

# Interactive Affordance Map Building for a Robotic Task

David Inkyu Kim

Gaurav S. Sukhatme

**Abstract**—We describe a technique to build an affordance map interactively for robotic tasks. Affordances are predicted by a trained classifier using geometric features extracted from objects. Based on 2D occupancy grid, a Markov Random Field (MRF) model builds an affordance map with relational affordance with neighboring cells. The quality of the affordance map is refined by sequences of interactive manipulations selected from the model to yield the highest reduction in uncertainty.

## I. INTRODUCTION

Consider a robot tasked to rearrange disorganized boxes (Fig. 1). Without prior knowledge of the environment, the robot must explore the world and identify target objects that need to be moved. It must also identify objects that it can move.

Object affordances [1] describe relative functionalities between the agent (robot) and the objects: a cup can be “graspable” and a box on the floor can be “pushable”. There are few particular characteristics of affordance. First, affordance depends upon the agent’s capability. A cup may be graspable to a humanoid robot with a manipulator arm, but it may not be to a wheeled mobile robot. Some boxes can easily be pushed by a big robot but not by a smaller robot. Second, affordance may vary with time depending on the relative pose of an object to the robot. A free box may be “pushable”, but it will not be if a wall is behind the box. Defining and understanding object affordance is clearly useful in the sense that such knowledge can be brought to bear directly on manipulation.

In order to recognize affordances, a robot first needs to utilize its sensors to perceive the environment. Input sensor data are processed to recognize object identities and corresponding affordances. However, visual cues are not sufficient to fully understand affordances for lack of interaction: whether a box is empty or not is often unknown unless the situation is disambiguated by the robot trying to push the box. Therefore, visual perception and interactive manipulation are both necessary for affordance modeling.

In this paper, we propose an algorithm to semantically map affordances for robotic tasks as follows: 1) explore and build a map of the environment, 2) predict affordances of the map based on sensor observations, 3) plan a sequence of interactive manipulations to evaluate and confirm the predicted affordances and reduce overall uncertainty of the affordance map, 4) show the experimental result on simulated warehouse.

The authors are with the Robotic Embedded Systems Laboratory, Dept. of Computer Science, University of Southern California, Los Angeles, CA, U.S.A. davidink@usc.edu, gaurav@usc.edu




Fig. 1: A PR2 robot situated in a simulated warehouse with disorganized boxes (left). In order to rearrange boxes at desired positions, the robot must push the boxes. But some boxes cannot be pushed because they are blocked either by walls or other boxes. An affordance represents “pushability” in a certain direction. In this paper, we explore how to learn an affordance map via interactive manipulation in order to plan a rearrangement (right).

## II. RELATED WORK

When a robot is given a task in an unknown environment, it must understand the configurations and characteristics of the objects to accomplish the task. For navigation toward a goal among movable obstacles [2], knowledge about physical characteristics of objects is necessary. Wu *et al.* [3] proposed a navigation algorithm by interacting with obstacles to learn characteristics of the objects. They used greedy approach to select an object to test whether it is static. Many techniques have been developed to reveal characteristics of objects like physical properties, colors, shapes, etc. However, object affordance, differing from general properties of objects, is a higher level abstraction and is directly suitable for manipulation.

Koppula *et al.* [4] learned human activities and object affordance from RGB-D videos. A human performing an action, like drinking, relates affordance label to the interacting object, like a cup. So an affordance could be extracted by observing human activities and modeling the relationship between a human and the object. Varadarajan *et al.* [5] built an affordance network to recognize unknown objects. By searching for the part with a specific affordance (*e.g.* a handle for graspable), a novel object could be recognized with the matched affordance. These affordances are focused on understanding objects in a human context. Here, we interpret affordances of objects from the robot’s view and use them to accomplish robotic tasks.

