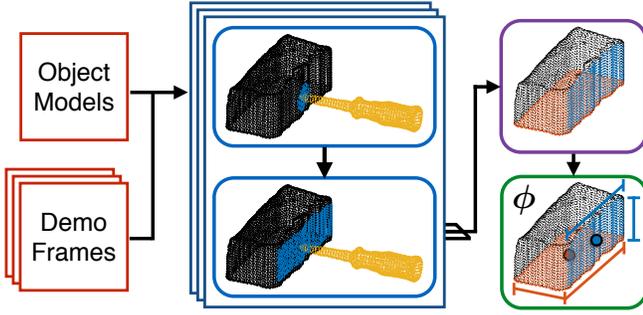Our feature generation process is based on decompos-

Figure 2. The figure illustrates the key steps of the feature generation process. (Red) The inputs to the process are a set of 3D point cloud models of the objects being manipulated, and a set of scene frames from a demonstration of the task. (Blue) For each frame and pair of objects, the robot estimates a contact region. Using this region as an initialization, the robot computes an estimate of the affordance-bearing part using GrabCut segmentation. (Purple) The individual estimates of the object parts are clustered in order to remove redundant and noisy part estimates. (Green) The robot uses a predefined set of rules to generate three position and three length features for each part. The x-y-z positions and lengths are defined in the task frame.

ing objects into parts. The segmentation is guided by the demonstrations of the skill. Several other works have also used actions in order to improve segmentation performance [16, 54, 15]. However, most of these methods focus on segmenting scenes into objects rather than parts. Decomposing objects into parts has largely been based on geometric features, such as convexity [46, 48, 14]. Lakani et al. recently proposed a method for segmenting objects into parts based on where it could be grasped by the robot [28]. Their approach uses an initial over-segmentation of an object model into supervoxels and merges these supervoxels into parts.

III. GENERATING FEATURES AND META FEATURES

In order to learn a skill that generalizes between scenarios, the robot first requires a suitable method for representing the scenarios that it encounters. In particular, the robot must generate a set of features describing the objects and their current configuration. Generalizing between objects is often easier when the representation explicitly models the manipulation-relevant parts of the objects [51, 24]. Our feature generation approach therefore begins by identifying parts of objects that may be relevant to the task. The part segmentation is based on the skill demonstrations in order to extract task-relevant parts. The part detection method is described in Section III-A. For each of the resulting parts, the robot generates a set of features describing the position and size of the part. An overview of the feature generation process is shown in Fig. 2.

Each of the resulting features is associated with a set of meta features, which are used to predict the relevance of the feature for the new task. The object features and their meta features are described in Section III-B.

A. Detecting Affordance-bearing Parts from Demonstrations

Affordance-bearing parts are important because they indicate regions of an object that can be used for manipulating or interacting with other objects [13], e.g., the edge of a knife for cutting or the handle of a cup for grasping. Most of these interactions involve direct physical contact or close proximity between the objects. The affordances of a part also depend on its shape, with similarly-shaped regions of objects having similar affordances.

Given these two insights, we propose extracting affordance-bearing object parts using the GrabCut segmentation algorithm [44]. GrabCut is commonly used to segment the foreground from the background in photo editing. The segmentation iterates between modeling the pixels' features in each segment and segmenting the pixels into two groups using a min-cut approach. The distribution over the pixel assignments is modeled as a Markov random field. In the photo-editing domain, the segmentation is initialized by a human user providing a coarse estimate of points that belong to the foreground.

In our approach, the robot uses GrabCut to segment the task-relevant part from the rest of the object. Given a 3D point cloud model, the robot computes the position, the normal, the curvature, and spectral features for each point based on its local neighbourhood [35, 3]. The part's affordances may depend on other properties, e.g., material and texture, but these are beyond the scope of this paper. The robot also creates a k-nearest neighbour graph connecting the points. The unary potentials are modeled using Gaussian mixture models over the local geometric features. For the pairwise potentials, we employ the commonly used Potts model [42]. GrabCut allows for observation-dependent pairwise potentials, but the model then does not represent a Markov random field prior [43].

An object part is relevant to the task if it is used to interact with other objects. The segmentation was therefore initialized by selecting points that were likely to be interacting with other objects during the demonstrations. As many interactions are based on direct physical contact, or close proximity, we considered a point to be interacting if it was within a threshold distance of the points of another object with opposing normals. This spatial proximity assumption has been shown to be applicable to a wide variety of manipulation tasks [1].

Given a set of object models and their trajectories during a demonstration, we begin by computing the GrabCut segmentation for each pair of objects for each frame in which they are in close proximity. This process generates multiple redundant estimates of the object's parts. In order to merge the part estimates across the frames, we cluster the part estimates using spectral clustering [7]. In our experiments, we computed the similarity between the parts using a Bhattacharayya kernel over the parts' positions and normals in the object frame [18]. The clustering was performed multiple times with different numbers of clusters, and we selected the clustering that maximized the intra-cluster kernel values minus the inter-cluster kernel values. The final part included all of the points that were present in the majority of the cluster's estimates.

B. Object Features and Meta Features

The next step is to generate a set of features describing the constellation of objects and parts in a scene. A subset of these features will ultimately be used by the robot to generalize the learned manipulation skill to different situations. Rather than defining specific features for one task, we want to define
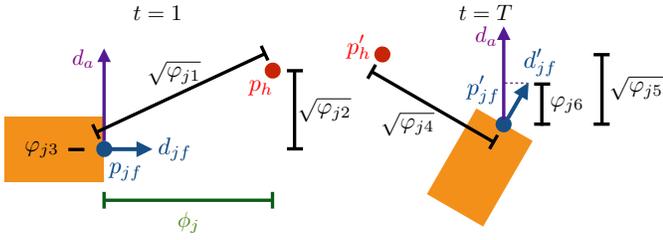
Figure 3. This figure illustrates the meta features (black) computed for a feature ϕ_j. This feature (green) describes the initial horizontal distance between the side part of the object (orange) and the robot's hand p_h (red). The feature is characterized by a feature position p_{jf} and a direction d_{jf} shown in blue. The prior is being computed for the movement in the vertical action direction d_a (purple). The left side of the figure shows the initial object configuration $t = 1$, while the right side shows the final configuration $t = T$. The black lines indicate the meta features. The seventh meta feature $\varphi_{j7} \in \{-1, 1\}$ (not shown) is set to $\varphi_{j7} = 1$ as the feature defines the position, and not the size, of a part.

a generic set of rules for generating features for arbitrary objects across different manipulation tasks. In order to create a consistent set of features between different scenes, we assume that the number of objects and the number of respective parts are the same across different instances of a given task.

To describe the objects in the scene, the robot needs features that capture the locations and shapes of the individual object parts. The positions and sizes of the parts are computed by first fitting an axis-aligned bounding box to each of the parts. The x-y-z position of the part is then given by the center of the bounding box. The position is defined relative to the initial position of the robot's hand. The x-y-z dimensions of the box are used as features to capture the size of the part. This bounding box representation is relatively coarse, and some of the geometric details of the parts are lost. However, most of the extracted parts will have basic shapes. By using a part-based approach, the proposed representation automatically creates more features to describe the task-relevant regions of the objects. In the future, the set of features could be extended to include other potentially relevant object properties, e.g., the principal axes of the parts and the objects' masses.

Each task will have its own set of object features due to the automatic feature generation. Hence, the robot needs a suitable method for computing the similarity between features in order to transfer the feature-relevance prior between tasks. We can define this similarity by using meta features $\varphi_{jh} \forall h \in \{1, ..., H\}$ to describe the features $\phi_j \forall j \in \{1, ..., M\}$ as well as their relation to the skill component being learned [30].

We describe each feature ϕ_j by first associating it with a 3D position p_{jf} and direction d_{jf}. The position p_{jf} corresponds to the initial position of the feature's part. The direction d_{jf} corresponds to the direction in which the feature is computed, e.g., $d_{jf} = [1\,0\,0]^T$ if the feature ϕ_j describes the position or size of a part along the x axis. In order to relate the feature to the skill learning, we also define the initial 3D position of the hand p_h and the direction d_a of the skill component being learned, i.e., $d_a = [0\,0\,1]^T$ when learning the z component of the manipulation skill.

