The 2016 AAAI Fall Symposium Series:
Artificial Intelligence for Human-Robot Interaction
Technical Report FS-16-01

An HRI Approach to Feature Selection

Kalesha Bullard[1], **Sonia Chernova**[1], and **Andrea L. Thomaz**[2]
[1]School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA
[2]Electrical and Computer Engineering, The University of Texas at Austin, Austin, TX
ksbullard@gatech.edu

Introduction

Our research seeks to enable social robots to ask intelligent questions when learning tasks from human teachers. We use the paradigm of Learning from Demonstration (LfD) to address the problem of efficient learning of task policies by example (Chernova and Thomaz 2014). In this work, we explore how to leverage human domain knowledge for task model construction, by allowing users to directly select a set of the salient features for classification of objects used in the task being demonstrated.

In the context of task learning, feature selection focuses on selecting a subset of relevant features to be used for building the task model. Within LfD, the problem of feature selection has been largely ignored with the premise that the most representative set of features will be selected manually a priori by the system developer or teacher (Chernova and Thomaz 2014). However if the goal is to enable robots to autonomously learn and perform general tasks, then autonomous construction of the task model state representation is an important step toward that end. With that, automatic feature selection is an active area of research; however the number of examples required to build an accurate model increases exponentially with the number of features in the state space representation. In LfD, only a small number of examples is typically provided by the human teacher to train the robot so this may be an insufficient amount of data for a statistical feature selection algorithm. Therefore, we explore an approach that leverages the human teacher's expertise in acquiring the set of relevant features to be used in constructing the task model, given a small number of object demonstrations.

In terms of prior work on learning task features, Cobo et. al. presented Abstraction from Demonstration, an LfD algorithm for learning state abstractions (subsets of relevant task features from the original state space) then subsequently using reinforcement learning on the abstract state spaces to generate an optimal policy (Cobo et al. 2011). Though this work also seeks to learn feature subsets through demonstration, we seek to enable the robot to actively query the human teacher for the relevant features associated with the task. The first step towards that end is presented here, a user study evaluating whether humans are able to *directly* provide this information to the robot, by enumerating features used to distinguish between object categories in a classification task.

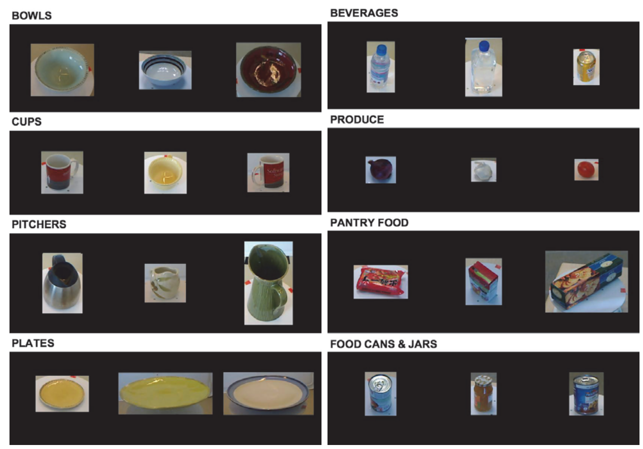

(a) Empty Dishwasher Task (b) Unpack Groceries Task

Figure 1: Illustration of one of the task demonstrations given for each of the tasks used in the *with images* condition of the user study.

Specifically, we conducted a user study on Crowdflower, directly requesting from human naive users the most useful features for determining which category an unknown object belongs to amongst four classes of objects given a priori, for two different multiclass classification tasks. In the study, we examine two conditions: (1) only class names were provided with *no* examples (images) of the objects given to illustrate each class or (2) images of objects used as examples to illustrate each class were provided to introduce bias and asses how this impacted user characterization of the object classes (see Figure 1). Our preliminary findings suggest that in both conditions: (1) users are not reliably able to *directly* select the set of useful features for classifying the objects in a task and (2) the ability of users to do this in a way that is comparable to baseline performance is task-dependent. These findings however are at least partially consistent with our hypotheses and to conclude, we discuss some next steps in this ongoing project.