Among many approaches for recognizing and learning about unknown characteristics of objects, there has been prior research on utilizing interaction between the robot and the environment. Chu *et al.* [6] used haptic feedbacks from

touching objects to learn haptic adjectives. Gupta *et al.* [7] studied interactive manipulation in the cluttered environment to plan rearrangement of objects to maximize expectation of a hidden target object being revealed. Van Hoof *et al.* [8] used informative gaining strategy to choose a robot action for better segmentation based of probabilistic distribution over segmentation. Ugur *et al.* [9] studied ‘traversability’ affordance with interactive learning, which is useful for a robot path planning.

Our research differs from previous work in that we obtain task specific affordances for a robot and build an affordance map. Instead of applying simple pre-assigned motions for interaction, we plan a manipulation sequence based on a classifier that predicts affordances. We use a MRF model to estimate the consequences of manipulation to maximize information gain from such interaction.

### III. AFFORDANCE PREDICTION AND MANIPULATION PLANNING

#### A. Defining affordance

Defining affordance for a robotic task is complicated since affordance has semantic meaning [10]. Four components need to be considered in defining affordance: When an *agent* perform a *behavior* on an *entity*, an *effect* occurs. Behavior (robot action) and effect (interaction result) should be described with parameters to utilize affordance in robotic application. For example, an affordance “pushable” can be expressed with parameters of direction and magnitude. Affordance can also be sampled from the distribution, considering uncertainty and errors in perception and manipulation. In our work, we define affordance as pushability over 4 direction for a unit grid. Simplified parameters of affordance enables predicting to be tractable. Also, derived affordance can be directly applied to the planning. Details about affordance derivation will be described in Section III-C.

#### B. Building a grid map and collecting visual data

In the work reported here, a PR2 robot with a RGB-D camera is used to obtain a 3D point cloud based on which a 2D occupancy grid map of the environment is created. Our discrete affordances can be utilized on this simplified 2D grid world. Here we assume a PR2 robot is situated in a simulated warehouse with randomly distributed boxes. First, the robot explores the world to build a grid map. To find out the occupancy of grid cells, we use a ray tracing technique (Fig. 3 left). As the robot traverses the environment and discovers unknown occupied cells, each incoming point cloud from a single view (Fig. 2 left) is merged into the global point cloud. Registration of point clouds is aided by an Extended Kalman Filter based on odometry data. A frontier-based approach is used to explore the map by setting a frontier on the boundary between known and unknown regions. The robot is driven through all approachable regions without any physical interaction with objects and covers the map. As a result, we get multiple point cloud snapshots, a merged and fully registered 3D world point cloud (Fig. 3 right) and discretized 2D occupancy grid map.




Fig. 2: An example point cloud of a single view (left). One example of the warehouse setting with randomly placed boxes (right).




Fig. 3: Obtained 2D occupancy grid using ray tracing. Black means occupied, light gray means empty, and dark gray means unknown (left). Registered 3D world point cloud by stitching all obtained point clouds with EKF (right).

#### C. Affordance prediction based on geometric features

Previously, we studied the semantic labeling of 3D point cloud with object affordance to predict object affordances based on RGB-D data [11]. In terms of object recognition, visual features like colors and textures are useful. In terms of affordance, the paper has been shown that geometric appearances are useful, explicitly expressing affordance of an object even without previous knowledge of an object identity. Here, we adapt the classifier that predicts affordance by utilizing geometric features. Point cloud segmentation is required to extract segments and to label them with an affordance. A segment does not necessarily correspond to a single object since the goal is to derive affordance, not object identity. Here, we use voxel segmentation as it clearly confers the advantage of discretized affordances.

As studied in [11], there are two types of geometric features, unary and pairwise. Unary features represent geometric properties of a single voxel segment while pairwise features describe geometric relationship between neighboring voxels. In the following, we describe geometric features used in the paper and provide the affordance information carried within each feature. Further details about feature extraction are described in [11].

- Saliency describes the shape of the point cloud by scatterness, linearity, and planarity. It is computed by the method inspired by tensor voting [12].
- Normals represent the planar direction where affordance for pushing can be inferred.
- Span and centroid imply how far and where each

segment reaches in the XYZ coordinates.

