

Active Articulation Model Estimation through Interactive Perception

Karol Hausman¹

Scott Niekum^{2,3}

Sarah Osentoski⁴

Gaurav S. Sukhatme¹

Abstract—We introduce a particle filter-based approach to representing and actively reducing uncertainty over articulated motion models. The presented method provides a probabilistic model that integrates visual observations with feedback from manipulation actions to best characterize a distribution of possible articulation models. We evaluate several action selection methods to efficiently reduce the uncertainty about the articulation model. The full system is experimentally evaluated using a PR2 mobile manipulator. Our experiments demonstrate that the proposed system allows for intelligent reasoning about sparse, noisy data in a number of common manipulation scenarios.

I. INTRODUCTION

When operating in human environments, robots must frequently cope with multiple types of uncertainty such as systematic and random perceptual error, ambiguity due to a lack of data, and changing environmental conditions. While methods like Bayesian reasoning allow robots to make intelligent decisions under uncertainty, in many situations it is still advantageous, or even necessary, to actively gather additional information about the environment. The paradigm of *interactive perception* aims to use a robot’s manipulation capabilities to gain the most useful perceptual information to model the world and inform intelligent decision making.

In this work, we leverage interactive perception to directly model and reduce the robot’s uncertainty over articulated motion models. Many objects in human environments move in structured ways relative to one another; articulation models describe these movements, providing useful information for both prediction and planning [1], [2]. For example, many common household items like staplers, drawers, and cabinets have parts that move rigidly, prismatically, and rotationally with respect to each other.

Previous works on detection of articulated motion models have generally used a passive maximum likelihood strategy to fit models, ignoring uncertainty [1]–[3]. However, passive, fit-based approaches can often lead to several problems. First, these methods give no indication if there are competing high-likelihood hypotheses, or if the fitted model is the only reasonable explanation of the data. Second, purely observational models only look at the fit of observed data points, but do not reason about rigidity constraints or certainty, leading them to be tricked in many situations. Third, without




Fig. 1. Top: The PR2 performing an action that reduces the uncertainty over articulation models. Bottom: Three of the most probable articulation model hypotheses: free-body(green), rotational(blue) and prismatic(pink). Selected action is depicted in yellow and lies on the plane of the cabinet

a confidence metric, the robot cannot reason about how to gather additional data intelligently when needed.

In this paper, we address all of the above challenges by tracking a distribution over articulation models, updating the model distribution based upon the outcome of manipulation actions, and selecting informative actions to converge quickly to accurate articulation models. An example of our method in action is shown in Fig. 1.

The key contributions of our approach are that we: a) develop a particle filter approach to keep track over uncertainty over different articulation models and their parameters, b) design a manipulation sensor model that updates model likelihoods based on the feedback from the robot’s manipulator, and c) introduce a probabilistic action selection algorithm that reduces uncertainty efficiently. The effectiveness of this approach is demonstrated through a series of experiments using a PR2 mobile manipulator.

¹Karol Hausman and Gaurav S. Sukhatme are with the Department of Computer Science, University of Southern California, Los Angeles, CA 90089, USA. hausman@usc.edu

^{2,3}Scott Niekum is with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA and the School of Computer Science, University of Massachusetts Amherst, Amherst, MA 01020, USA.

⁴Sarah Osentoski is with Robert Bosch Research and Technology Center, Palo Alto, CA 94304, USA.

II. RELATED WORK

Interactive perception has enjoyed success in many different applications, including object segmentation [4]–[6], object recognition [7], [8], object sorting [9], [10], and object search [11]. Some classes of problems have state information that cannot be observed passively, requiring the use of interactive perception. One instance of this is the estimation of kinematic models of articulated objects.

The task of estimating articulation models has been addressed in various ways. Sturm et al. [1] use human demonstrations to obtain visual data from tracked markers, which is then used to characterize single or multiple joints in a probabilistic fashion. Katz et al. [12], [13] and Pillai et al. [3] were able to move away from a marker-based system by using visual features to detect and track rigid parts. The observed motion of these rigid bodies is then used to characterize various joint types. Huang et al. [14] use different camera view-points and make use of structure-from-motion techniques to recover and classify the joint properties. In their recent work, Martin and Brock [2] present an algorithm that is able to perform articulation model estimation on the fly, in contrast to previous offline algorithms. Jain and Kemp [15] present an alternate approach in which an equilibrium point control algorithm is used to open doors and drawers without extensive use of visual data.