Feature Selection from Demonstration

Problem Statement and System Overview

We seek to explore the following research questions:

- *Learning:* Using only a small number of examples, is a human teacher able to characterize the most salient features of a classification task, such that a robot can better learn the task interactively than it would be able to autonomously?
- *Interaction:* If so, how do we enable the robot to extract this information from a human teacher? And does physical embodiment impact how humans characterize the task?

A diagram of the high-level system we are developing to address these questions is illustrated by Figure 2. This is ongoing work, and the experiment presented in this paper is the *first* step in exploring both the *Learning* question and the first *Interaction* question, as highlighted in red on the diagram.

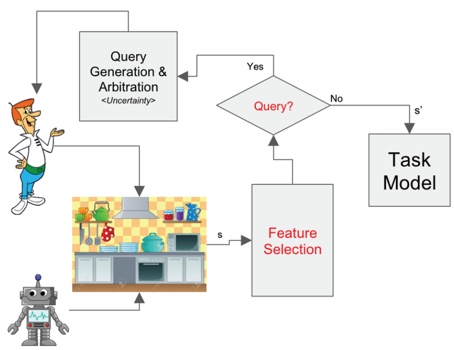

Figure 2: Diagram of System Overview

In our learning scenario, a robot and a human are perceiving the same environment, and there is a set of perceptual features associated with the observed state of the world s. The goal is to perform feature selection in order to prune the state representation to only the set of features relevant for the task, s'. The process of feature selection using human input is intended to be iterative and the outputted set of features is used to construct the state representation for the task model.

Approach

Our approach uses demonstrations from human users to learn sets of relevant features, for each class of objects required for a task. As our running example, we situate the robot within a kitchen setting, learning tasks which require commonly encountered objects in that setting. We experiment with two tasks: *emptying the dishwasher* and *unpacking the groceries*, each a multiclass classification problem involving four different object classes.

We define a high-level task, $t \in T$, as a sequence of primitive actions, $a_1, a_2, \ldots, a_n \in A$, where A is the set of actions in the task. We assume that tasks are object-centric and therefore A is used to manipulate a set of objects in the situated environment, when performing task t. In order to successfully perform the task, a robot must *first* be able to select objects in its environment which are members of each target object class in the task for use in task execution; then subsequently act on the selected objects.

In terms of selecting feature subsets, in the feature selection literature, features are typically described as being *useful* or *relevant* (Blum and Langley 1997; Kohavi and John 1997) and feature subset selection algorithms can be partitioned into *filters*, *wrappers*, and *embedded methods* (Guyon and Elisseeff 2003).

Feature relevance measures correlation between inputted feature values and outputted class labels and is therefore only dependent upon the dataset given. Filters are the least computationally intensive class of feature selection algorithms; they are employed as a preprocessing step before classification and are used to eliminate all irrelevant features from the original feature set. Feature usefulness, in contrast, measures how useful a feature is in contributing to the learning goal; it is dependent on both the data and the classifier being used for learning. Wrappers and embedded methods both conduct a search in the space of possible feature subsets, using learning performance as a metric for evaluating feature subsets, and are used to compute feature usefulness. Whereas wrappers conduct an exhaustive search through the feature subset space, embedded methods use a greedy search strategy to incrementally add or remove features.

Even when employing wrappers or embedded methods though, filters are typically used as a baseline feature subset selection approach. Therefore in this work, we use filters for the baseline computational feature subset selection approach at every iteration of the learning episode. The goal here is to first examine whether human selection of features is comparable to even a baseline statistical feature selection algorithm; other computational feature selection algorithms can always be added later. Toward that end, there are two different methods we use to autonomously compute feature subsets. Both compute a new subset of *relevant* features each time new object demonstrations are given.