- Accessibility influences affordance as well. If sensor reading shows that there is no obstacle between the robot and the object, spatial knowledge of occupancy is used to estimate accessibility.
- Hierarchical Matching Pursuit 3D (HMP3D) features are extracted from point cloud for object recognition [13]. In each voxelized point cloud, a dictionary is learned by the K-SVD algorithm to express a point cloud with a sparse vector. The sparse vectors are encoded with hierarchical matching pursuit where feature components are derived from the point cloud voxels with various scales. In our work, we apply a simplified feature vectors by collecting surface normals of voxels, each with different sizes.
- Occupancy of the neighboring voxel segments can be encoded in pairwise features. For example, a single object in a free space is pushable to any nearby empty cells. But two contiguously positioned objects cannot be pushed towards each other unless a robot is strong enough to push both together. Also, if a wall is blocking an object, the object cannot be pushed towards the wall. By looking at neighboring voxel segments, we can estimate the influence of neighbors to affordances.

With the extracted unary and pairwise geometric features, a classifier is trained to predict affordance labels. We define the robotic task affordances as follows: pushable up, down, left, and right in 2D. This notation is useful as a robot can directly use the affordances as pushing primitives. In addition to the direction of the affordance, the affordance parameter about how far an object can be pushed needs to be defined. For simplicity and lower complexity, we set up affordance unit values to be one grid cell size.

Among many classifiers, we choose a logistic regression classifier for its simplicity and fast computation. We trained the classifier with the sampled voxel segments along with geometric features. These point cloud examples are gathered by driving the PR2 around the simulated warehouse among various object configurations. Through training, we attain parameters  $\beta^l = \{\beta_0^l, \beta_1^l, \dots, \beta_m^l\}$ , where  $\beta_i^l$  stands for  $i$ th parameter over  $m$  features for affordance label  $l$ . The logistic regression is expressed as follows:

$$p^l(\mathbf{f}) = \frac{e^{g^l(\mathbf{f})}}{e^{g^l(\mathbf{f})} + 1}, \quad g^l(\mathbf{f}) = \beta_0^l + \sum_{i=1}^m \beta_i^l f_i \quad (1)$$

where  $\mathbf{f} = \{f_1, f_2, \dots, f_m\}$  and  $f_i$  stand for the  $i$ th feature value of the segment. A function  $g^l(\mathbf{f})$  is a linear combination of parameter  $\beta^l$  and features. An outcome of prediction  $p^l(\mathbf{f})$  is computed from the logit function of  $g^l(\mathbf{f})$  and the resulting prediction value lies between 0 and 1, expressing how probable an object is to have a particular affordance.

#### D. Manipulation planning for affordance mapping

Based upon the predicted affordance, we need a plan to verify the affordance of each occupying cell by interactive manipulation. Here, we consider a simple manipulation

primitive—pushing. After pushing, we need to validate the affordance whether the prediction was correct. For validation, one might acquire haptic feedback from the end effector [6] or check the torque feedback from the mobile base wheels since pushing an overweight object might result in slippage. In our approach, we perform change detection of the point cloud for verification. After executing one unit of pushing in a given direction, two candidate point clouds are compared, one being the hypothesized state of the environment after the unit push and the other being the observed point cloud. By computing average Euclidean distance error between expected point cloud and observed point cloud, the affordance can be confirmed.

