

# Designing Robot Learners that Ask Good Questions

Maya Cakmak and Andrea L. Thomaz

School of Interactive Computing, Georgia Institute of Technology  
801 Atlantic Dr., Atlanta, GA, USA  
{maya,athomaz}@cc.gatech.edu

## ABSTRACT

Programming new skills on a robot should take minimal time and effort. One approach to achieve this goal is to allow the robot to ask questions. This idea, called Active Learning, has recently caught a lot of attention in the robotics community. However, it has not been explored from a human-robot interaction perspective. In this paper, we identify three types of questions (label, demonstration and feature queries) and discuss how a robot can use these while learning new skills. Then, we present an experiment on human question asking which characterizes the extent to which humans use these question types. Finally, we evaluate the three question types within a human-robot teaching interaction. We investigate the ease with which different types of questions are answered and whether or not there is a general preference of one type of question over another. Based on our findings from both experiments we provide guidelines for designing question asking behaviors on a robot learner.

## Categories and Subject Descriptors

I.2.9 [Artificial Intelligence]: Robotics; H.1.2 [Models and Principles]: User/Machine Systems

## Keywords

Learning from Demonstration, Active Learning

## 1. INTRODUCTION

Learning from Demonstration (LfD) is aimed at allowing end-users to program new skills on a robot with minimal effort [1]. To this end, LfD techniques try to maximize the generalizability of the learned skill to unseen situations with a minimal number of demonstrations provided by humans. Recently, Active Learning (AL) methods have been explored to achieve this goal. The idea in AL is to improve learning rates by giving the learner more control over what examples it receives. The AL process involves the learner making a




Figure 1: Participants demonstrating three different skills to our robot Simon.

*query*<sup>1</sup>, which is a request for certain information, and the teacher provides this information. The requested information can have various forms, although most commonly it is a category label for an unlabeled data instance. Different strategies are available for choosing such queries (*e.g.* reducing uncertainty or reducing variance).

Our focus, in this paper, is how the incorporation of AL methods in LfD impacts a robot’s interaction with its user. In prior work looking at using AL in an HRI setting, we learned that compliance with the robot’s questions is not guaranteed [2, 5]. Human teachers sometimes ignore the robot’s questions, particularly when they judge it to be a *bad* question. Thus, in this work we try to operationalize what it means for a robot learner to ask *good* questions.

First we look at the notion of good questions from the machine learning perspective. We identify three types of queries from AL (label, demonstration and feature queries) and discuss how they can be applied in an LfD setting (Sec. 3). Then we look at the notion of good questions from a human learning perspective, with an experiment into human question asking (Sec. 4). This allows us to characterize people’s usage of the identified query types. Finally, we evaluate the three query types in an HRI setting, where people teach skills to a humanoid robot via LfD (Sec. 5). We find that feature queries are perceived as the smartest and label queries are the easiest to answer.

## 2. RELATED WORK

The potential of AL for making learning efficient has been noticed by the Robotics community. AL has been applied to a number of LfD frameworks. For example, *confidence based autonomy* [6] and *dogged learning* [12], use the principle of uncertainty sampling to select states in which

<sup>1</sup>We use the terms (making a) *query* and (asking a) *question* interchangeably throughout the paper.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HRI’12, March 5–8, 2012, Boston, Massachusetts, USA.  
Copyright 2012 ACM 978-1-4503-1063-5/12/03 ...\$10.00.

the learning agent requests a demonstration while learning a policy. In [10], the robot actively selects points outside the region of stability of a learned policy, and requests demonstrations from these states.

The evaluation of robotic systems that ask questions to humans has been limited, particularly in LfD settings. Our earlier work compares passive and active task learning [5] and addresses the question of *when* to ask questions in a mixed-initiative AL setting [2]. Rosenthal et al. investigate how augmenting questions with different types of additional information improves the accuracy of human teachers’ answers [21]. In later work, they explore the use of humans as information providers in a real-world navigation scenario [22]. Other related work explores the idea of learning actions on a robot through dialog [4].