The first method, we call the *Dynamic Binary-Union Feature Set*. Given n object classes, this approach involves splitting the original multiclass dataset into n binary class datasets, where each dataset D_i has samples of object class o_i as positive samples and all other samples as negative samples and where $\{i \in \mathbb{Z} \,|\, 0 \leq i \leq n-1\}$. To construct this, each object class o_i, in the set of object classes O used in task t, can be characterized by a subset of features F_{o_i} that enables an agent to correctly recognize a previously unseen member of o_i. Each subset of relevant features generated for a class is computed using a filtering algorithm, which ranks features by information gain, as shown in Equation 1 with $f_j \in F$, the set of all features. At each iteration of the learning episode (after new object demonstrations have been added), F_{o_i} is initialized as an empty set, and all features with a positive information gain are added to F_{o_i}. The subset of features F_b selected for the task, called the *Binary-Union Set*, is the union of all such subsets F_{o_i}, such that $\{\forall o_i \in O \,|\, F_b = \bigcup F_{o_i}\}$.

$$IG(o_i, f_j) = H(o_i) - H(o_i|f_j) \tag{1}$$

The second method for automatically generating feature subsets, we call the *Dynamic Multi-Class Feature Set*, F_O.

Here again, a filtering algorithm, which ranks features by information gain is used. However, in contrast to the previous method described, this method ranks features according to their information gain for all n object classes *simultaneously* At each iteration of the learning episode, F_O is initialized as an empty set, and all features with a positive information gain are added to F_O, such that $IG(O, f_j) = \sum_{o_i} IG(o_i, f_j)$.

After validating learning performance with three different classifiers (k-nearest neighbors, support vector machines, and random forests), we observed comparable performance and selected a kNN classifier for the remainder of this work.

Evaluation

The research questions we examine in the experiment presented explore (1) whether humans are able to characterize the most salient features of object classes used in a task, given only a small number of examples (which may include no examples at all), and (2) how a robot should query or extract this information from a human teacher. Toward that end, we have two hypotheses that we are testing: *(1) Humans intuitively understand and are able to characterize the salient features of a task, and (2) humans can more effectively communicate this information indirectly (selecting representative instances) than directly (enumerating relevant features).*

In this work, we test the first hypothesis (user ability to effectively characterize a task) using only the second condition of the second hypothesis (by enumerating all salient features of the task). The baselines for comparison are the statistical feature selection algorithms described in the Approach section, and the primary evaluation metrics are learning performance and sample complexity.

We evaluate our approach on two different tasks. The *emptying the dishwasher* task has four classes: (1) bowls, (2) cups, (3) pitchers, and (4) plates, as illustrated in Figure 1a. Similarly, the *unpacking the groceries* task has four classes: (1) beverages, (2) produce, (3) pantry food, and (4) food cans and jars, as illustrated in Figure 1b.

Data Collection and Training

Image Dataset We used the University of Washington RGB-D Object Dataset to obtain images of common household objects as input for task model construction. There are over 300 objects in the dataset, organized into 51 categories and within each category (e.g. soda can), there are multiple object instances (e.g. pepsi can, mountain dew can, etc.). For each object instance, there are several hundred images captured from different viewpoints and distances from the camera, and some objects in the dataset have been captured under more than one lighting condition (Lai et al. 2011). Each image has a unique label that includes the object category name and object instance id number, used for ground truth assignment of images in the dataset to object classes associated with the task being learned by the robot. Additionally, the dataset includes a cropped version of each rgb and depth image, whereby the background scene has been cropped and the object isolated. Given this as input, we fit a rotated bounding box[1] to the isolated object in the image and extract the following features: {*x, y, z, orientation, r, g, b, bounding_box_volume, bounding_box_area, bounding_box_length, bounding_box_width, bounding_box_height, bounding_box_aspect_ratio, surface-area_volume_ratio, compactness, number_SIFT_features*}

The image dataset includes over 200,000 images in total. For each task, we only consider the subset of images corresponding to objects used in that task. So for the emptying the dishwasher task, the task-relevant dataset includes only images of bowls, cups, pitchers, and plates. The task-relevant dataset is partitioned into a 60/40 split, such that 60 percent of the dataset can be randomly sampled for training demonstrations. For this experiment, we randomly sampled 10 different training sets from the 60 percent partition (similar to teaching the same tasks in ten different environments); from the remaining 40 percent, we generated a test set of 2000 instances used to evaluate all classifiers.