In order to complete affordance mapping, a strategy is needed to plan a sequence of manipulations and execute it. The basic approach would be a naive one, which examines all occupied cells exhaustively. This is expensive since a robot must travel all around the map and manipulate (push) every occupied cell. We decrease the number of manipulations by estimating reduction in uncertainty resulting from each manipulation. The uncertainty of the affordance map can be expressed by entropy where lower entropy indicates higher certainty. The planning strategy searches for manipulation that reduces entropy the most. Additionally, we predict the influence of manipulation by adapting a Markov Random Field (MRF). A probabilistic graphical structure is constructed by setting each occupied cell to be a node and connecting edges between the neighboring cells.

$$P(x_1, x_2, \dots, x_n) = \frac{1}{Z} \prod_{i=1}^n \psi_i(x_i) \prod_{i,j} \psi_{ij}(x_i, x_j) \quad (2)$$

The probability distribution of occupied cells' affordances  $P(x_1, x_2, \dots, x_n)$  is computed from the product of unary and pairwise potentials where  $n$  is the number of occupied cells and  $Z$  is the partition function. The unary potential  $\psi_i(x_i)$  is a prior, which is the predicted affordance from the classifier. The pairwise potential  $\psi_{ij}(x_i, x_j)$  encodes the mutual influence between cell  $x_i$  and  $x_j$ . We define the pairwise potential function as follows:

$$\psi_{ij}(x_i, x_j) = \exp(Jx_i x_j)$$

where  $J > 0$  is a parameter that describes how compatible cell  $i$  and  $j$  are. This pairwise factor expresses the similarity between two cells with scale of  $J$ . This is based on the idea that neighboring segments tend to have similar affordances. For example, cells constituting a wall at the top of the grid map would all share “not pushable up” affordance properties. Also, a directional pair of affordances, like pushable up and down, of neighboring cells have similar affordance values by symmetry. That is, if a cell is not pushable upward, the cell above should be not pushable down. This symmetry occurs as the affordances encode both accessibility and space availability: A “pushable up” cell should be accessible from the lower cell and the upper cell should be vacant. From the examples shown above, the pairwise factor expressing

similarity of the neighbors can “smooth” the beliefs of the affordances.

By using an MRF model, we can estimate the entropy reduction from the manipulation by both confirming the affordance of a cell and estimating the smoothing influence on the probability distribution. We use this model to plan a sequence of manipulations to reveal the affordance map. The problem can be framed as the maximum Traveling Salesman Problem, where we have distances (costs) between nodes (target cells) and rewards (entropy reduction) for visiting each node. As the TSP problem is proven to be NP-hard, here we apply a simple heuristic by choosing the maximum information gaining manipulation at every step. We compute the information gain by each possible manipulation as follows:

$$G = \alpha R - (1 - \alpha)C$$

$$R = \sum_{l=1}^4 p^l \Delta E(x_i = \text{true}) + (1 - p^l) \Delta E(x_i = \text{false})$$

$$C = D(\text{pose}_{\text{current}}, \text{pose}_{\text{access}})$$

Information gain  $G$  is the scaled sum of the reward  $R$  and the cost  $C$  with scaling factor  $\alpha$ . Reward  $R$  is the expected entropy reduction by confirming cell  $x_i$ ’s all four affordance labels to be true or false, each associated with predicted affordance  $p^l$  from equation 1. Cost  $C$  is a distance the robot needs to travel from current position  $\text{pose}_{\text{current}}$  to the goal position  $\text{pose}_{\text{access}}$  where it can perform the manipulation. The goal position depends on the type of affordance a robot wants to examine. From the formula above, a robot chooses maximum information gaining movements iteratively.

#### IV. EXPERIMENTAL RESULTS

The experiments have been conducted in the Gazebo simulator warehouse with the PR2 robot. The PR2 is equipped with the Microsoft Kinect to obtain a 3D point cloud and two 7 DOF manipulator arms to interact with objects. To train the affordance classifier, we built 20 maps of a 12 x 12m virtual warehouse each with different configurations of boxes (Fig. 2 right). The dimension of the box is 1 x 1 x 2m. Each box occupies a single 1 x 1m grid cell. The numbers of boxes in the scenes vary from 10 to 30, making various training inputs of relatively positioned boxes and walls. The boxes and walls are set to have same textures and heights so that they cannot be distinguished by visual cues alone. The ground truth for training examples are automatically generated based upon the assumption that walls and two aligned boxes cannot be pushed. An example of ground truth for the 4 affordances, pushable up, down, left, and right are shown in Fig. 4 left. We assume that robot in a cell can move to one of 4 directions to the neighboring cells. The robot is driven around the map by approaching nearest frontier (non-explored area) to cover 2D occupancy grid map and to collect snapshots of point clouds. The robot travels all reachable space and the gathered point clouds are stitched into a 3D world point cloud. For each view, obtained point cloud is segmented into 1 x 1 x 1m voxels and geometric features are extracted to predict