Otte et al. [16] present related work that addresses the problem of physical exploration to estimate the articulation model. However, the problem is formulated differently in this case—the robot is searching for a high-level strategy (in what order should the robot interact with objects?) in order to estimate the structure of the whole environment.

In our work, we probabilistically combine visual observations with the outcomes of the robot’s manipulation actions in a recursive state estimation framework. We use fiducial markers to simplify the perceptual challenge, as the main focus is on informative action selection and the representation of multiple hypotheses. By using this approach, it is possible to improve any of the above-mentioned visual articulation model frameworks by taking advantage of the robot’s manipulation capabilities to estimate the underlying model in a principled way.

III. APPROACH

A. Articulation Model Representation

Our approach takes advantage of robot manipulation capabilities to actively reduce the uncertainty over articulation models and their parameters. This requires fusing different sensor modalities—vision-based object tracking and outcomes of manipulation actions performed by the robot. Both sensor measurements are fused using a recursive Bayesian update formulation which is implemented as a particle filter [17] (Fig. 2).

Articulation Models and their Parameters: We consider four types of articulation models: rigid, prismatic, rotational and free-body, parametrized similarly to Sturm et al. [1]. The rigid model has 6 parameters specifying a fixed relative




Fig. 2. Particles representing different hypothesis of the underlying articulation model. Blue is a rotational model, pink is a prismatic model and green is a free-body model.




Fig. 3. Probabilistic graphical model for articulation recognition.

transformation between two objects or parts; the prismatic model has 2 more parameters than the rigid model that indicate the direction of the prismatic axis; the rotational model has 3 more parameters than the rigid model that indicate the axis of rotation and the radius; the free-body model does not contain any parameters, as the objects can move freely with respect to each other.

Probabilistic Graphical Model for State Estimation: Our estimate of the articulation model can be represented as the probabilistic graphical model shown in Fig. 3, where we assume that the model and its parameters do not change over time. The model consists of the type of articulation model $\mathcal{M}$, its parameters θ , sets of 6-DOF object pose measurements $\mathbf{z}_t^m$ which are represented in a plate notation (there exist N_t measurements for each time step t), action sensor measurements z_t^a and actions A_t .

Given this probabilistic graphical model the posterior probability of the articulation model and its parameters can be factorized:

$$P = p(\mathcal{M}, \theta | z_{1:t}^a, \mathbf{z}_{1:t}^m, A_{1:t-1}) \propto \underbrace{p(\mathbf{z}_t^m | \mathcal{M}, \theta)}_{\text{visual sensor model}} \underbrace{p(z_t^a | \mathcal{M}, \theta, A_{t-1})}_{\text{action sensor model}} \underbrace{p(\mathcal{M}, \theta | z_{1:t-1}^a, \mathbf{z}_{1:t-1}^m, A_{1:t-2})}_{\text{previous belief}}. \quad (1)$$

This formulation enables us to fuse all the sensor measurements coming from vision as well manipulation data. In the following subsections, we describe the respective parts of the above factorization.

B. Visual Sensor Model

In order to evaluate how the observed poses of the objects can be explained by the model and its parameters, a visual sensor model is needed. Having obtained the current noisy measurement z_t^m , similarly to Sturm et al. [1], it is projected onto the model that is being evaluated. The projection is achieved by using the model's inverse kinematics function:

$$\hat{\mathbf{q}} = f_{\mathcal{M},\theta}^{-1}(z_t^m). \quad (2)$$

The resulting configuration estimate $\hat{\mathbf{q}}$ is the configuration (i.e. a scalar joint position or rotational angle in the prismatic and rotational models, and a constant for the rigid model) to the original observation z_t^m that lies on the surface of the model. This configuration is then used with the model's forward kinematics function to obtain the 6-DOF pose corresponding to the projected configuration:

$${}^{\mathcal{M}}\hat{z}_t^m = f_{\mathcal{M},\theta}(\hat{\mathbf{q}}). \quad (3)$$