Evaluations of AL systems with naive teachers in Human-Computer Interaction (HCI) are also relevant. For instance, our finding in previous work that people do not enjoy a constant stream of questions, was shown in a recommendation system domain [13]. Other studies investigated different HCI issues such as feasibility of novel query types (*e.g.* feature queries) [7] or cost of different query types [23]. Similar to this paper, Gervasio *et al.* investigated question asking in procedure learning with a range of question types [8, 9].

### 3. QUERY TYPES IN ACTIVE LEARNING

The conventional *query* in the AL literature involves choosing an unlabeled instance and requesting a label for it. The instance can be chosen from a pool of unlabeled instances or instantiated by the learner in some way. Such queries have been used in learning skills on a robot, where a skill is represented as a policy that maps a state to a discrete action. In this context, a query consists of asking for an action in a chosen state [6, 17, 12]. This is useful in an HRI setting when the robot’s actions are discrete and the human has a way to refer to each action. However, robot skills often involve continuous actions and the input from the teacher to the learner is a sequence of state-action pairs (*i.e.* trajectories). In these cases it is impractical to ask for an isolated state-action pair (*e.g.* asking for the motor commands of a given arm configuration). Thus we need to re-think the way that a query is made. We consider the following alternatives.

#### 3.1 Label queries

Robot skill learning involves modeling/encoding a skill from a set of demonstrations, such that it can be reproduced correctly. The demonstrations provided by the teacher are inherently labelled as *positive*. One way the robot could make queries is to execute a motion and ask whether the skill was performed correctly. We call this a *label query*.

The information provided by label queries depends on the answer. If the response is positive, then the motion can be used as another demonstration. What to do with negative examples is not as straightforward. LfD methods are designed to encode a skill from positive examples only, since it is unnatural to demonstrate “what not to do”. One way to make use of negative examples that arise from label queries, is to update the learned model such that the probability of the negative data being generated by the model is minimized while the probability of the positive data being generated is maximized. The main issue with this idea is the attribution of negativity to the whole trajectory, while only parts of the trajectory might be responsible for the failure. A second ap-

proach for making use of negative examples is to guide the learner’s future queries towards positive examples (*e.g.* [11]).

Methods for generating label queries depend on the particular framework used for representing the skill. However a general approach applicable to most frameworks is to sample trajectories from the learned skill and evaluate them with a certain criterion. For instance, the robot can choose to query the trajectory that it is least certain about or the trajectory that is most likely to increase the applicability of the skill.

#### 3.2 Demonstration queries

The second type of query, which we call *demonstration* or *demo queries*, involves creating a new scenario and requesting a demonstration from the teacher. Demo queries give less control over what information is acquired from a query, as compared to label queries. In label queries, the learner specifies the whole trajectory and the teacher only provides a label. In demonstration queries, the learner only specifies certain constraints, while the trajectory is still produced by the teacher. Nevertheless, this gives some control to the learner such that useful demonstrations can be acquired. Demo queries are analogous to a method known as *active class selection* [16], which consist of requesting an example from a certain class.

One way to constrain trajectories provided by the teacher is to specify the starting state. Trajectories are often represented with a sequence of end effector configurations relative to a goal object. Thus the robot can configure its end effector in a certain way relative to the goal and request the demonstration to start in this configuration. A larger range of queries can be made by manipulating the target object. A different way of constraining the demonstrations provided by the teacher is to allow the teacher to control only a subset of the robot’s joints while the robot executes a certain trajectory on the rest of the joints, as in [3].

#### 3.3 Feature queries

The third type of query is inspired from a relatively recent technique in AL that involves asking whether a feature is important or relevant for the target concept that is being learned [20, 7]. A particularly successful application of this method is document or email classification. Learning approaches to this problem involve the teacher reading documents one by one and providing category labels for them. Even in an AL setting, providing a label for the queried documents is cumbersome and time consuming. On the other hand a feature query directly asks whether a word is a strong indicator of a certain category. This allows the learner to directly modify its model for categorizing new documents and drastically reduces the time spent by the teacher.