Given these variables, the first meta feature describes the initial proximity between the feature and the robot's hand

$\varphi_{j1} = ||p_{jf} - p_h||^2$. The second meta feature describes the proximity in the direction of the skill component $\varphi_{j2} = ||d_a^T(p_{jf} - p_h)||^2$, and the third meta feature describes the alignment between the feature and the skill component $\varphi_{j3} = |d_a^T d_{jf}|$. The next three meta features, φ_{j4} to φ_{j6}, are computed in the same manner as the first three, but using the feature position p'_{jf}, feature direction d'_{jf}, and hand position p'_h at the end of the skill demonstration. The skill direction d_a is kept constant while the feature direction d_{jf} is assumed to move together with object part, as illustrated in Fig. 3. The seventh meta feature φ_{j7} indicates whether the feature represents the position $\varphi_{j7} = 1$ or the size $\varphi_{j7} = -1$ of a part. The final meta feature $\varphi_{j8} = 1$ is a bias term.

Additional meta features could be explored in the future, e.g., the feature's alignment with the part's mean normal, or if the part is currently interacting with another part. Meta features could also be used to describe other types of object features, e.g., features generated by a convolutional neural network [31]. The current set of meta features allows the robot to capture the proximity and alignment of the various object features. The meta features also adapt to the direction of the skill component being learned and which hand is being used. In Section IV-B, we explain how these meta features can be used to compute a prior over the relevance of the feature.

IV. LEARNING RELEVANT FEATURES FOR GENERALIZING MANIPULATION SKILLS

In order to perform manipulations under different situations, the robot must learn a manipulation skill that adapts to the extracted object features. We use dynamic motor primitives (DMPS) to represent the skills. Our reformulation of the parameterized DMPs with object features is explained in Section IV-A. In order to select a relevant set of features for generalizing the skill, the robot uses stochastic search variable selection (SSVS) with a meta-level prior to transfer information regarding feature relevance between skills. The probabilistic model over the skill parameters is explained in Section IV-B. The feature selection process, based on Gibbs sampling, is described in Section IV-C.

A. Motor Primitive Skill Representation

The robot's skills are represented using dynamic motor primitives (DMPs) [17, 47]. A DMP consists of linear dynamical systems of the form

$$\ddot{y} = \alpha_z(\beta_z \tau^{-2}(y_0 - y) - \tau^{-1}\dot{y}) + \tau^{-2}\sum_{j=1}^{M} \phi_j f(x; \boldsymbol{w}_j),$$

where y is the state, y_0 is the initial state, α_z and β_z are constants that define the spring and damper coefficients, τ is a time coefficient, x is the state of the canonical system, and f is a forcing function. Each skill component, e.g., the hand's movement in the z direction, is modeled by a separate linear system. The canonical state x acts as a timer for synchronizing multiple linear systems. It starts at $x = 1$ and decays according to $\dot{x} = -\tau x$. In our framework, the amplitude parameters ϕ_i correspond to the object features in

$$j \in \{1,...,M\} \quad k \in \{0,...,K\}$$
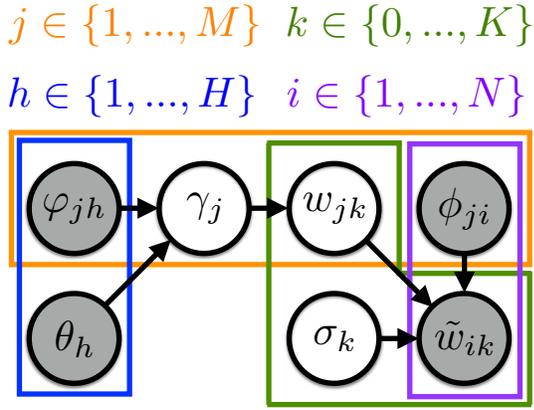$$h \in \{1,...,H\} \quad i \in \{1,...,N\}$$

Figure 4. The graphical model for the skill parameters. Shaded nodes indicate observed variables, and white nodes represent latent variables. Repeated variables are indicated by the square plates and the variables' indices. The figure is best viewed in color.

order to allow the skill to adapt to different scenarios. The adaptation of the skill's trajectory to an object feature ϕ_j is defined by its corresponding forcing function $f(x; \boldsymbol{w}_j)$, where the k^{th} element of the vector $\boldsymbol{w}_j \in \mathbb{R}^K$ is given by the parameter $[\boldsymbol{w}_j]_k = w_{jk-1}$. The forcing function f is represented using a locally weighted regression of the form

$$f(x; \boldsymbol{w}_j) = \alpha_z \beta_z \left(\frac{\sum_{k=1}^{K} \psi_k(x) w_{jk} x}{\sum_{k=1}^{K} \psi_k(x)} + w_{i0}\psi_0(x) \right),$$

where $\psi_k \forall k \in \{1,...,K\}$ are Gaussian basis functions, and ψ_0 is a basis function that follows a minimum jerk trajectory from 0 to 1. The shape of the DMP's trajectory is defined by the parameters w_{ij} together with their corresponding basis functions. The standard formulation of the DMPs, with the explicit goal-state parameter g, can be obtained by setting $w_{1k} = 0 \forall k \in \{1,...K\}$, $w_{10} = 1$, $\psi_0(x) = 1$, and $\phi_1 = g - y_0$. In our reformulation, the first term of the linear system is defined entirely by the initial state of the robot and its environment, and the the goal states are absorbed into the object features ϕ_i.

B. Probabilistic Model for Skill Learning with Sparse Features

The next step is to learn the shape parameters w_{jk}, which determine how the DMP adapts to changes in the object features. In theory, the robot could learn a skill that adapts to every object feature extracted from the scenes. However, in practice, most of the object features will be irrelevant for generalizing the skill and the corresponding skill parameters should be close to zero. The robot should therefore learn to select the relevant object features and ignore the irrelevant ones.

Our approach to learning the relevant object features is based on the stochastic search variable selection (SSVS) approach [12]. The graphical model of our framework is shown in Fig. 4. As the relevant object features may be different for each linear system of the DMP, we perform a separate feature selection for each skill component. This probabilistic model captures not only the distribution over the skill parameters $\boldsymbol{w}_j$, but also the relevance $\gamma_j \in \{0,1\}$ of the object feature ϕ_j.

The relevance γ_j is distributed according to a Bernoulli distribution $\gamma_j \sim \text{Bern}((1 + \exp(-\boldsymbol{\varphi_j}^T\boldsymbol{\theta}))^{-1})$ where the meta features $\boldsymbol{\varphi}_j \in \mathbb{R}^H$ and the hyperparameters $\boldsymbol{\theta} \in \mathbb{R}^H$ define the prior distribution over the feature's relevance. The hyperparameters $\boldsymbol{\theta}$ are known as the *meta prior* [30]. Since the prior on γ_j has a logistic regression form, the values of the meta prior are learned from previous skills using iteratively reweighted least squares.

The relevance γ_j determines the prior variance over the shape parameters $\boldsymbol{w}_j$ for the j^{th} object feature ϕ_j. If the feature is relevant $\gamma_j = 1$, then the distribution over the shape parameters $w_{jk} \forall k \in \{0,...,K\}$ is given by $p(w_{jk}|\gamma_j = 0, \check{s}, \hat{s}) = \mathcal{N}(0, \check{s}^2)$, where the standard deviation $\check{s}$ defines a broad Gaussian. If the feature is irrelvant $\gamma_j = 0$, then the distribution is given by $p(w_{jk}|\gamma_j = 1, \check{s}, \hat{s}) = \mathcal{N}(0, \hat{s}^2)$, where $\hat{s} \ll \check{s}$. In this manner, the shape parameters corresponding to an irrelevant feature are more likely to be close to zero and, hence, have less effect on the generalization of the skill.

The distribution over the skill parameters is to be inferred from a set of N training demonstrations. The i^{th} training demonstration includes a set of values for the object features $\phi_{ji} \forall j \in \{1,..M\}$ and a trajectory ξ_i representing the robot's actions. Rather than working directly with the trajectory data ξ_i, we transform each trajectory into a separate set of target shape parameters $\tilde{w}_{ik}$. These target parameters are equivalent to the shape parameters if the DMP were learned using only a single constant object feature $\phi_1 = 1$. The goal is thus to learn a set of skill parameters $\boldsymbol{w}_j$ that approximate these target values $\tilde{w}_{ik} \approx \sum_{j=1}^{M} w_{jk}\phi_{ji}$ for all of the training trajectories. We model the distribution over the target values using a normal distributions $\tilde{w}_{ik} \sim \mathcal{N}(\sum_{j=1}^{M} w_{jk}\phi_{ji}, \sigma_k^2)$, where σ_k^2 is the output variance in the shape parameters. We model the distribution over these variances using an inverse-gamma distribution $\sigma_k^2 \sim \text{Inv-Gamma}(a, b)$. We set the shape and scale parameters a and b to constants, which worked well in the experiments. Alternatively, the hyperparameters can depend on the number of relevant features $\sum_j \boldsymbol{\gamma}_j$ such that the variance is smaller if the model uses more features [12].