Fig. 4: Ground truth for 4 directions of affordances. Green arrow represents a cell being pushable and red represents not (left). Predicted pushing up affordance values for each voxel, color coded in heatmap style: 0 (not pushable) being blue and 1 (pushable) being red (right).




Fig. 5: Prediction result for affordance “pushing up” with heatmap color (left). Affordance prediction after MRF smoothing (right).

the affordance (Fig. 4 right). Since a single object might span over multiple voxels, the predicted affordance values are averaged over one grid cell to represent estimated affordance of object standing in the cell (Fig. 5 left).

Logistic regression (using WEKA tool [14]) was used to learn the parameters  $\beta_0^l, \dots, \beta_m^l$  for each label  $l=1, \dots, 4$  with number of features  $m = 42$ . 10-fold cross-validation was used to evaluate the classifier. The prediction precision and recall for pushing up and down was 81.5% and 79.7% and pushing left and right was 73.1% and 74.1%, respectively. As mentioned in the Section III, the affordance predictions for directional pairs are same as they are symmetric.

After obtaining the 2D occupancy grid map and affordance predictions for each occupied cell, a robot starts to plan a manipulation sequence by computing the information gain. To perform a pushing primitive, PR2 uses its two manipulator arms by aligning the end effectors vertically to maximize contact surface with the object. Accessibility needs to be considered in selecting candidate pushes. Accessibility can be exploited via sequences of moving objects, but we define it to be a single step approachability. For a given cell and affordance direction, if there exists a path to the target position where certain the affordance can be examined, the accessibility is set to be true.

The parameter  $\alpha$  for computing information gain is set to be 0.8 and cost of moving one cell is set to be 1. LibDAI library [15] is used to build up a MRF graph and to run belief propagation method to derive the probability



distribution  $P(\mathbf{x})$  in equation 2. A factor graph is generated by setting each occupied cell to be a node and connecting the neighboring up, down, left, and right cells by edges. The unary potentials are directly imported from the prediction result. The pairwise potentials are computed using the Ising model with  $J=0.4$ . An example of predicted affordance of “pushing up” is shown in Fig. 5 left. The prediction values between 0 and 1 are shown in heatmap color where 0 being blue (not pushable) and 1 being red (pushable). The result of MRF smoothing in the first manipulation step is shown in Fig. 5 right. The MRF successfully reduces the overall entropy by reflecting the assumption of affordance similarity between neighbors. Sometimes, applying MRF can “over-smooth” the affordance prediction by treating separate objects to be similar (*e.g.* Fig. 5 top left pushable up object). But its influence degrades when future manipulation is applied to confirm the true affordance.

In Fig. 7, a sequence of manipulations is shown. All candidate cells with 4 directional affordances are examined for their estimated information gain by manipulation. The robot selects the most information gaining manipulation and moves to the position where it can examine that affordance (*e.g.* A robot reaches to the ‘right’ vacant cell to confirm pushable ‘left’). After performing a pushing action, the robot returns back to the position where it started pushing and detects the change to confirm the affordance. As affordance depends on the relative positions of environment, affordance might be changed when an object gets pushed. Therefore, affordances should be modified according to the new configuration with the occupancy map the robot has. For example, in the second and third manipulation steps in Fig. 7, the object that has been pushed to the right ‘was’ pushable, but it becomes not pushable after the push, due to the wall. Other 3 directional affordances are also affected by successful pushing and those affordances are updated as well. At the end of manipulation step, the overall affordance beliefs are updated through MRF with newly obtained affordance information. In every iteration, an accessibility is recomputed according to the configuration change. The experiment is conducted by iterating manipulation selection, action, and affordance confirmation until the maximum information is reached (further information gain will be negative: movement costs more than reducing entropy).