Finally, the visual sensor model is approximated with a 2-dimensional multivariate normal distribution (Eq. 5) that is centered around 2-dimensional vector of zeros and depends on the translational and rotational distance d (Eq. 4) between the projected ${}^{\mathcal{M}}\hat{z}_t^m$ and the observed z_t^m measurement. Since the visual measurements are conditionally independent given the model and its parameters, we are able to write the final sensor model:

$$d(z_t^m, {}^{\mathcal{M}}\hat{z}_t^m) := [trans_d(z_t^m, {}^{\mathcal{M}}\hat{z}_t^m), rot_d(z_t^m, {}^{\mathcal{M}}\hat{z}_t^m)] \quad (4)$$

$$p(\mathbf{z}_t^m | \mathcal{M}, \theta) = \prod_{z_t^m \in \mathbf{z}_t^m} \mathcal{N}(d(z_t^m, {}^{\mathcal{M}}\hat{z}_t^m); \mathbf{0}, \Sigma), \quad (5)$$

where Σ is the estimated covariance of the observations. In the case of a free-body model, the above formulas do not hold, as it does not have forward and inverse kinematics functions defined. The likelihood update for the free-body model corresponds to a constant probability density value that is defined by a uniform distribution over all possible poses in the robot's workspace. This follows from the intuition that a free-body is not constrained by any joint and is roughly equally likely to be anywhere in the workspace.

C. Action/Manipulation Sensor Model

The second sensor modality that is incorporated in the model is the outcome of manipulation actions. The actions are defined as short translational movements performed on a scale of 10cm, where a stable grasp of the articulated object is assumed. Each action is represented as a vector ${}^{\mathcal{W}}\mathbf{V}_t^a$ in the world frame $\mathcal{W}$.

The observed quantity in the manipulation sensor model is defined as a binary variable that is equal to 0 if the manipulation was unsuccessful (the manipulator cannot complete the motion because the joint friction is too high, it attempted motion in a rigid direction, etc.) and 1 otherwise (the expected motion is observed).

For each model, its parameters, and the selected action, we can compute the likelihood of the action being successful as:

$$p(z_t^a | \mathcal{M}, \theta, A_{t-1}) = e^{-\gamma\beta}, \quad (6)$$




Fig. 4. The angle β between the action vector ${}^{\mathcal{W}}\mathbf{V}_t^a$ and the tangent vector to the current model ${}^{\mathcal{W}}\mathbf{V}_t^m$ (rotational in this case).

where γ is the weighting factor dependent on stiffness of the arm controller used in the robot; β is the angle between the vector ${}^{\mathcal{W}}\mathbf{V}_t^m$ describing the motion prediction at the current grasping point ${}^{\mathcal{W}}z_t^m$ for a given model $\mathcal{M}$ in the world frame $\mathcal{W}$ and the action vector ${}^{\mathcal{W}}\mathbf{V}_t^a$ that is the direction of the movement of the manipulator. It worth noting that one can use other functions to estimate the likelihood of the action being successful as long as the likelihood is inversely proportional to the angle β . To estimate ${}^{\mathcal{W}}\mathbf{V}_t^m$, the most recent noisy observation ${}^{\mathcal{W}}z_t^m$ is used as a point estimate of the pose of the articulated part at the grasping location. We introduce a function f that, given the model $\mathcal{M}$ and the current estimate of the pose ${}^{\mathcal{W}}z_t^m$, returns a vector that is a tangent to the predicted motion path of the object at the current point (the blue vector depicted in Fig. 4):

$${}^{\mathcal{W}}\mathbf{V}_t^m = f(\mathcal{M}, {}^{\mathcal{W}}z_t^m) \quad (7)$$

$$\beta = \arccos \left(\frac{{}^{\mathcal{W}}\mathbf{V}_t^m \cdot {}^{\mathcal{W}}\mathbf{V}_t^a}{\|{}^{\mathcal{W}}\mathbf{V}_t^m\| \|{}^{\mathcal{W}}\mathbf{V}_t^a\|} \right). \quad (8)$$

This formulation results in a decrease of the likelihood of successful motion as the angle between predicted motion vector and action vector increases. In the case of the rigid and the free-body models, we assume that the only reason for $z_t^a = 0$ is the configuration of the articulated object (there are no singularities and the robot's arm movement is in its workspace), which leads to the definition of the manipulation sensor model:

$$p(z_t^a = 0 | \mathcal{M}^{rigid}, \theta, A_{t-1}) = 1 \quad (9)$$

$$p(z_t^a = 1 | \mathcal{M}^{rigid}, \theta, A_{t-1}) = 0 \quad (10)$$

$$p(z_t^a = 0 | \mathcal{M}^{free}, \theta, A_{t-1}) = 0 \quad (11)$$

$$p(z_t^a = 1 | \mathcal{M}^{free}, \theta, A_{t-1}) = 1. \quad (12)$$

D. Best Action Selection

In order to select the best action (e.g. the action that reduces the uncertainty over the model the most), we compare two different strategies.