C. Inferring Relevant Object Features

Given a matrix $\tilde{\boldsymbol{W}} \in \mathbb{R}^{N \times M}$ of target shape parameters $[\tilde{\boldsymbol{W}}]_{i,j} = \tilde{w}_{ij}$ and a matrix $\boldsymbol{\Phi} \in \mathbb{R}^{M \times N}$ of object features $[\boldsymbol{\Phi}]_{j,i} = \phi_{ji}$ from N demonstrations, the next step is to estimate the relevant features $\boldsymbol{\gamma}$ for generalizing the manipulation skill to new scenarios. In particular, we want to determine which feature are likely to be relevant $\gamma_j = 1$ under the posterior distribution $p(\boldsymbol{\gamma}|\tilde{\boldsymbol{W}}, \boldsymbol{\Phi})$. The posterior distribution $p(\boldsymbol{\gamma}|\tilde{\boldsymbol{W}}, \boldsymbol{\Phi})$ is given by

$$p(\boldsymbol{\gamma}|\tilde{\boldsymbol{W}}, \boldsymbol{\Phi}) = \int \int p(\boldsymbol{\gamma}, \boldsymbol{W}, \boldsymbol{\sigma}|\tilde{\boldsymbol{W}}, \boldsymbol{\Phi}) \mathrm{d}\boldsymbol{W} \mathrm{d}\boldsymbol{\sigma}$$

where $\boldsymbol{W} \in \mathbb{R}^{M \times K+1}$ is a matrix of all of the shape parameters $[\boldsymbol{W}]_{j,k} = w_{jk}$, and $\boldsymbol{\sigma}$ is the set of all standard deviations σ_k for the target weights. Although this distribution cannot be computed in closed form, it can be efficiently approximated using a Gibbs sampling approach [12, 11]. Gibbs sampling is a Markov chain Monte Carlo method that allows us to
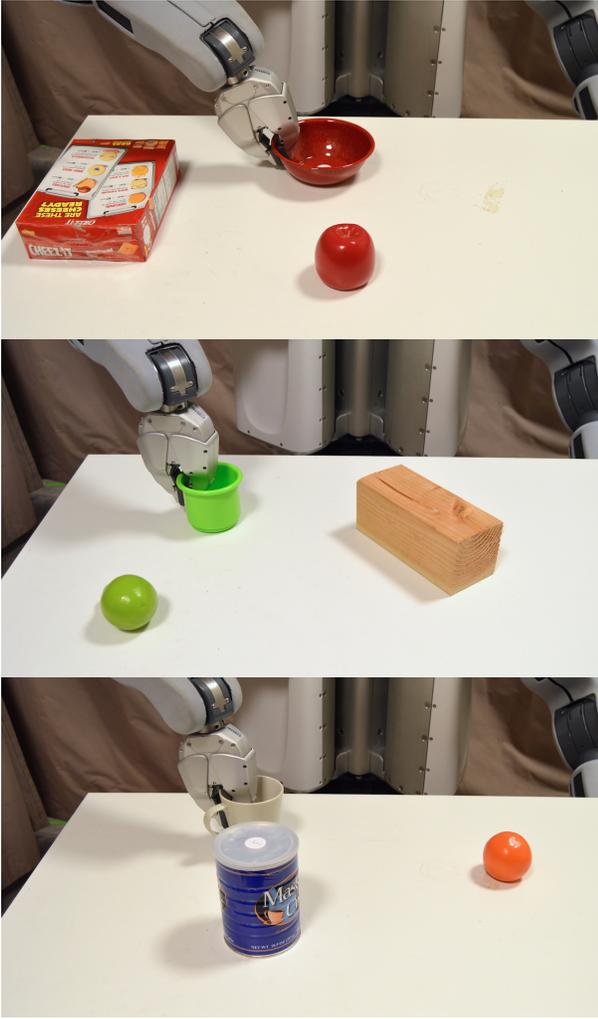
Figure 5. Example scenes for the placing task. The task is to place the grasped object centered on the box, block, or can depending on the scene.

Figure 6. The image shows the manipulated objects from the experiments. Some of the objects were used for learning more than one skill

Bernoulli distribution

$$\gamma_j \sim \text{Bern}\left(Z_j^{-1}\mathcal{N}(\boldsymbol{w}_j|0,\check{s}^2 I)p(\gamma_j=1)\right),$$

where

$$Z_j = \mathcal{N}(\boldsymbol{w}_j|0,\hat{s}^2 I)p(\gamma_j=0) + \mathcal{N}(\boldsymbol{w}_j|0,\check{s}^2 I)p(\gamma_j=1),$$

and the meta prior $\boldsymbol{\theta}$ and meta features $\boldsymbol{\varphi_j}$ are used to compute the prior

$$p(\gamma_j=1) = 1 - p(\gamma_j=0) = (1 + \exp(-\boldsymbol{\varphi_j}^T\boldsymbol{\theta}))^{-1}.$$

This sampling process continues to iterate between the different components to generate more samples from the posterior.

The Gibbs sampling provides the robot with a distribution over the relevance parameters. In order to ultimately select a set of relevant features, the robot computes the maximum a posteriori (MAP) estimate of the the relevance parameters. Hence, the robot selects a feature ϕ_j to be relevant iff the majority of the samples from the Gibbs sampling were $\gamma_j = 1$ [12]. This MAP estimate also corresponds to the Bayes estimate of the relevance parameters under a $0/1$ loss function. The skill parameters $\boldsymbol{w}_j$ for the selected features are then computed using linear regression. When presented with a new scenarios, the robot extracts the object features and uses the learned skill parameters to perform the manipulation task.

V. EXPERIMENTS

The proposed approach was evaluated on a series of manipulation tasks using a PR2 robot, as shown in Fig. 1. The robot has two arms with seven degrees of freedom each. Each arm is equipped with a parallel-jaw gripper. The robot observes the scene using a head-mounted kinect. The extraction of the object parts based on the skill demonstrations is detailed in Section V-A. In Section V-B, we discuss a benchmark experiment comparing the performance of the feature selection when using a meta prior versus a standard uniform prior. In a second benchmark experiment, we evaluated the accuracy of the predicted goal states of the learned skills, as described in Section V-C. The evaluation of the robot's ability to learn skills using feature selection is discussed in Section V-D.

A. Extracting Object Parts

The robot was initially shown demonstrations, using kinaesthetic teaching, of each of the following tasks: placing an object on top of another object (place), pushing an object with an elongated object (push), tilting a box between two adjacent
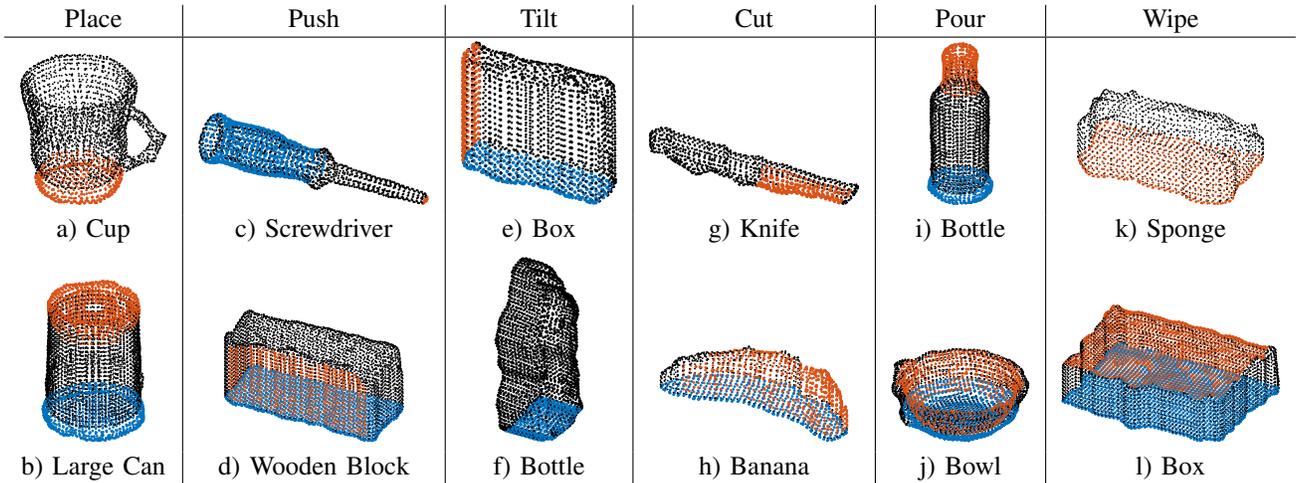
draw samples from the joint distribution $p(\boldsymbol{\gamma}, \boldsymbol{W}, \boldsymbol{\sigma}|\tilde{\boldsymbol{W}}, \boldsymbol{\Phi})$ by iteratively sampling over each of the individual components, i.e., $\boldsymbol{\gamma}$, $\boldsymbol{W}$, and $\boldsymbol{\sigma}$, given all of the other components.