The result of two approaches, exhaustive approach and maximum information gain method, are compared in Fig. 6. The experiments were executed over 20 different maps and the average entropy plot is shown. The exhaustive approach selects a manipulating cell based on proximity: the nearest occupied cell is selected on each step. The exhaustive method doesn’t utilize MRF model, since it doesn’t consider how influential each manipulation can be. Second approach selects manipulation by choosing maximizing information gain. Each manipulation gains new knowledge of affordance and the effects propagate through MRF. As shown on the plot, the information maximizing approach reduces overall entropy faster than exhaustive method.




Fig. 6: Average entropies between two approaches: exhaustive and maximum information gaining approach. As the maximum information gaining approach converges faster as it computes the effect of affordance confirmation on neighboring cells.

## V. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a method to build a task oriented affordance map by interactive manipulation. Through robot exploration, a 2D occupancy grid map was built with a 3D registered world point cloud. The point cloud was voxelized and affordance was predicted using unary and pairwise geometric features. Based upon the prediction, interactive manipulation was planned and executed using a MRF model to reduce the affordance map uncertainty. As a result, a complete affordance map was constructed for the robotic task. In the next step, we plan to improve the accuracy of the classifier as well as the efficiency of the manipulation scheme.

- Affordance definition can vary specifically to the task the robot needs to solve. The occupancy grid map with simplified directional affordances were suitable for rearrangement task. But for the other application, for example of grasping an object, the affordance should contain information about how and where it can be grasped according to the specific robot hand configuration. For more generalized affordance, it can be described as a distribution. More complex method should be involved to correctly infer specific types of affordance distributions. Also other features should be considered to enrich the representation of the world. Other than flat surfaces, many other geometric characteristics can be represented with various shapes. Other features like colors, textures, materials, etc. can also enrich the classifier for better prediction of affordance. The result of classification can also be improved by adapting non-linear classifiers.
- The affordances used in this paper were dedicated ones to the PR2 platform. As various robots have different power, manipulability, and sensors, the classifier for PR2 might not be adaptable to other robots. The classifier must consider robot parameters and specifications for predicting and planning. This adaptation can be enhanced by techniques such as reinforcement learning where robot platforms can utilize and modify the

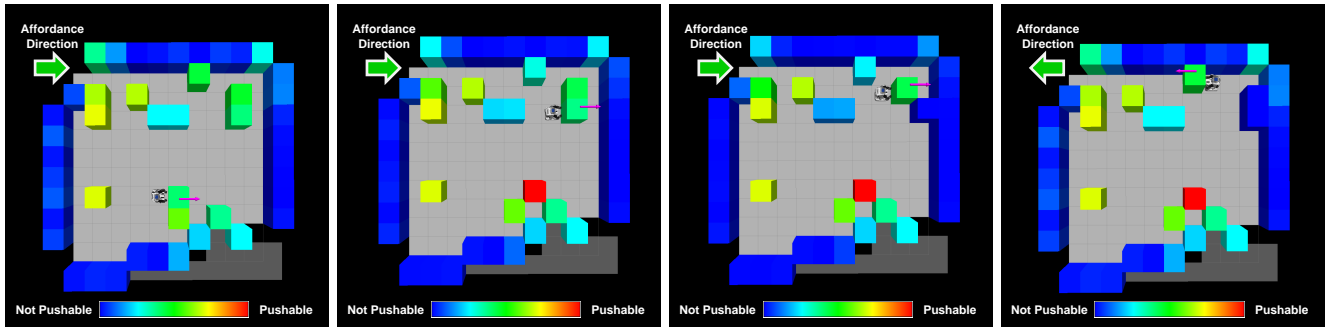


Fig. 7: A sequence of manipulation steps. At each step, a robot chooses the manipulation that has highest information gain which means larger entropy reduction with less movement. For example in the first step, robot chooses to push a box in the middle to the right. After pushing the box, the robot confirms affordance of the box to be pushable right (colored as red in the second picture). Second push confirms the box to be pushable, but it is not pushable anymore due to the wall behind (colored as blue in the third picture).

classifier learned from other robots.