Note that the crucial element in the success of feature queries in the given example is that the features (words) with which the instances (documents) are represented are meaningful for humans and the way they contribute to the classification of the instance (being a strong indicator) is intuitive. We believe that robot skill learning is at a unique position for taking advantage of this method: while features might not be as human friendly (*i.e.* feature names might be too technical and feature values might be arbitrary numbers) the robot’s embodiment can be used to refer to them.

Methods for choosing feature queries are also dependent on the framework used for representing and learning skills. One framework that allows feature queries to be directly

integrated is task space selection [18, 14]. This involves representing demonstrations in high dimensional spaces that involve features that might or might not be relevant (*e.g.* relative position and orientation of different points on the robot to different objects in the environment). Methods try to identify a subspace or assign weights to each feature such that the skill is best represented. In this context a feature query is to directly ask whether a feature is important for a skill. These queries can also be used for directly manipulating the skill representation or for guiding other queries.

## 4. HUMAN QUERIES DURING LEARNING

In designing good questions for a robot learner, the Active Learning literature gives us one perspective, but we are also interested in taking human question asking behavior as inspiration. Our intuition is that questions people typically use will be easy for them to answer. While human question asking has been investigated and modeled in different cognitive tasks such as text understanding [15, 19], we did not find any studies on human questions in the context of skill or task learning by demonstration. Thus, we conducted an experiment to investigate how humans ask questions in scenarios similar to the ones robots are faced with. Our experiment involves providing the participant with demonstrations of a task, letting them ask questions about it and then asking them to perform it in a new situation.

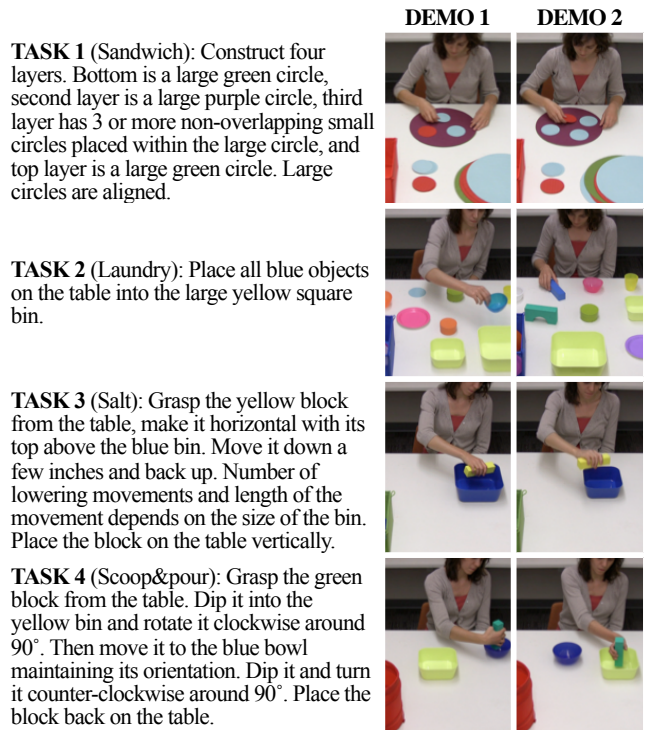
### 4.1 Experimental Design

We consider four tasks described in Fig. 2. The tasks are chosen to cover the two types of LfD problems in robotics. The first two tasks are focused on *learning a task goal* (desired end state) whereas the latter two emphasize *learning skills* to perform the task (desired actions or movements to accomplish the task). Each of these are designed to be an abstraction of a real-world task (*e.g.* making a sandwich), to let the experimenter answer participants’ questions based on the real-world counterpart. The abstraction also avoids the influence of prior knowledge about the tasks.