We initially set all of the relevance parameters $\boldsymbol{\gamma}$ to one and the standard deviations $\boldsymbol{\sigma}$ to a predefined value. We also compute the priors over each relevance parameter $p(\gamma_j)$ based on their respective meta features φ_{jh} and the meta prior θ_h. We then sample a new set of shape parameters $\boldsymbol{w}_k \in \mathbb{R}^M \forall k \in \{0,...,K\}$ according to

$$\boldsymbol{w}_k \sim \mathcal{N}\left((\boldsymbol{\Phi}\boldsymbol{\Phi}^T + \sigma_k^2 \boldsymbol{R}^{-1})^{-1}\boldsymbol{\Phi}\tilde{\boldsymbol{w}}_k|(\boldsymbol{R}^{-1} + \sigma_k^{-2}\boldsymbol{\Phi}\boldsymbol{\Phi}^T)^{-1}\right)$$

where $\tilde{\boldsymbol{w}}_k \in \mathbb{R}^N$ is a vector containing the N samples' target shape parameters for the k^{th} basis function of the DMP, and $\boldsymbol{R} \in \mathbb{R}^{M \times M}$ is a diagonal matrix with elements $[\boldsymbol{R}]_{jj} = \hat{s}^2$ if $\gamma_j = 0$ and $[\boldsymbol{R}]_{jj} = \check{s}^2$ if $\gamma_j = 1$. Given these updated shape parameters, we then sample a new set of variance terms $\sigma_k^2 \forall k \in \{0,...K\}$ using an inverse gamma distribution

$$\sigma_k^2 \sim \text{Inv-Gamma}(a, b + (2(\tilde{\boldsymbol{w}}_k - \boldsymbol{\Phi}^T\boldsymbol{w}_k)(\tilde{\boldsymbol{w}}_k - \boldsymbol{\Phi}^T\boldsymbol{w}_k))^{-1}).$$

Finally, using the updated shape and variance parameters, we sample a new set of relevance parameters γ_j according to the

Place	Push	Tilt	Cut	Pour	Wipe
a) Cup	c) Screwdriver	e) Box	g) Knife	i) Bottle	k) Sponge
b) Large Can	d) Wooden Block	f) Bottle	h) Banana	j) Bowl	l) Box

Figure 7. Examples of object parts extracted from demonstrations of manipulations. The points show the objects' 3D models. The different colored points indicate parts of the point cloud. The top row consists of the grasped objects, while the bottom row indicates the objects that they interacted with using the red parts. For the tilting task, the red part corresponds to the supporting surface at the end of the task, and the mustard bottle (f) is irrelevant to the task.

sides (tilt), cutting an object with a knife (cut), emptying a bottle into a container (pour), and using an object to wipe the surface of another object (wipe). Each task was demonstrated five times with three different sets of objects, resulting in a total of $6 \times 5 \times 3 = 90$ demonstrations. The variety in the task scenarios is illustrated in Fig. 5, which shows example scenes for the placing task. Most of the objects used in the experiment are from the YCB object set [5].

Each demonstration was performed with three objects in the scene, although only one or two were relevant depending on the task. The manipulated objects are shown in Fig. 6. Coarse 3D models of the objects were computed from a single view point by extruding their shapes. Objects that moved during the demonstrations were tracked using the point cloud library's particle filter tracker [10, 45]. We assumed that grasped objects moved together with the robot's hand in order to track them.

Given the set of demonstrations and object models, the robot extracted the object parts as explained in Section III-A. The GrabCut segmentation was performed using a coefficient of one for the pairwise potentials, and mixtures of three Gaussians for the unary potentials. Points in the point cloud model were considered to be interacting with other objects if they were within 2cm of the other object's points and the inner product between their normals was less than -0.5. The trajectories were subsampled at 1Hz in order to reduce the computation time. The spectral clustering was performed using a Bhattacharyya kernel with the additional covariance hyperparameters set to 0.025m for the positions and 0.25 for the normals [18, 23]

Due to variations in the demonstrations and the stochasticity in the tracking and segmentation, the demonstrations resulted in varying numbers of object parts. Therefore, one demonstration that had the same number of object parts as the mode was selected as a prototypical example for each task. The parts of the other demonstrations were then matched to these demonstrations by selecting the most similar part. The similarity between parts was again computed using a Bhattacharyya kernel [18, 23]. In this manner, we created a consistent set of matched object parts across the different demonstrations.

From the demonstration's 270 object instances, 87.4% already had the same number of parts as their respective modes, 11.1% had more parts, and 1.5% had less parts. Demonstrations with additional parts were a minor issue as the matching process removed redundant and erroneous parts. In one pushing demonstration, the object was pushed close to the irrelevant object. This interaction generated additional parts which were then removed by the matching process.

Demonstrations with fewer parts are the result of interactions not being detected due to poor demonstrations or tracking. For example, while most of the pouring demonstrations were performed close to the container's rim, one demonstration involved pouring from a larger height of about 15cm. This interaction was not detected and hence no corresponding parts were generated. The interaction could be detected by observing the fluid [1], but currently only the containers are being tracked. For these demonstrations, the matching process chose another part to replace the undetected parts.