Entropy-based Action Selection: In this case, we define the best action as the action that will maximally reduce the entropy over the model distribution in expectation. Because the outcome z_{t+1}^a of the future action A_t has not been observed,

we compute the expected entropy, H , of the distribution over articulation models after the action is executed:

$$A_t^* = \underset{A_t}{\operatorname{argmin}} \mathbb{E}_{z_{t+1}^a \sim p(z_{t+1}^a | \mathcal{M}, \theta)} H[\mathcal{M}, \theta | Z_{1:t+1}^a, z_{1:t}^m, A_t]. \quad (13)$$

Intuitively, the action that maximally reduces the entropy should disambiguate the current hypotheses of articulation models in the most efficient way. However, as we will show later, this intuition is incorrect in many situations.

Information Gain-based Action Selection: Our second approach is based on an information gain metric, rather than on entropy itself. Similar to entropy-based action selection, first, the expected posterior is calculated:

$$Q = \mathbb{E}_{z_{t+1}^a \sim p(z_{t+1}^a | \mathcal{M}, \theta)} p(\mathcal{M}, \theta | Z_{1:t+1}^a, z_{1:t}^m, A_t). \quad (14)$$

Next, we look for the action that will maximize the information gain between the current distribution and the expected distribution, given the action that is being evaluated. Information gain is defined as the Kullback-Leibler divergence [18] between these two distributions:

$$A_t^* = \underset{A_t}{\operatorname{argmax}} D_{KL}(P || Q). \quad (15)$$

IV. IMPLEMENTATION

A. Particle Filter

Since the joint distribution of a model and its parameters has no simple parametric form and may often be multimodal, we approximate it using a particle filter. Each particle represents an articulation model type (e.g. rigid, free-body, rotational, or prismatic) and the associated model parameters. An example showing the state of the particle filter is depicted in Fig. 2. Since the particle filter has to cover a high dimensional space, we propose a particle initialization algorithm that ensures that the highest likelihood areas of the joint distribution are well-covered.

The action selection algorithm is then applied after updating the particle weights according to the visual sensor model. For particle resampling, stratified resampling [19] is used because it better ensures that the full support of the underlying distribution is covered adequately.

Initialization of Particles: We take advantage of human demonstrations or prior observation data in order to initialize the particles, though this step is not strictly required. An equal number of particles are used for each articulation model. In order to fit model parameters to the data, the algorithm from Sturm et. al [1] is employed, in which the authors use MLESAC (maximum likelihood sample consensus) [20] to fit and refine parameters. For more details we refer the reader to [1]. It is important to emphasize that a small number of MLESAC iterations are used to estimate the parameters of the model. This ensures that more of the high-dimensional parameter space is covered, rather than just the highest likelihood settings, given the noisy visual data.




Fig. 5. Automatically generated candidate actions.

B. Visual Sensor Model

We record all available object pose data during demonstrations, as well as during robot manipulation actions. Every time a correction step of the particle filter step occurs, we make use of all the new data that was not previously used to update the weights of the particles according to Eqn. 5.

C. Manipulation Sensor Model

In our robot experiments, we use a PR2 mobile manipulator that does not contain a force/torque sensor to directly measure the forces in the robot's wrist. In order to obtain manipulation feedback from the robot, we implemented a compliant controller that stops whenever the controller effort exceeds a certain threshold. In this way, we are able to detect the sensor output z_t^a that specifies if the interaction with the articulation model was successful (i.e. if the part moved as expected or resisted motion).

D. Best Action Selection

To select the optimal action, a set of candidate actions is first generated. The actions are sampled such that they cover the range of directions that the robot can act in order to move the articulated object from a known grasp point. One set of candidate actions is depicted in Fig. 5. Next, either the expected entropy (post-action) or the relative entropy (compared to pre-action) of the particle filter is calculated for each generated action. Finally, based on the chosen optimality criteria, the action optimal with respect to the selected criteria is chosen.