The tasks or skills involve a fixed set of objects. Two demonstrations, varied in certain aspects, were recorded for each skill. The recordings were used for demonstrating the skill to the participants. This is done to mitigate unintended variances between the demonstrations given to different participants. Participants learn all four tasks and the order is counterbalanced across participants. The interaction between the participant and the experimenter is recorded by a camera that oversees the table (Fig. 5).

#### 4.1.1 Procedure

Throughout the experiment, the participant sits in front of a table with a computer screen to their right for displaying the demonstration videos. For each task, the experimenter first brings the involved objects onto the table. Then the participant watches the video. Participants are told that they can watch the video as many times as they want, at any time during the experiment. This is to reflect the fact that a robot has access to the demonstrations provided by a teacher before they ask questions and forgetting what happened in the demonstration is not an issue. Thus we would like to avoid questions that address information present in the videos. After watching the video at least once, the participant is probed for questions. There is no upper limit on number of questions, but a lower limit of three questions.



**Figure 2: Description of the four tasks used in the experiment and snapshots from the two demonstration videos shown to the participants.**

When the participant has no more questions, the experimenter configures the environment in a certain way and asks the participant to perform the task. The participant is reminded that they can keep asking questions, during and after their execution. The execution is repeated twice with different starting configurations.

In the instructions, participants are encouraged to ask any type of question that they can think of and that they are not restricted to yes/no questions. They are encouraged to refer to the objects on the table in their questions. Participants are told that the purpose of the study is to take inspiration from their questions to allow robots to ask questions. Thus we tell them not to ask questions about the underlying purpose of the tasks since explaining these to a robot would be impossible. Such questions are not answered by the experimenter if they occur.

#### 4.1.2 Evaluation

All recordings were later transcribed by the experimenter using ELAN<sup>2</sup>. This includes annotations of the participant’s questions, the answers given by the experimenter and the participant’s interaction with the objects and gestures. In addition we annotated the demonstration views by the participant and the task executions.

The participants’ questions are categorized into the three query types described in Sec. 3 based on the following criteria. If the question asks the experimenter to evaluate an action or a complete task execution it is considered a *label query*. If the question requests an action, an action suggestion or a complete task execution from the experimenter, it

<sup>2</sup>[www.lat-mpi.eu/tools/elan/](http://www.lat-mpi.eu/tools/elan/)




Figure 3: (a) Distribution of questions into three types of queries for each task in the first experiment averaged across participants. (b) Distribution of feature queries into sub-categories. (c) Average distribution of common actions and gestures that physically ground questions across tasks.

is categorized as a *demonstration query*. Lastly, all questions that address a certain aspect of the task in general terms is considered a *feature query*.

We also categorized the questions in terms of their form based on the taxonomy established in [15]. This taxonomy is shown in Fig. 4 with examples from our experiment for each form of question.

## 4.2 Results

12 participants (3 female, 9 male) completed our first experiment. The interactions took about 25 minutes and involved an average of 40.58 (SD=13.09) questions.

### 4.2.1 Query types

The distribution of questions asked by participants across the three query type categories is given in Fig. 3(a). We observe that a significant majority of questions asked by humans (82%) are feature queries. We notice a large variety of query types within feature queries in terms of the information they attempt to gather from the teacher. We establish the following subcategories of feature queries (a similar distinction was made in [8]).

- *Feature relevance tests* directly ask whether a feature is important for a task: *Does f matter?* (Fig. 5(b))
- *Feature invariance tests* ask whether a feature has to have a certain value: *Does f have to be  $f_1$ ?* (Fig. 5(d))
- *Feature value tests* ask whether a feature is allowed to have a certain value: *Can f be  $f_1$ ?* (Fig. 5(a))
- *Feature value requests* directly asks the value range that a feature is allowed to have: *What can f be?* (Fig. 4 (SIMPLE))

Note that  $f$  can refer to an attribute of the involved objects, relations between objects, a property of the involved movements, a parameter of the task description or the dependency between any two of these. The distribution of feature queries into these subcategories is shown in Fig. 3(b). We observe a rather uniform distribution across the first three categories, while the open-ended value requests are less common.