Examples of extracted parts for each task are shown in Fig. 7. The figure shows that the demonstration-based part segmentation successfully detects task relevant parts such as the blade of the knife for cutting, the opening of the bottle for pouring, and the supporting surface of the large can for placing. There are some variations in the sizes and shapes of the detected parts, which will influence the resulting features. For example, the shaft of the screwdriver, Fig. 7C, was incorporated into the tip part for some of the demonstrations. In order to estimate these variations, we computed the object-relative positions of the parts and the size of their bounding boxes. We then computed the standard deviations in these values across the five demonstrations for each object. The average standard deviation of the values was 3mm. It should be noted that the actual features, which are computed in the task frame and not the object frame, exhibit much greater variations across the demonstrations.
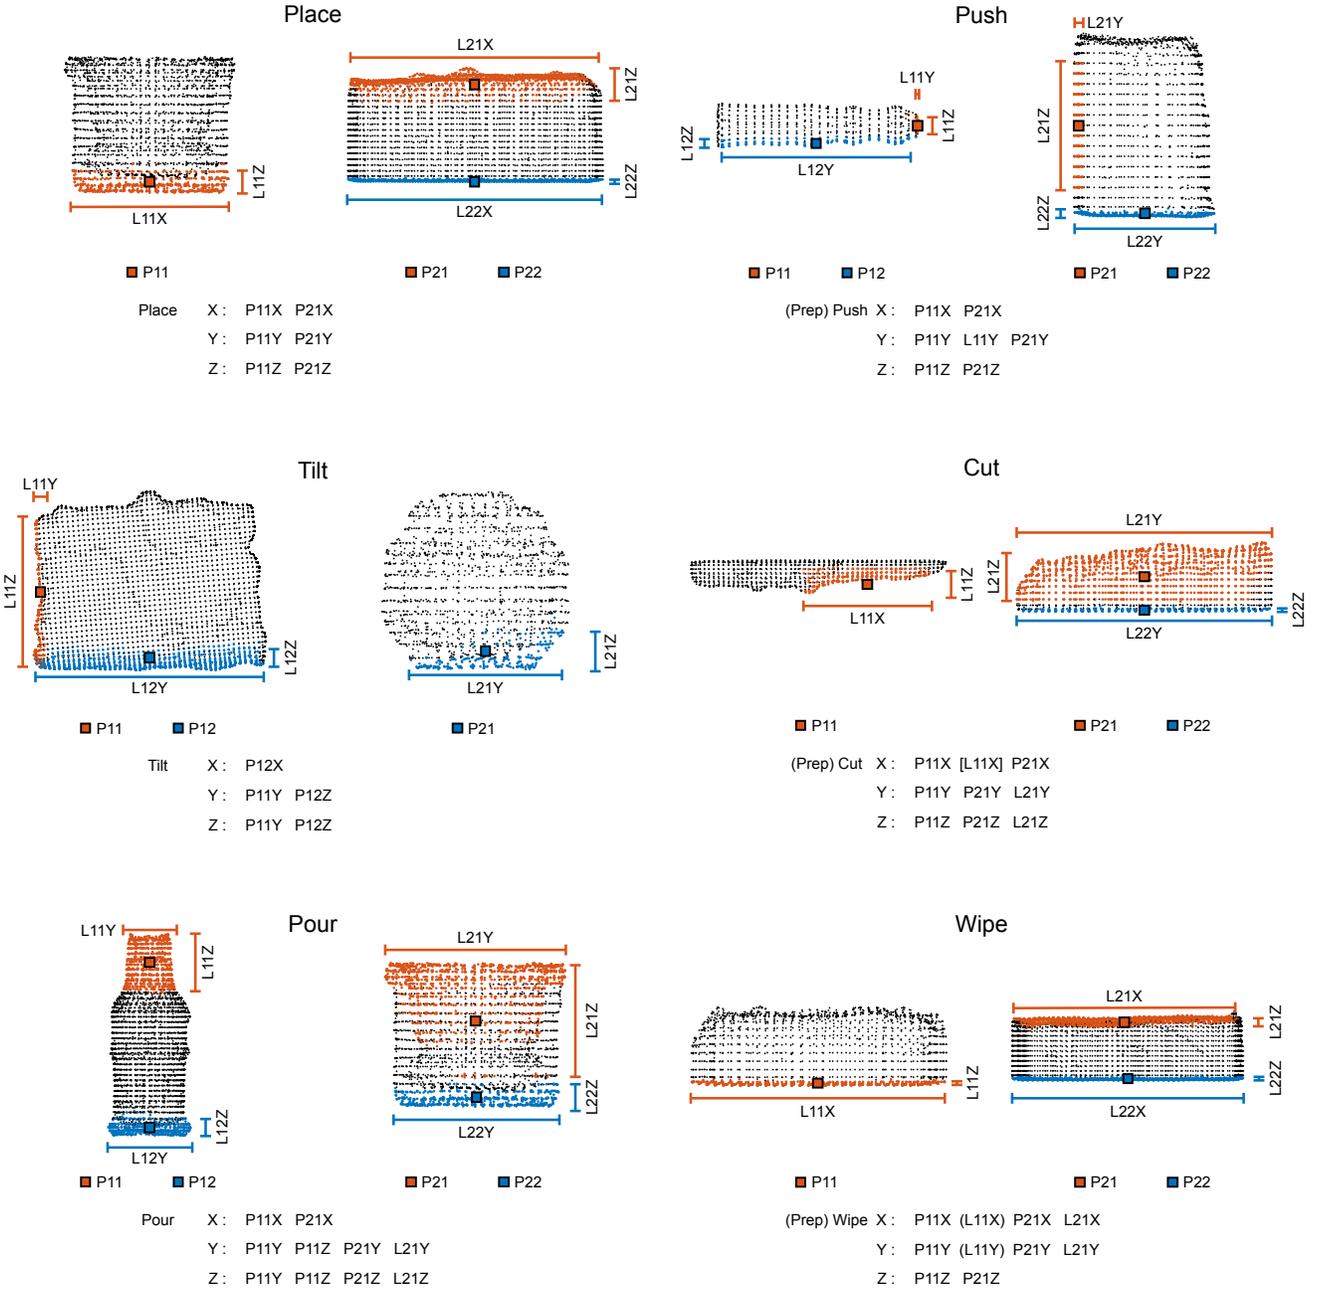
Figure 8. Parts of objects and their corresponding features. The features are aligned with the task frame. The third object in each task is not shown The feature naming convention indicates position/length, object number, part number, and x-y-z dimension. The human selected features for each skill components are listed below the picture. Features in square brackets only apply to the manipulation skill and features in parenthesis only apply to the preparation skill

B. Benchmarking Priors for Feature Selection

Given the extracted parts, the next step is to select the relevant object features for learning the versatile manipulation skill. The goal of this experiment is to evaluate using a meta-level prior to transfer knowledge about feature relevance between different skills. This benchmark experiment therefore compares the performance of the SSVS when using a uniform prior versus a meta prior.

Rather than learning one motor primitive per task, the pushing, wiping, and cutting tasks were each split into two skills. The first skill corresponds to moving the held object into contact with the second object. The second skill then performs the pushing, wiping, or cutting manipulation. We indicate the first skill by *prep* (e.g., cut prep) as it is used to prepare the objects for the manipulation [25]. The set of relevant object features and the meta features' values may vary between the skills and the corresponding prep skills.

The relevant features were selected using the SSVS approach described in Section IV. After the temporal scaling of the target DMPs, the rotational components of the skills exhibited little variation between the different demonstrations, even for the pouring and tilting tasks. These components were therefore modeled using a single constant feature, and our feature selection evaluation was applied to only the three

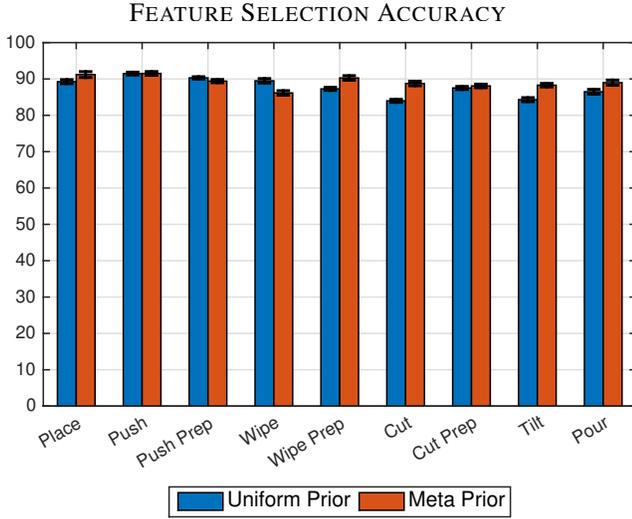## FEATURE SELECTION ACCURACY

FEATURE SELECTION RECALL

Figure 9. The figure shows the accuracy of the SSVS feature selection process. The blue bar indicates the accuracy when using a standard uniform prior over the features' relevance. The red bar shows the the accuracy when using a meta prior that was trained on the feature relevance of the other tasks' skills. The error bars indicate +/- one standard deviation.
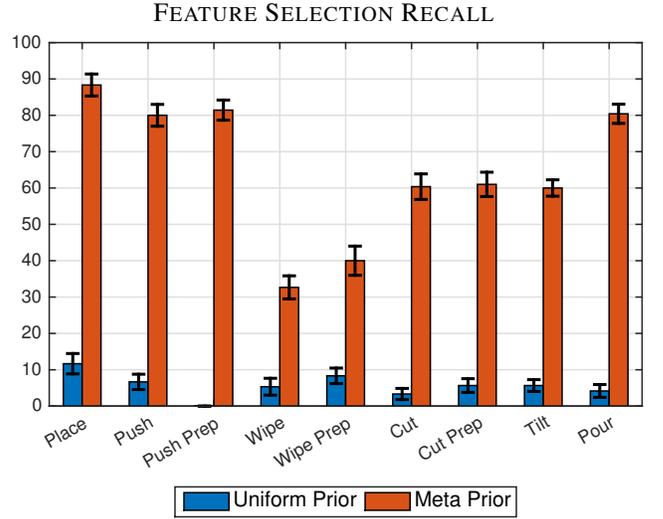
Figure 10. The figure shows the recall of the SSVS feature selection process. The blue bar indicates the recall when using a standard uniform prior over the features' relevance. The red bar shows the the recall when using a meta prior that was trained on the feature relevance of the other tasks' skills. The error bars indicate +/- one standard deviation.

translational components of the skill. The hyperparameters were set to $a = 5$ and $b = 5$ for the prior on the standard deviations σ_k. For the shape parameters w_{jk}, we set the hyperparameters to $\hat{s}^2 = 0.15^2$ and $\check{s}^2 = 50\hat{s}^2$, such that the prior variances on the relevant features are fifty times greater than those of the irrelevant features. The DMPs were learned using five shape parameters per feature per skill component, resulting in an average of $5 \times 27 \times 3 = 405$ parameters to be learned per skill. As part of the preprocessing, both the target shape parameters $\tilde{w}_{ik}$ and the object features ϕ_{ji} were normalized across the N training samples.