Entropy-based Action Selection: In order to estimate the entropy of the distribution represented in the particle filter, taking only particles' weights into account is not sufficient. For example, if the particles are concentrated around the true estimate of the articulation model, their weights will be similar, making the entropy large. To solve this problem, the position of the particles must be taken into account by estimating the underlying differential entropy of the distribution.

For each model we perform kernel density estimation [21] over its parameters with Gaussian kernels:

$$\hat{f}_{\mathbf{H}}(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n K_{\mathbf{H}}(\boldsymbol{\theta} - \boldsymbol{\theta}_i) \quad (16)$$

$$K_{\mathbf{H}}(\boldsymbol{\theta}) = |\mathbf{H}|^{-1/2} K(\mathbf{H}^{-1/2} \boldsymbol{\theta}) \quad (17)$$

$$K(\boldsymbol{\theta}) = (2\pi)^{-d/2} e^{-1/2 \boldsymbol{\theta}^T \boldsymbol{\theta}}. \quad (18)$$

The bandwidth matrix $\mathbf{H}$ is estimated according to Scott's rule [22]:

$$\mathbf{H}_{ij} = 0, i \neq j \quad (19)$$

$$\sqrt{\mathbf{H}_{ii}} = n^{-\frac{1}{d+4}} \sigma_i, \quad (20)$$

where σ_i is the standard deviation of the i -th variable, d is the dimensionality of the parameters, and n is the number of samples.

By sampling the kernel density estimate at the (weighted) locations of the original particles, we are able to estimate the differential entropy of the joint distribution represented in the particle filter with the resubstitution estimate proposed in [23]:

$$\mathbf{H} = -\frac{1}{n} \sum_{i=1}^n \ln \hat{f}_{\mathbf{H}}(\boldsymbol{\theta}_i). \quad (21)$$

At the end of this process, the expected entropy of the joint distribution of models and their parameters is estimated as the sum of the entropies (according to Eq. 21) of all the model types for each action. The action with the smallest expected entropy is considered optimal with respect to the entropy-based criteria.

Information Gain Action Selection: In this case, we look for the action that will maximize the information gain (e.g. relative entropy) of the joint distribution approximated with the particle filter. Information gain is defined as the Kullback-Leibler divergence between the current distribution and the expected distribution post-action:

$$IG(P, Q) = D_{KL}(P||Q) = \sum_i P(i) \ln \frac{P(i)}{Q(i)}. \quad (22)$$

It is worth noting that in this case, kernel density estimation is not needed since the algorithm can operate on the particle weights directly, making it more computationally efficient than the previous method.

V. EVALUATION AND DISCUSSION

A. Experimental Setup

In the following experiments, we use a PR2 mobile manipulator with a Microsoft Kinect sensor attached to its head. Since object recognition is not the focus of this work, we use RGB-D data to track the pose of objects in the world with AR tags, a type of visual fiducial. All code used in the following experiments is open source and available online as a ROS package¹.

¹http://wiki.ros.org/active_articulation



Fig. 6. Four experimental scenarios. Top-left: rotational cabinet door with an incomplete demonstration; top-right: two free-body erasers moved apart in a straight line; bottom-left: locked drawer as a rigid joint; bottom-right: stapler as a rotational joint.

We designed four experimental scenarios in which the visual demonstration data leaves room for ambiguity about the articulation model type and parameters. In the first scenario, the underlying model is rotational and involves a cabinet door, in which the demonstration consists of a single opening and closing of the door to an angle of approximately 30 degrees. The second scenario is a free-body model and consists of a whiteboard eraser that the demonstrator moves along a straight line with respect to another eraser. The third scenario consists of a rigid model and involves the robot initializing the filter without a demonstration by attempting to open a locked drawer, tracking only the fiducial markers on this static object. The final scenario is a rotational model and consists of a stapler that is dynamically moved around the scene during the demonstration. Fig. 6 shows all four scenarios.

For the visual sensor model, we set Σ to a diagonal matrix with $0.04m$ for the translational component and 30 degrees for the rotational component based on empirical data. Throughout the experiments, the number of MLESAC iterations was set to 3 when initializing the particle filter.

B. Action Selection

In order to evaluate our system, we first compare different action selection strategies on the cabinet door scenario. After choosing the most efficient action selection strategy, we perform the remainder of our experiments in all four scenarios. In all experiments, 500 particles were used to cover the space of the probability distribution.