In addition we observe a distinction between *full* and *partial* demonstration and label queries. A full demonstration query requests a complete demonstration while a partial one requests a single step of the demonstration. Similarly, a full label query asks for an evaluation of a complete task execution while a partial label query asks whether a step was

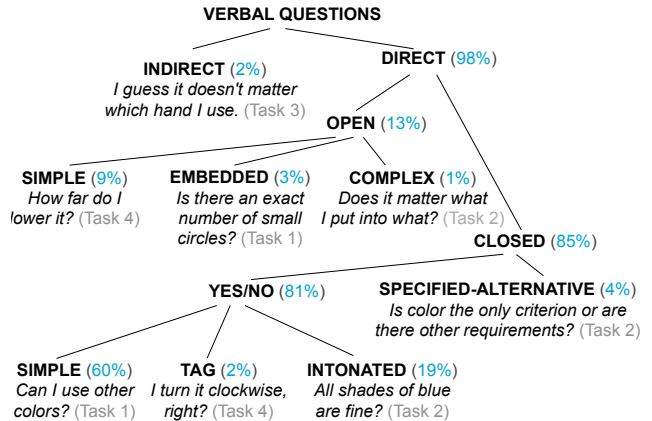


Figure 4: Question forms [15] and example instances from our experiment (percentage of average occurrence shown in parentheses).

correct. Such partial label queries seem to be a useful mechanism for avoiding the credit assignment problem that arise when a negative label is encountered (Sec. 3.3).

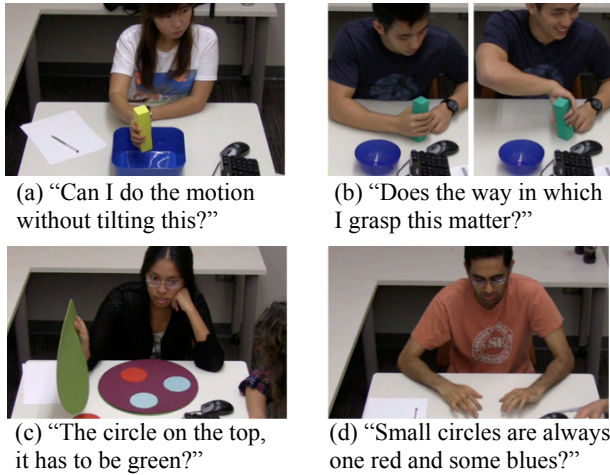
### 4.2.2 Question forms

The average usage percentage of each question form across tasks is shown in Fig. 4. We observe a clear preference for yes/no questions (81%), more often posed in the simple form (60%). We notice that the choice between the three types of yes/no questions reflect whether the person has an expectation about what the answer is, and indicates their confidence in this expectation. While *simple* yes/no questions are neutral about either answer, *tag* yes/no questions (*e.g. right?, correct?*) indicate strong expectation of the answer being “yes”. One example is the participant saying “I can use either hand, right?” when he has already started a movement with his alternate hand.

### 4.2.3 Physical grounding of questions

We observe that human questions during task learning are largely embodied and physically grounded. On average 62% of the questions asked by a participant involved interacting with the environment. Label queries are inherently embodied questions since they require the robot to perform the actions to be labelled. Similarly demo queries require the person to create the scenario in which a demonstration is requested. It is, however, possible for the person to verbally





**Figure 5: Examples of the four common types of physical groundings used with questions.**

describe a sequence of actions (or a hypothetical scenario) and request a label (or a verbal demonstration).