The meta prior's parameters θ_h were trained using iterative reweighted least squares. The training data was obtained from the other tasks' skills, i.e., all skills excluding the current skill and the corresponding prep skill when applicable. The relevant features of the prior skills were hand-labelled for the purpose of this evaluation. The manually selected features for each skill are shown in Fig. 8. In ambiguous cases, we selected the features that minimized the leave one out cross validation error for the trajectory prediction.

The meta-prior training set was created by randomly selecting half of the positive samples $\gamma = 1$ and an equal number of negative samples $\gamma = 0$ from the pool of previous skills' features. The feature relevance prior, as described in Section IV-B, can be computed from a single demonstration. In order to compute the prior for the entire skill, we took the average over the priors computed from the meta features of the individual training demonstrations. The uniform prior assumes that each feature has the same a priori probability of being selected, i.e., $p(\gamma_j) = c \forall j \in \{1, ..., M\}$ where c is a constant. The constant probability was obtained by computing the average number of relevant features using all of the features from the previous skills. Over all of the skills, 9.97% of the features are considered to be relevant.

The set of relevant features were estimated using the Gibbs
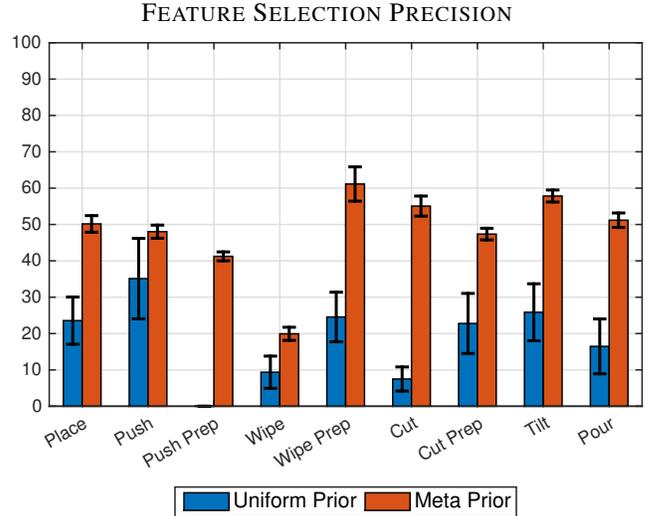
FEATURE SELECTION PRECISION

Figure 11. The figure shows the precision of the SSVS feature selection process. The blue bar indicates the precision when using a standard uniform prior over the features' relevance. The red bar shows the the precision when using a meta prior that was trained on the feature relevance of the other tasks' skills. The error bars indicate +/- one standard deviation.

sampling approach outlined is Section IV-C. For each trial, the robot randomly selected $N = 5$ demonstrations from the current skill. The SSVS approach was applied to these five trials. All of the relevance parameters were initialized to one $\gamma_j = 1$. We used 200 samples for the burn-in stage of sampling, followed by another 1000 samples for the actual sampling. Using a Bayes estimator approach, a feature was considered relevant iff it was assigned a value of $\gamma_j = 1$ for the majority of the samples. These estimates of the relevant features were then compared to the ground truth human labels. The results for using the meta prior, as well as the uniform prior are shown in Fig. 9, 10, and 11. The evaluations were repeated ten times for each skill, with different training
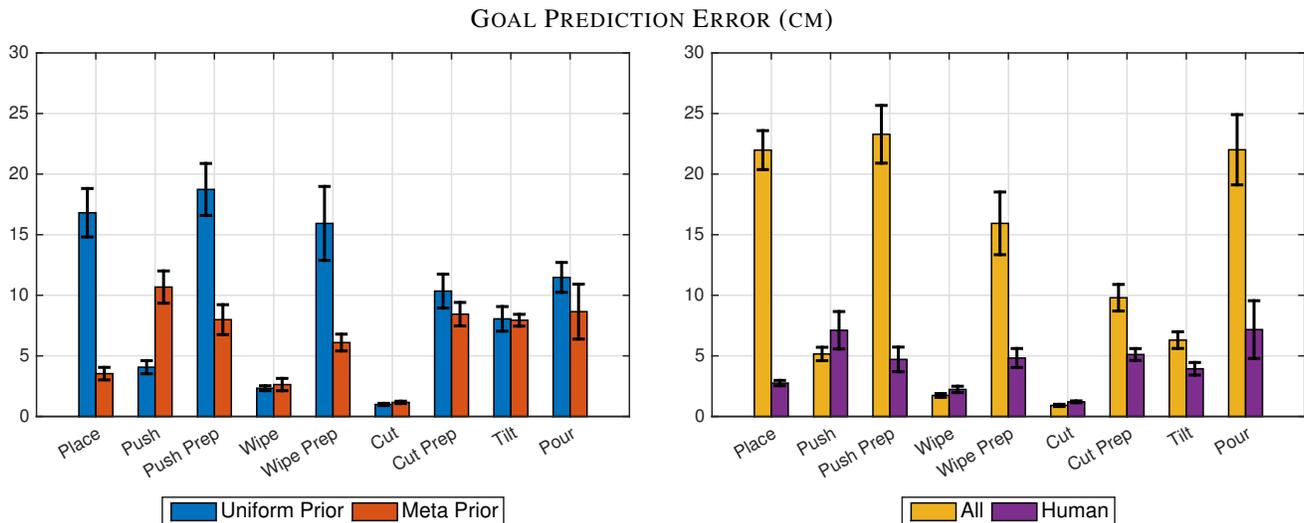
GOAL PREDICTION ERROR (CM)

Figure 12. The plots show the average errors in centimeters for the goal prediction evaluation. The left plot shows the results when using SSVS with a uniform prior and a meta-level prior respectively. The plot on the right shows the two baseline methods of using all of the features and human selected features. The errorbars indicate +/- one standard error.

sets for the meta priors and demonstrations for the feature selection. The errorbars indicate the standard deviations over these evaluations.

Both of the priors result in high accuracies for the feature selection, with the meta prior approach having an average accuracy of 89.2%. However, due to the sparsity of the relevant features, a high accuracy can be achieved by treating all of the features as irrelevant $\gamma_j = 0 \forall j \in \{1, ...M\}$. The resulting manipulation skills would however not generalize between different situations. We therefore need to consider the precision and recall to determine how well the two models capture the relevant features.

The recall is shown in Fig. 10. While the meta prior approach achieves an average recall of 64.9%, the standard uniform prior only has an average recall of 5.7%. The recall values can be improved by including additional demonstrations. This result indicates that the meta-level prior guides the sampling process to capture most of the relevant features. In contrast, given only a few demonstrations, many of the relevant features are missed when using a uniform prior. A high recall is particularly important if the robot subsequently uses reinforcement learning to optimize the skill, as the robot is less likely to explore variations in features that it already considers irrelevant.

The benefit of the meta prior can also be seen in the precision, as it increase the average precision from 18.36% to 48.9%. The precision values are lower due to the lack of variations in the demonstrations. For example, the vertical position of the bottom of the grasped object is a relevant feature for the place task. However, as all of the objects are initially resting on the table, the bottoms of all three objects are at the same height. As a result, the Gibbs sampling often considered these features to be relevant as well, which allowed it to use them to estimate the height of the table more accurately. However, if the task were performed with the objects placed on different surfaces, then these additional feature would lead to worse generalization. The precision values could therefore

be improved through more samples, either from additional demonstrations or through self exploration, in novel scenarios. Once the robot has mastered a skill, the set of selected features could be used to update the meta prior for future tasks.

As a benchmarking experiment, the results show that the meta-level prior more than doubles the precision and recall of the feature selection compared to the standard uniform prior. The results thus show that the meta prior can autonomously transfer prior relevance knowledge between distinct skills.

C. Goal State Prediction

In our reformulation of the DMPs, the goal state of the movement is defined as a linear combination of the extracted object features. The robot is therefore also selecting features for predicting the goal state of the movement. In this experiment, we evaluated the robot's ability to predict the goal state of the skill from the object features. We again compare using a uniform prior to the proposed meta-level prior. We additionally compare to using all of the features, as well as using the manually selected features.

In each trial, one of the demonstrations was selected as the test sample, and $N = 5$ other demonstrations from the skill were randomly selected as the training samples. For the uniform prior and meta prior approaches, the robot selected a relevant set of object features for generalizing the skills using SSVS, as described in Section IV. The robot then used the selected features to predict the goal state of the test DMP. The root mean squared error (RMSE) between the predicted and the demonstrated goal states were computed to evaluate the approaches.