Evaluation: We evaluate three different action selection strategies: minimum expected entropy and maximum KL-divergence as described in Sec. III-D, and random action selection. In the cabinet door scenario, each action selection method was tested by having the robot perform 7 actions with that method. The results are averaged over 5 runs of the experiment. We measure the success of each action selection strategy by its reduction in uncertainty (i.e. minimizing the entropy over models) and by the number of particles




Fig. 7. Top: Entropy over articulation model types after each action for the cabinet door experiment. Bottom: Probability of the true articulation model (rotational) after each action.

representing the correct model—a rotational joint in this case. In order to eliminate the influence of unsuccessful action execution (e.g. the gripper slips off the door handle) on the action strategy comparison, we simulate the output of the manipulation sensor model by analyzing data from executions where manipulation performed well; empirically, we found that the manipulation action was generally successful when the angle between the action vector and the tangent to the model was less than 30 degrees. Using this, we simulated a binary success criteria based on a threshold of 30 degrees.

Results: Fig. 7 presents the obtained entropy over articulation model types and probability of the correct model after each action. It can be seen that the KL-divergence action selection method reduces entropy much faster than the other evaluated methods. Somewhat surprisingly, the expected entropy strategy is almost as slow as selecting an action randomly.

To gain some insight into this, recall that the goal of maximizing the KL-divergence is to change the distribution as much as possible, whereas minimizing the expected entropy aims to obtain as ‘peaky’ of distribution as possible. In the

presented scenarios, it is the case that in order to minimize the entropy of the articulation models, the algorithm needs to increase it initially to find the right model, which eventually, leads to a significant reduction of the entropy. In other words, the minimum entropy strategy fails because it reduces entropy greedily at each step and gets stuck in local minima; by contrast, the information gain approach aims to change the distribution as much as possible in any direction, allowing escape from minima. One can observe this effect in the bottom graph of Fig. 7. After one action, the entropy-based method yields better results (in terms of the probability of the correct model as well as the entropy) than the KLD-based algorithm, but in all following steps, the KLD-based algorithm outperforms the other methods.

Fig. 8 emphasizes the difference between the two presented action selection methods, showing the probability distribution and chosen actions for both methods. Although both methods chose a similarly good first action, the KLD algorithm selects a much better second action. The entropy-based method chose an action that is perpendicular to the motion predictions of two competing articulation models, resulting in a minimal entropy reduction. The KLD approach selects an action that shifts the probability density over articulation models towards rotational—a model type with a larger entropy in its parameters. This causes the overall entropy of the joint distribution to increase temporarily. In the long run, it is desirable to move belief towards the correct model type, so that in the next step, the entropy over its parameters can be reduced. However, an entropy reduction method is unlikely to select this action because of the 1-step entropy increase.

C. Articulation Model Estimation

Evaluation: We now use the superior KL-divergence action selection method to test the overall performance of our system. For each of the previously described experimental scenarios, the algorithm is run 5 times and the resulting statistics are averaged over all the runs. A different number of actions were taken in each scenario based on the convergence rate of the algorithm.

Results: The second row of Fig. 8 shows the distribution of the articulation model types in the cabinet door experiment before any action was taken (based only on the visual data from the demonstration), after the first action, and after the second action, respectively. One can see that the algorithm has shifted most of the probability mass to the correct model (rotational) after only two actions. The first action (in the door plane, Fig. 8, top-left) was selected in order to eliminate the free-body model hypothesis while the second action (towards the normal of the door plane, Fig. 8, top-right) was chosen to distinguish between the prismatic and rotational hypotheses, as they act on different angles.

It is worth noting that once the rotational articulation model type has converged, the standard deviation of each of the parameters of that model was smaller than 0.04 which demonstrates that both the model type and its parameters have converged.




Fig. 8. Outcomes of performing actions on the cabinet door. Top-left: the current state of the particle filter with most probable hypotheses of each articulation model. The yellow arrow shows the optimal action in terms of maximizing the information gain. Top-right: the state of the particle filter after performing the first action ($z_1^a = 0$, door did not move), together with the currently chosen most informative action. Second row-left: Percentage of each of the models' particles in the particle filter before any action, after the first action, and after the second action ($z_2^a = 1$, door did move). The two bottom rows depict the same idea, but the action is selected based on minimization of the expected entropy. The outcomes of the actions in this case were both $z_2^a = 1$.