- As mentioned earlier, affordances are not static properties. According to the various conditions how an object is situated, affordance can be dynamic. Understanding how affordance changes over time will boost robot's capability such as creating a desired affordance for specific tasks in novel settings.

#### ACKNOWLEDGMENT

This work was supported in part by the US National Science Foundation Robust Intelligence Program (grant IIS-1017134), and by the Defense Advanced Projects Research Agency under the Autonomous Robot Manipulation program (contract W91CRBG-10-C-0136).

#### REFERENCES

- [1] J. Gibson, "The concept of affordances," *Perceiving, acting, and knowing*, pp. 67–82, 1977.
- [2] M. Stilman and J. Kuffner, "Navigation among movable obstacles: Real-time reasoning in complex environments," *International Journal of Humanoid Robotics*, vol. 2, no. 04, pp. 479–503, 2005.
- [3] H. Wu, M. Levihn, and M. Stilman, "Navigation among movable obstacles in unknown environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1433–1438, 2010.
- [4] H. Koppula, R. Gupta, and A. Saxena, "Learning human activities and object affordances from RGB-D videos," *International Journal of Robotics Research (IJRR)*, 2013.
- [5] K. Varadarajan and M. Vincze, "Afnets: The affordance network," in *Computer Vision ACCV 2012*, 2012.
- [6] V. Chu, I. McMahon, L. Riano, C. McDonald, Q. He, J. Martinez Perez-Tejada, M. Arrigo, N. Fitter, J. Nappo, T. Darrell, and K. Kuchenbecker, "Using robotic exploratory procedures to learn the meaning of haptic adjectives," in *Robotics and Automation (ICRA), IEEE International Conference on*, 2013.
- [7] M. Gupta, J. Mueller, and G. Sukhatme, "Using manipulation primitives for object sorting in cluttered environments," *IEEE transactions on Automation Science and Engineering*, 2013. (to appear).
- [8] H. van Hoof, O. Kroemer, and J. Peters, "Probabilistic segmentation and targeted exploration of objects in cluttered environments," *Robotics, IEEE Transactions on*, vol. 30, pp. 1198–1209, Oct 2014.
- [9] E. Uğur and E. Şahin, "Traversability: A case study for learning and perceiving affordances in robots," *Adaptive Behavior*, 2010.
- [10] E. Şahin, M. Çakmak, M. Doğar, E. Uğur, and G. Üçoluk, "To afford or not to afford: A new formalization of affordances toward affordance-based robot control," *Adaptive Behavior*, vol. 15, no. 4, pp. 447–472, 2007.
- [11] D. Kim and G. Sukhatme, "Semantic labeling of 3d point clouds with object affordance for robot manipulation," in *IEEE/RSJ International Conference on Robotics and Automation (ICRA)*, 2014.
- [12] J. Lalonde, N. Vandapel, D. Huber, and M. Hebert, "Natural terrain classification using three-dimensional ladar data for ground robot mobility," *Journal of Field Robotics*, vol. 23, pp. 839–861, 2006.
- [13] K. Lai, L. Bo, and D. Fox, "Unsupervised feature learning for 3d scene labeling," in *IEEE International Conference on on Robotics and Automation (ICRA)*, 2014.
- [14] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. Witten, "The weka data mining software: An update," *SIGKDD Explorations*, vol. 11, no. 1, 2009.
- [15] J. Mooij, "libDAI: A free and open source C++ library for discrete approximate inference in graphical models," *Journal of Machine Learning Research*, vol. 11, pp. 2169–2173, Aug. 2010.