The results of the experiment are shown in Fig. 12. We have included the results for all of the tasks for completeness, although not all of the tasks require a specific goal location. For example, the pushing distance varied randomly between demonstrations. The average errors over all of the skills are 4.33cm for human-selected features, 6.34cm for meta-prior
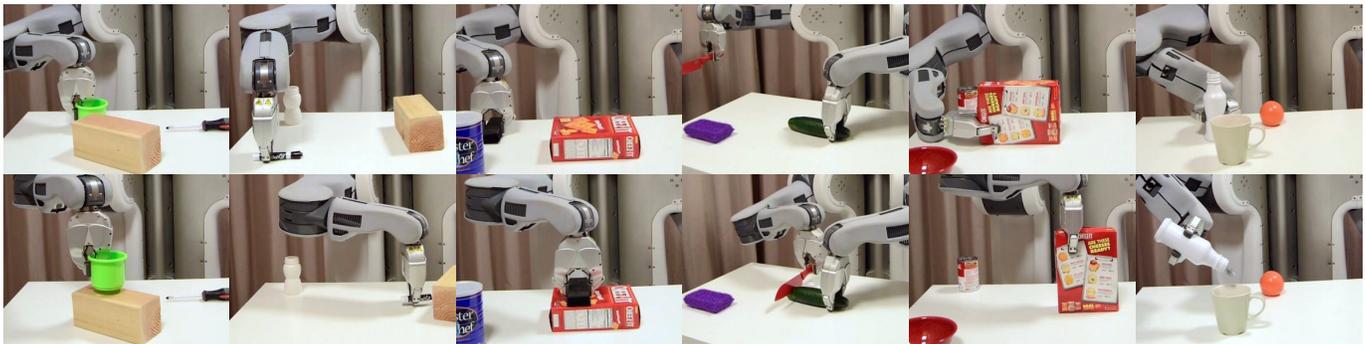
Figure 13. The pictures show the robot's skill executions. The top row shows the initial states for the task. The bottom row shows the resulting manipulations being performed. From left to right, the columns show executions from the placing, pushing, wiping, cutting, tilting, and pouring tasks

features, 9.86cm for uniform-prior features, and 11.9cm for all of the features.

The meta and uniform priors achieved similar results for the wiping, cutting, and tilting tasks. These tasks require relatively small movements and the skills therefore do not need to adapt to large variations in the object features. The meta prior outperformed the uniform prior for the placing, pouring, and prep skills. These skills involve moving the grasped object into close proximity to a part of another object. In addition to requiring larger movements, the object features for these tasks also exhibit more variations as the initial scenes are less constrained than in the cutting, wiping, and tilting tasks. As a result, selecting the wrong features can lead to larger errors. The meta-level prior guides the robot in selecting the relevant features given the limited number of training demonstrations.

The benefit of the meta-level prior is also greater for these tasks due to the similarities in the skills' relevant features, i.e., the relevant features for these skills are generally axis-aligned and in close proximity to the hand at the end of the demonstration. The meta-level prior is therefore better at predicting the relevant features as it has more training samples. In contrast, the tilting task includes relevant features across dimensions, e.g., the y component of the skill depends on the z position of a part. Pouring is the only other skill from the set that includes these kinds of relevant features. As a result, the benefit of the meta-level prior may be smaller for these skills, but it could be increased by learning additional skills.

Pushing is the only skill for which the meta-level prior performed substantially worse than the uniform prior. The manually selected features also performed worse for this task. This result demonstrates that an erroneous prior can also decrease performance. The meta-level prior and the human both expected certain features to be relevant, and the robot therefore attempted to model the random variations in the pushing distance based on these features. The influence of the prior will decrease as the robot obtains more training samples.

As one would expect given the limited amount of training data, the manually selected features have the best performance. The meta-prior approaches performance could be further improved by using additional meta-features to differentiate between relevant and irrelevant features, as well as learning from more previous skills. Overall, the experiments have shown that the meta-level prior, in comparison to the uniform prior or no

feature selection, can substantially increase the performance of skills learned from only a few demonstrations.

D. Learning Manipulation Skills

In the final experiment, the learned skills were executed on the real robot to verify that the skills can adapt to different situations and perform the intended manipulations. The skills were trained using the relevant set of features learned using the meta prior. The skills were learned using all $N = 15$ demonstrations. In order to test the skills with different object features, the evaluations were performed using new combinations of objects, in varying positions and orientations, and with different grasps. Each skill was executed three times. If a prep skill failed in such a manner that the subsequent skill could not be performed, then it was marked as a failure and another execution was performed to evaluate the following manipulation skill.

Examples of the skill executions can be seen in Fig. 13, as well as in the attached video. The video also includes examples of failed skill executions. The robot achieved an overall success rate of 77.8% during the evaluations. Learning these skills from $N = 15$ demonstrations is not a trivial task, as each skill had 375 to 465 parameters to be learned. Therefore, although the skills capture the coarse movements that allow the robot to perform the manipulations, the executions could be further refined and made more robust. In particular, the amplitudes of the pushing and cutting movements tended to be slightly too large, the pouring was performed off-center, the tilting resulted in horizontal shifts of the boxes, and half of the failed executions were the result of prep skills stopping a couple of centimeters before making contact. These issues could be alleviated by adjusting the robot's trajectories by a few centimeters.

This skill refinement could be achieved through additional self exploration using reinforcement learning. As the skills only require minor adjustments, the robot could employ a policy search approach [38, 32]. The skill executions could also be refined by adding visual/tactile/haptic feedback into the DMPs [39, 6, 55]. These feedback terms can be easily incorporated as features ϕ that depend on the current state rather than the initial state. The feature selection process could thus also be extended to select relevant feedback terms.

The experiments have shown that the robot can use the SSVS approach together with a meta-level prior to efficiently learn versatile manipulations from relatively few demonstrations. Having demonstrated the feasibility and benefits of this approach, we plan on extending the skill learning framework in the future to allow for different action frames and larger variations in object orientations. We will also investigate explicitly incorporating the interaction sites in the feature generation process. These interactions, which are currently used to initialize the part segmentation, could provide additional details for generalizing manipulation skills.

VI. Conclusion

In this paper, we presented a learning-from-demonstration method for selecting the relevant object features for generalizing manipulation skills. The proposed approach is based on stochastic search variable selection (SSVS). We extend the standard SSVS model to incorporate a meta-level prior, which allows the robot to transfer knowledge regarding feature relevance between different skills. The transfer was performed using meta features that capture the proximity and alignment of the object features. We also explained how the GrabCut segmentation method can be used to extract affordance-bearing parts of objects from demonstrations. The extracted parts formed the basis for the automatic feature generation.

The proposed method was evaluated using a PR2 robot. The robot learned placing, pushing, cutting, wiping, tilting, and pouring skills that generalize between different sets of objects and object configurations. The results show that the meta prior allows the robot to more than double the feature selection's average precision and recall compared to a uniform prior.

References

[1] E. Aksoy, A. Abramov, J. Dörr, K. Ning, B. Dellen, and F. Wörgötter. Learning the semantics of object-action relations by observation. *International Journal of Robotic Research (IJRR)*, 30(10):1229–1249, 2011.

[2] J. Bohg, A. Morales, T. Asfour, and D. Kragic. Data-driven grasp synthesis - a survey. *IEEE Transactions on Robotics (TRo)*, 2014.

[3] A. Boularias, O. Kroemer, and J. Peters. Learning robot grasping from 3d images with markov random fields. In *International Conference on Intelligent Robot Systems (IROS)*, 2011.

[4] S. Calinon, T. Alizadeh, and D. G. Caldwell. On improving the extrapolation capability of task-parameterized movement models. In *International Conference on Intelligent Robots and Systems (IROS)*, 2013.

[5] B. Çalli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar. Benchmarking in manipulation research: The YCB object and model set and benchmarking protocols. *Computing Research Repository*, abs/1502.03143, 2015.

[6] Y. Chebotar, O. Kroemer, and J. Peters. Learning robot tactile sensing for object manipulation. In *International Conference on Intelligent Robots and Systems (IROS)*, 2014.

[7] F. R. K. Chung. *Spectral Graph Theory*. American Mathematical Society, 1997.

[8] B.C. da Silva, G. Baldassarre, G.D. Konidaris, and A.G. Barto. Learning parameterized motor skills on a humanoid robot. In *International Conference on Robotics and Automation (ICRA)*, 2014.

[9] M. P. Deisenroth, P. Englert, J. Peters, and D. Fox. Multi-task policy search for robotics. In *International Conference on Robotics and Automation (ICRA)*, 2014.

[10] D. Fox. Kld-sampling: Adaptive particle filters. In *Advances in Neural Information Processing Systems (NIPS)*. MIT Press, 2001.

[11] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 6(6):721–741, November 1984.