Fig. 9 presents a similar comparison for the other two experiments. As before, the first column of the graph shows the averaged distribution before any action was taken, and the consecutive columns show the results after each respective action. In case of the eraser experiment, it is apparent that one action was enough to converge on the correct free-body model (Fig. 9, top row). This is because the KL-divergence action selection algorithm chose the action that is perpendicular to the previously demonstrated linear motion. Since the eraser can move freely in this direction, the prismatic hypothesis is pruned out and the free-body model is the only hypothesis that can explain both the visual and the manipulation data. The rigid model is not taken into consideration since the demonstration movements were too large to support this hypothesis. Note that a passive, fit-based




Fig. 9. Top: Probability distribution over models before and after the action in the eraser experiment. Bottom: Analogous statistics for the rigid-joint experiment. Results of 3 actions are presented.




Fig. 10. Rotational particles in the stapler experiment before (left) and after (right) a manipulation action.

method would be tricked in this scenario by almost certainly fitting a prismatic model to the visual data.

The rigid-body experiment required three actions to shift the majority of the belief to the correct model (Fig. 9, bottom row). The explanation of this behavior lies in the measurement noise of the fiducial marker tracking. Since the marker appears to slightly oscillate in the demonstration, the particle initialization algorithm is able to fit many different models to the data, including prismatic joints with axes in many different directions. Manipulation actions help to change this initial belief by moving along the axis of the most likely prismatic hypotheses. Since the manipulation output confirms that there is no movement in this direction, the prismatic hypotheses are pruned out. However, due to a large number of different prismatic hypotheses, many actions are needed to be certain that it is actually a rigid model.

The final experiment illustrates uncertainty over model parameters caused by systematic observation noise, rather than over model type. Fig. 10 shows the state of the particle filter in the stapler scenario before and after an action was taken. In this case, the articulation model type is clear from the demonstration (rotational joint), but there is significant ambiguity in the parameter space of this model. Since there is systematic noise in the marker detection (resulting from the changing viewing angle of the markers), it is crucial to use the particle filter in order to model all potential articulation

model hypotheses and use manipulation to prune out the incorrect ones. In this experiment, one action was sufficient in order for the algorithm to converge.

VI. CONCLUSIONS AND FUTURE WORK

We introduced a probabilistic interactive perception system to represent and reduce uncertainty over various types of articulated motion models. The system includes a visual sensor model to deal with perceptual noise, a manipulation model to probabilistically reason about the outcomes of actions, and an action selection metric that allows for efficient reduction of uncertainty over models and parameters.

Experiments demonstrated, somewhat counter-intuitively, that the most efficient way of reducing the entropy of the distribution over candidate models and parameters is *not* achieved by greedily minimizing the expected entropy after each action. Instead, an information gain approach based on KL-divergence was shown to be significantly more effective than minimizing entropy or performing random actions.

We then showed that the full system can actively cope with ambiguity and reduce uncertainty over models and their parameters in a variety of situations. The PR2 was able to actively manipulate objects to disambiguate a limited cabinet door demonstration; to prevent being fooled by prismatic motion between two free-body erasers; to eliminate alternate hypotheses of a rigid locked drawer caused by random visual noise; and to narrow the distribution of rotational parameters of a stapler that were corrupted by systematic perceptual noise.

There are several clear areas for future work in this domain. At present, we restrict our method to look for the best action over a one-step horizon. It may be useful to investigate this problem with a further lookahead in the future. Unfortunately, in this work we could not use direct force measurements due to the lack of force-torque sensors on the PR2 robot. Nevertheless, our probabilistic framework allows to easily integrate these measurements, which is a possible extension of the current system. We also plan to move away from using initial demonstrations, instead taking advantage of previously learned priors in order to initialize the particles. Additionally, removing fiducial markers and using visual features to track objects (similarly to [2]) is a natural extension of this work. Finally, it is desirable to be able to apply this method to more sophisticated joint types like helical screws, loose joints, and compositions of multiple joints, allowing the use of this method in more complex tasks, such as assembly.

ACKNOWLEDGEMENT

This work was funded in part by the NSF under grant IIS-1208497. We would like thank Jörg Müller for his very helpful input.