Learning Episode In a typical LfD interaction, human teachers provide only a small number of demonstrations; therefore, we collect up to a maximum of twenty demonstrations of each object, such that by the end of a learning episode, the training set is a uniformly distributed sample of $20 * m$ object images, where m is the number of object classes in the task. In the experiments presented, $m = 4$ for both tasks. The test set is also sampled uniformly.

At each iteration of the learning episode, a set of object demonstrations are given (1 example per class) and the learner can then reassess the set of features to be used in task model construction, based upon the new information. The statistical feature selection algorithms compute information gain from the instances in the updated training set and generate new subsets of relevant features. From the user study, we obtain one set of features for each {task, training set} pair, based upon the examples provided at the beginning of the learning episode and keep the same human-selected feature subset throughout the learning episode. A classifier is trained and tested for each feature subset selected.

Crowdsourced User Study We conducted a user study on Crowdflower to collect data from humans about what features they would use to differentiate between object classes given in household task. There were two parts to the basic instructions, in both conditions, provided below as follows:

(a) You are teaching a robot to ⟨task⟩. In order to do this, the robot must learn from you how to distinguish between: ⟨classes⟩. (b) If the robot is doing the task and encounters a new object, what features below would be most useful for the robot in determining which one of the above categories the new object belongs to? (Check all that apply)

There were two conditions examined in the study: (1) Given no images to illustrate each category and (2) Provided with images as examples of each class. In condition two (images), there was one additional part to the instructions. After part a of the above instructions and before part b, we added one sentence: *Below you provide the following examples of*

[1]We assume the rotation is only with respect to the countertop normal

each type of object, though there may be others. This was followed by an image like the one shown in either Figure 1a or Figure 1b for dishwasher and groceries tasks respectively.

For both conditions, we collected 20 user responses per questionnaire. The images condition had ten questionnaires since there are ten different training sets being used as input. In the images condition however, after pruning users who were not able to successfully answer the test questions, we were only left with 10 user responses per questionnaire. In order to include a feature in the selected subset for a {task, training set} pair, at least half of the users who completed that questionnaire had to select the feature as relevant.

Preliminary Results

As an important note, in the experiments run, the test set is approximately two orders of magnitude larger than any training set provided to the robot, so there are many examples of objects the learner will be tested on that it may never have seen in training. This makes generalization more challenging. However in this work, our goal is to *compare* the different feature subset selection approaches and thereby assess *relative* (rather than absolute) learning performance.

Figures 3a and 3b show the learning curves (bias) for the Emptying Dishwasher and Unpacking Groceries tasks respectively. Here the yellow and orange curves represent the performance of the human selected feature sets, and other curves are all baselines we use for comparison, described in the Approach section. On the Emptying the Dishwasher task, the human selected feature sets perform on average worse than even the baseline with no feature selection. On the Unpacking Groceries task, the human selected feature sets both perform on par with the other baselines. The orange learning curve is an average performance across all ten of the training sets for each task.

Overall our results indicate that humans are not reliably good at selecting the low-level features *directly*. At best, users are able to perform comparably to the "baseline" computational feature selection algorithms. We also observe that performance of the human-selected feature sets is not consistent across tasks. People were able to better characterize the object categories used in the unpack groceries task than in the empty the dishwasher task.