Feature queries on the other hand are more amenable to purely verbal question asking, however they require the ability to talk about features and feature values. More specifically, they require tokens that have a common physical grounding for the two parties (person asking the question and person being asked) to replace  $f$  and  $f_1$  in Sec. 4.2.1. Even though such tokens might exist, we observe that they are often replaced or accompanied by actions or gestures that communicate them – we saw four common types:

- A *single value instantiation* refers to actions that generate a certain feature value while asking a question about it. For instance in Fig. 5(a) the person asks if the yellow block can be lowered without being tilted while providing an instance of non-tilted orientations.
- *Multiple value instantiations* refer to actions that generate several different values of feature during a feature query. An example is the participant showing several different grasps while asking whether the way in which the block is grasped matters (Fig. 5(b)).
- *Referential gestures* include pointing, touching or holding up a certain object in the environment while making a feature query that relates to a feature of the object. For instance, in Fig. 5(c) the participant holds up a green large circle while asking whether the top layer has to be green.
- *Iconic gestures* involve arm and hand movements that communicate a certain feature or feature value referenced in a question. An example, shown in Fig. 5(d), is placing two hands at a certain distance from each other to indicate a certain size.

The average distribution of these actions and gestures across subjects is given in Fig. 3(c). We observe that 59% of feature queries made by participants involve some form of physical ground. Single value instantiations and referential gestures are more common than the other two types. Also we observe interesting differences in how they are used in different tasks, which we discuss in Sec. 4.2.4.

#### 4.2.4 Learning Goals versus Skills

We do not see major differences in the distribution of query types or question forms between tasks that require learning goals (Tasks 1 and 2) versus tasks that involve learning skills (Task 3 and 4). The main difference in the questions asked during goal versus skill learning is seen in the physical grounding of the questions (Fig. 3(c)). We find that questions during skill learning are more likely to be physically grounded (74%) than during goal learning (46%). This is not surprising as movements are harder to describe verbally. Secondly, we see that value instantiations are more common in skill learning (41%) as compared to goal learning (15%). Correspondingly, a larger portion of the questions during goal learning use referential and symbolic gestures. This is also an intuitive finding as the features involved in skill learning are about the movements, while the features involved in goal learning are about objects in the world. In a sense, referring to features of an object in the world, or using a gesture to indicate a feature value are both instantiations of the feature.

## 5. ROBOT QUERIES DURING LEARNING

Next, we investigate the use of the three types of queries described in Sec. 3 with an HRI experiment that involves teaching skills to a robot and answering its questions.

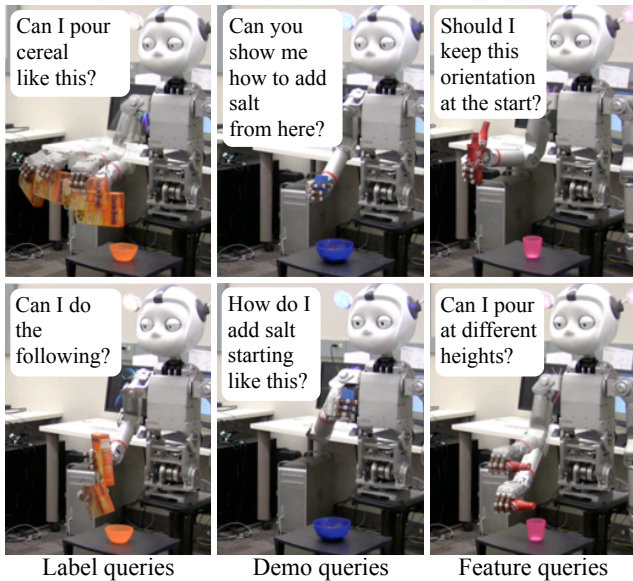
### 5.1 Platform

The robot platform used in our experiment is Simon, an upper torso humanoid robot with 7 Degree of Freedom (DoF) compliant arms, 4 DoF hands that allow object manipulation and gestures, and a socially expressive head (Fig. 1). The arms can be operated in a gravity compensation mode with low stiffness to allow easy kinesthetic interactions. Simon’s behaviors are implemented as a finite state machine whose transitions are controlled with speech commands. The Microsoft Speech API is used for speech recognition.