REFERENCES

- [1] J. Sturm, C. Stachniss, and W. Burgard, "A probabilistic framework for learning kinematic models of articulated objects." *J. Artif. Intell. Res. (JAIR)*, vol. 41, pp. 477–526, 2011.
- [2] R. M. Martn and O. Brock, "Online interactive perception of articulated objects with multi-level recursive estimation based on task-specific priors," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014.
- [3] S. Pillai, M. Walter, and S. Teller, "Learning articulated motions from visual demonstration," in *Proceedings of Robotics: Science and Systems*, Berkeley, USA, July 2014.
- [4] H. van Hoof, O. Kroemer, and J. Peters, "Probabilistic segmentation and targeted exploration of objects in cluttered environments," 2014. [Online]. Available: <http://www.ias.tu-darmstadt.de/uploads/Publications/hoof2014probabilistic.pdf>
- [5] K. Hausman, F. Balint-Benczedi, D. Pangercic, Z.-C. Marton, R. Ueda, K. Okada, and M. Beetz, "Tracking-based interactive segmentation of textureless objects," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 1122–1129.
- [6] J. Kenney, T. Buckley, and O. Brock, "Interactive segmentation for manipulation in unstructured environments," in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*. IEEE, 2009, pp. 1377–1382.
- [7] D. Schiebener, J. Morimoto, T. Asfour, and A. Ude, "Integrating visual perception and manipulation for autonomous learning of object representations," *Adaptive Behavior*, vol. 21, no. 5, pp. 328–345, 2013.
- [8] K. Hausman, C. Corcos, J. Mueller, F. Sha, and G. S. Sukhatme, "Towards interactive object recognition," in *3rd Workshop on Robots in Clutter: Perception and Interaction in Clutter, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Chicago, IL*, September 2014.
- [9] L. Chang, J. R. Smith, and D. Fox, "Interactive singulation of objects from a pile," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 3875–3882.
- [10] M. Gupta and G. Sukhatme, "Using manipulation primitives for brick sorting in clutter," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 3883–3889.
- [11] M. Gupta, T. Ruhr, M. Beetz, and G. S. Sukhatme, "Interactive environment exploration in clutter," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013, pp. 5265–5272.
- [12] D. Katz, M. Kazemi, J. Andrew Bagnell, and A. Stentz, "Interactive segmentation, tracking, and kinematic modeling of unknown 3d articulated objects," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 5003–5010.
- [13] D. Katz, A. Orthey, and O. Brock, "Interactive perception of articulated objects," in *Experimental Robotics*. Springer, 2014, pp. 301–315.
- [14] X. Huang, I. Walker, and S. Birchfield, "Occlusion-aware reconstruction and manipulation of 3d articulated objects," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 1365–1371.
- [15] A. Jain and C. C. Kemp, "Pulling open doors and drawers: Coordinating an omni-directional base and a compliant arm with equilibrium point control," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*. IEEE, 2010, pp. 1807–1814.
- [16] S. Otte, J. Kulick, M. Toussaint, and O. Brock, "Entropy-based strategies for physical exploration of the environments degrees of freedom," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014.
- [17] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT press, 2005.
- [18] S. Kullback and R. A. Leibler, "On information and sufficiency," *The Annals of Mathematical Statistics*, pp. 79–86, 1951.
- [19] G. Kitagawa, "Monte carlo filter and smoother for non-gaussian nonlinear state space models," *Journal of computational and graphical statistics*, vol. 5, no. 1, pp. 1–25, 1996.
- [20] P. H. Torr and A. Zisserman, "Mlesac: A new robust estimator with application to estimating image geometry," *Computer Vision and Image Understanding*, vol. 78, no. 1, pp. 138–156, 2000.
- [21] M. Rosenblatt *et al.*, "Remarks on some nonparametric estimates of a density function," *The Annals of Mathematical Statistics*, vol. 27, no. 3, pp. 832–837, 1956.
- [22] D. W. Scott, "On optimal and data-based histograms," *Biometrika*, vol. 66, no. 3, pp. 605–610, 1979.
- [23] I. Ahmad and P.-E. Lin, "A nonparametric estimation of the entropy for absolutely continuous distributions (corresp.)," *Information Theory, IEEE Transactions on*, vol. 22, no. 3, pp. 372–375, 1976.