As a follow-up question, we asked users to list any additional features they believed would be helpful in distinguishing between the provided object categories for the task. We got responses such as: *shape, weight, sound of the object, hardness, texture and material the object is made of, presence of a handle, smell, and taste of the object*. Some of these attributes are more abstract and difficult to quantify. This additional data suggests that directly communicating the low-level features involved in concept characterization is insufficient to capture all of the prior domain knowledge and high-level features people use in determining which objects are appropriate for serving a particular purpose in a task.

Our second hypothesis was that humans can more effectively communicate feature information *indirectly* (through the selection of representative instances) than directly (enumeration of features). Our next step is to conduct a user study where people are allowed to select representative instances of a particular object class used in a task; in this case, the computational feature selection algorithm will use the instances selected to infer what the human was trying to communicate through selecting them. We also plan to expand to more tasks in order to understand if there is a detectable trend or pattern that allows us to understand/predict which tasks people are inherently better at characterizing.

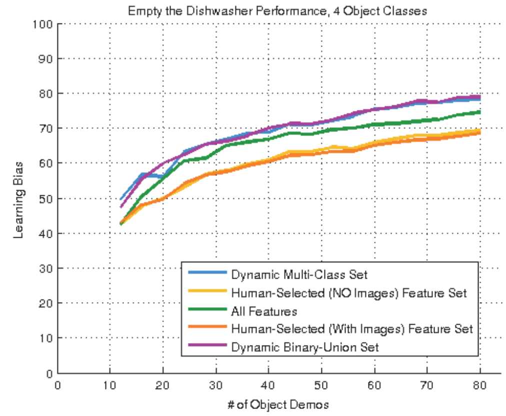

(a) Emptying Dishwasher Task

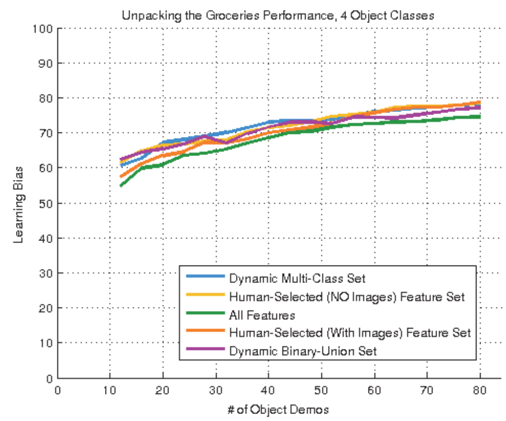

(b) Unpacking Groceries Task

Figure 3: Learning Performance of Tasks

Conclusion

Enabling robots to request the most useful features for characterizing a task is an important step toward autonomous task model construction. In this work, we conducted a user study to explore whether naive users are able to characterize the most salient features of a classification task, such that the robot can better learn the task interactively than it would be able to autonomously. Our findings indicate that users are not reliably able to *directly* select a set of useful features for classifying objects in a task and that the ability of users to do this in a way comparable to the baseline is task-dependent. However, this is ongoing work and in our next user study, we will explore whether allowing user to *indirectly* communicate the salient features will prove to be a more successful way of extracting this information from a human teacher.

References

Blum, A. L., and Langley, P. 1997. Selection of relevant features and examples in machine learning. *Artificial intelligence* 97(1):245–271.

Chernova, S., and Thomaz, A. L. 2014. *Robot Learning from Human Demonstration*. Morgan and Claypool Publishers.

Cobo, L. C.; Zang, P.; Isbell Jr, C. L.; and Thomaz, A. L. 2011. Automatic state abstraction from demonstration. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22, 1243. Citeseer.

Guyon, I., and Elisseeff, A. 2003. An introduction to variable and feature selection. *Journal of machine learning research* 3(Mar):1157–1182.

Kohavi, R., and John, G. H. 1997. Wrappers for feature subset selection. *Artificial intelligence* 97(1):273–324.

Lai, K.; Bo, L.; Ren, X.; and Fox, D. 2011. A large-scale hierarchical multi-view rgb-d object dataset. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, 1817–1824. IEEE.