### 5.2 Experimental Design

The experiment involves teaching three different skills to Simon: (i) pouring cereal into a bowl, (ii) adding salt to the salad and (iii) pouring soda into a cup. Each skill is paired with one query type and the pairings are not varied, since the three skills are very similar in nature. The particular queries made by the robot are pre-scripted, to avoid the large variance that would occur in queries automatically generated based on the teacher’s demonstrations. This is due to the stochasticity inherent to query methods as well as the unpredictable variance in the provided demonstrations.

The skill-query type pairing and the two particular queries made for each type of query are shown in Fig. 6. The first label query involves a trajectory for pouring cereal that is likely to be viewed as a negative example since the cereal box is tipped too early. The second label query involves a valid cereal pouring trajectory from a direction that is less likely to be demonstrated by the participants (because of the difficulty in posing the joints to make the cereal box properly oriented). Demo queries involve one starting position that is slightly outside the range of expected common starting points closer towards the robot, and another starting position that is high above the bowl (which is unlikely to be a starting point in the humans’ demonstrations). Feature queries involve one feature invariance test about starting orientation of the coke bottle, and one feature value test for



**Figure 6:** The queries made by the robot for each type of query in the second experiment.

heights of the bottle at the pouring location. Both feature queries are grounded using multiple instantiations of the referred features. Participants teach all three skills and the order of the skills is counter-balanced.

### 5.2.1 Procedure

Each participant first watches an instructional video that describes the study. The video introduces the speech commands and goes over an example interaction with a skill that is different from the ones taught in the experiment. After the video, participants are equipped with the microphone and referred to the whiteboard next to Simon for reminders of the speech commands. They are also asked to move Simon’s arm around to get a sense of how the arm behaves.

The experiment proceeds with one session for each skill. At the beginning of each session the experimenter places the related object in Simon’s hand. Then the participant provides two demonstrations of the skill, using the commands “New demonstration” and “End of demonstration” to denote the start and end of the movement. Simon uses speech and head nods to indicate that the person’s commands are heard. Additionally, Simon looks at his hand during a demonstration, and gazes back to the teacher when the demonstration is complete.

The participant is instructed to step away from the robot after the two demonstrations and use the command “Do you have any questions?” to trigger Simon’s query. Participants are told there is no restricted vocabulary or grammar for answering and they can say anything in response to Simon’s questions. Participants are not instructed about demo queries in advance, however they are prompted by the experimenter if they do not know what to do in response to demo queries. The experimenter advances the interaction with a keyboard stroke once the participant answers Simon’s question. This is repeated twice for each skill, thus each participant sees two queries for each query type. Then they rate the questions in a pen-and-paper questionnaire.

**Table 1:** Forced ranking results presented in terms of the number of participants out of 18 who voted the query type as best or worst.

|               | Smartness                |              | Informativeness           |              | Ease                      |              |
|---------------|--------------------------|--------------|---------------------------|--------------|---------------------------|--------------|
|               | <i>best</i>              | <i>worst</i> | <i>best</i>               | <i>worst</i> | <i>best</i>               | <i>worst</i> |
| Label         | 3                        | 6            | 4                         | 6            | 10                        | 4            |
| Demo          | 2                        | 11           | 8                         | 8            | 3                         | 7            |
| Feature       | 13                       | 1            | 6                         | 4            | 5                         | 7            |
| Friedman test | $\chi^2=13$ ,<br>p=0.001 |              | $\chi^2=0.12$ ,<br>p=0.94 |              | $\chi^2=3.11$ ,<br>p=0.21 |              |

### 5.2.2 Evaluation

Participants rate Simon’s questions on two 7-point Likert scale questions addressing *informativeness* of the question for the robot and *ease of answering* for the teacher. Each question has a comment box for explaining the particular ratings. Due to the sequential nature of the task these questions do not provide a reliable comparison of the three query types, however they allow the participant to elaborate about the evaluation criteria. For a more informed comparison of the query types, we ask three forced-ranking questions addressing *smartness* of the question, *informativeness* for the robot and *ease of answering* for the teacher at the end of the experiment.

In addition to the subjective evaluation by participants, we analyze the variance in the provided demonstrations in comparison to the information acquired through queries. We also provide a comparison of the time taken to pose and respond to queries. We examine the variety of answers provided by the participants in response to Simon’s queries.