[12] E. I. George and R. E. McCulloch. Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889, 1993.

[13] J. J. Gibson. *The Ecological Approach To Visual Perception*. Lawrence Erlbaum Associates, 1986.

[14] C. Goldfeder, P.K. Allen, C. Lackner, and R. Pelossof. Grasp planning via decomposition trees. In *International Conference on Robotics and Automation (ICRA)*, 2007.

[15] K. Hausman, C. Bersch, D. Pangercic, S. Osentoski, Z. C. Marton, and M. Beetz. Segmentation of cluttered scenes through interactive perception. In *ICRA Workshop on Semantic Perception and Mapping for Knowledge-enabled Service Robotics*, 2012.

[16] T. Hermans, J. M. Rehg, and A. F. Bobick. Guided pushing for object singulation. In *International Conference on Intelligent Robots and Systems (IROS)*, 2012.

[17] A. Ijspeert, J. Nakanishi, and S. Schaal. movement imitation with nonlinear dynamical systems in humanoid robots. In *International Conference on Robotics and Automation (ICRA)*, pages 1398–1403, 2002.

[18] T. Jebara and R. Kondor. Bhattacharyya and expected likelihood kernels. In *Conference on Learning Theory (COLT)*, Lecture Notes in Computer Science, 2003.

[19] M. Kalakrishnan, L. Righetti, P. Pastor, and S. Schaal. Learning force control policies for compliant manipulation. In *International Conference on Intelligent Robots and Systems (IROS)*, 2011.

[20] C.C. Kemp, A. Edsinger, and E. Torres-Jara. Challenges for robot manipulation in human environments. *IEEE Robotics Automation Magazine*, 14(1):20–29, 2007.

[21] J. Kober, K. Muelling, O. Kroemer, C.H. Lampert, B. Schoelkopf, and J. Peters. Movement templates for learning of hitting and batting. In *International Conference on Robotics and Automation (ICRA)*, 2010.

[22] G.D. Konidaris, S.R. Kuindersma, R.A. Grupen, and A.G. Barto. Autonomous skill acquisition on a mobile manipulator. In *Conference on Artificial Intelligence (AAAI)*, 2011.

[23] O. Kroemer and J. Peters. Predicting object interactions from contact distributions. In *International Conference on Intelligent Robots and Systems (IROS)*, 2014.

[24] O. Kroemer, E. Ugur, E. Oztop, and J. Peters. A

kernel-based approach to direct action perception. In *International Conference on Robotics and Automation (ICRA)*, 2012.

[25] O. Kroemer, C. Daniel, G Neumann, H. van Hoof, and J. Peters. Towards learning hierarchical skills for multi-phase manipulation tasks. In *International Conference on Robotics and Automation (ICRA)*, 2015.

[26] E. Krupka, A. Navot, and N. Tishby. Learning to Select Features using their Properties. *Journal of Machine Learning Research (JMLR)*, 9:2349–2376, October 2008.

[27] A.G. Kupcsik, M.P. Deisenroth, J. Peters, L. Ai Poh, V. Vadakkepat, and G. Neumann. Model-based contextual policy search for data-efficient generalization of robot skills. *Artificial Intelligence*, 2015.

[28] S. R. Lakani, M. Popa, A. J. Rodríguez-Sánchez, and J. H. Piater. CPS: 3d compositional part segmentation through grasping. In *Conference on Computer and Robot Vision (CRV)*, pages 117–124, 2015.

[29] K. H. Lee, J. Lee, A.L. Thomaz, and A.F. Bobick. Effective robot task learning by focusing on task-relevant objects. In *International Conference on Intelligent Robots and Systems (IROS)*, 2009.

[30] S. Lee, V. Chatalbashev, D. Vickrey, and D. Koller. Learning a meta-level prior for feature relevance from multiple related tasks. In *International Conference on Machine Learning (ICML)*, pages 489–496. ACM, 2007.

[31] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies. In *ArXiv*, 2015.

[32] S. Levine, N. Wagener, and P. Abbeel. Learning contact-rich manipulation skills with guided policy search. In *International Conference on Robotics and Automation (ICRA)*, 2015.

[33] J. Ernesti M. Do, J. Schill and T. Asfour. Learn to wipe: A case study of structural bootstrapping from sensorimotor experience. In *International Conference on Robotics and Automation (ICRA)*, 2014.

[34] J. H. Metzen, A. Fabisch, L. Senger, J. de Gea Fernandez, and E. A. Kirchner. Towards learning of generic skills for robotic manipulation. *German Journal of Artificial Intelligence*, 28:15–20, 2014.

[35] D. Munoz, N. Vandapel, and M. Hebert. Onboard contextual classification of 3-D point clouds with learned high-order markov random fields. In *International Conference on Robotics and Automation (ICRA)*, 2009.

[36] S. Niekum. An integrated system for learning multi-step robotic tasks from unstructured demonstrations. In *Conference on Artificial Intelligence (AAAI)*, 2013.

[37] A. L. Pais, K. Umezawa, Y. Nakamura, and A. Billard. Learning robot skills through motion segmentation and constraints extraction. HRI Workshop on Collaborative Manipulation, 2013.

[38] P. Pastor, M. Kalakrishnan, S. Chitta, E. Theodorou, and S. Schaal. Skill learning and task outcome prediction for manipulation. In *International Conference on Robotics and Automation (ICRA)*, 2011.

[39] P. Pastor, M. Kalakrishnan, L. Righetti, and S. Schaal. Towards associative skill memories. In *International Conference on Humanoid Robots (Humanoids)*, 2012.

[40] J. Peters and S. Schaal. Policy gradient methods for robotics. In *International Conference on Intelligent Robotics Systems (IROS)*, 2006.

[41] J. Piater, S. Jodogne, R. Detry, D. Kraft, N. Krueger, O. Kroemer, and J. Peters. Learning visual representations for perception-action systems. *International Journal of Robotics Research (IJRR)*, 2011.

[42] R. B. Potts. Some generalized order-disorder transformations. *Mathematical Proceedings of the Cambridge Philosophical Society*, 48:106–109, 1 1952.

[43] S.J.D. Prince. *Computer Vision: Models Learning and Inference*. Cambridge University Press, 2012.

[44] C. Rother, V. Kolmogorov, and A. Blake. GrabCut -interactive foreground extraction using iterated graph cuts. *ACM Transactions on Graphics (SIGGRAPH)*, 2004.

[45] R. B. Rusu and S. Cousins. 3D is here: Point Cloud Library (PCL). In *International Conference on Robotics and Automation (ICRA)*, 2011.

[46] R. B. Rusu, A. Holzbach, M. Beetz, and G. Bradski. Detecting and segmenting objects for mobile manipulation. In *ICCV, S3DV Workshop*, 2009.

[47] S. Schaal, J. Peters, J. Nakanishi, and A. Ijspeert. learning movement primitives. In *International Symposium on Robotics Research (ISRR)*. Springer, 2004.

[48] S.C. Stein, F. Wörgötter, M. Schöler, J. Papon, and T. Kulvicius. Convexity based object partitioning for robot applications. In *International Conference on Robotics and Automation (ICRA)*, 2014.

[49] M. Stolle, H. Tappeiner, J. E. Chestnutt, and C. G. Atkeson. Transfer of policies based on trajectory libraries. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 2981–2986, 2007.

[50] J. Stuckler and S. Behnke. Adaptive tool-use strategies for anthropomorphic service robots. In *International Conference on Humanoid Robots (Humanoids)*, 2014.

[51] J. Sung, S. H. Jin, and A. Saxena. Robobarista: Object part-based transfer of manipulation trajectories from crowd-sourcing in 3D pointclouds. In *International Symposium on Robotics Research (ISRR)*, 2015.

[52] B. Taskar, M. F. Wong, and D. Koller. Learning on the test data: Leveraging unseen features. In *International Conference on Machine Learning (ICML)*, pages 744–751. AAAI Press, 2003.

[53] M. Tenorth, S. Profanter, F. Balint-Benczedi, and M. Beetz. Decomposing CAD models of objects of daily use and reasoning about their functional parts. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 5943–5949, 2013.

[54] H. van Hoof, O Kroemer, and J. Peters. Probabilistic segmentation and targeted exploration of objects in cluttered environments. *IEEE Transactions on Robotics (TRo)*, (5): 1198–1209, 2014.

[55] A. Yamaguchi, C. G. Atkeson, and T. Ogasawara. Pouring skills with planning and learning modeled from human demonstrations. *International Journal of Humanoid Robotics (Humanoids)*, 12(3), 2015.