## 5.3 Results

### 5.3.1 Subjective evaluation

Our experiment was completed by 18 participants (4 female, 14 male). Results from the forced ranking questions comparing the three query types are given in Table 1. We find a significant ranking in terms of *smartness* of the questions asked by Simon. 72% of participants thought that feature queries were the smartest, while 61% thought that demonstration queries were the least smart questions. Reasons stated by participants for choosing feature queries as the smartest include: being “more abstract”, demonstrating that “Simon understands task constraints at a high level, rather than just trajectories” and not involving “repeating the whole process of pouring coke”.

Although the query types were not found to be significantly different in terms of *ease of answering*, we see that 56% of participants thought label queries were the easiest to answer. A number of participants noted the fact that Simon’s questions were yes/no questions and thus easy to answer (7 mentioned for label queries, 4 for feature queries). Two participants explained why they thought feature queries were not as easy to answer as label queries: “The first [label] one was a yes-no (obvious), the coke [feature] question actually made me think” and “I could not answer the second question [pouring coke at different heights] with yes/no, it depended.” We find that a large fraction of the participants

**Table 2: Average time (in seconds) taken for the robot to make a query and for the participants to answer the query for the three types of queries.**

|         | Query           | Answer          |
|---------|-----------------|-----------------|
| Label   | 14.11 (SD=2.88) | 3.31 (SD=2.96)  |
| Demo    | 7.36 (SD=1.08)  | 24.07 (SD=5.38) |
| Feature | 9.56 (SD=0.49)  | 7.17 (SD=11.99) |


answered the robot’s question with a plain yes/no (17 for label queries, 15 for feature queries out of 18), while few gave explanations, *e.g.* “No, because the cereal would spill” or added emphasis, *e.g.* “Yes, orientation matters.” Some mentioned in the survey that they did not trust the robot would understand them and provided additional information about their answer within the survey (“I said yes, but that would only work for a full box of cereal”).

We do not see any consistency in the ratings in terms of informativeness of the questions for the robot. This seems to be a result of different interpretations in informativeness, which is revealed by participant comments: (label) “good question, in that the timing of when the box is tipped is important”, (demo) “the salt question gives more freedom for the robot to choose new actions”, (feature) “more general information, than a 3rd or 4th demonstration.”

### 5.3.2 Task measures

The problem of choosing between different types of queries in AL is often addressed by defining a *utility* for each query. This reflects the payoff between the benefits of the query (*e.g.* informativeness) and the costs of the query. The most common measure of cost used in the literature is wall-clock time. For queries in the LfD setting, this will depend on how fast the robot can move and how fast it speaks. Nevertheless, we present the time taken to make each type of query, as well as time taken for the teacher to complete their answers, for a comparison in Table 2.

We see that the ordering of the three types in terms of making the query is as expected: label queries take the most time as they require a complete execution of the skill, feature queries take less time (determined by the longest of the verbal question or the action performed for physical grounding), and demo queries take the least (determined by the longest of the verbal questions and the movement to reach the desired query start pose). For the answer we see that demo queries take much longer than the other two types of queries, as they require a whole demonstration. We observe that feature queries have a longer response time than label queries, even though both are yes/no questions and a very large variance for feature queries. This was due to a subset of the participants having trouble understanding Simon’s question and the interruption by the experimenter to repeat the question. This happened with four participants on the first feature query and one time on both feature queries. When these outliers are removed the average time to answer the query is 4.07 (SD=5.54) seconds which is comparable to the response time for label queries. We believe part of the reason was that the participants in all interruption cases were not native English speakers. However, this supports the result from the subjective evaluation that feature queries are



**Figure 7: The end-effector trajectories provided by all participants (thin lines). (a) (Pouring cereal) Superimposed with the label queries made by the robot (thick lines). (b) (Adding salt) Simon’s pose shows the second demo query starting point. Starts of the demonstrations are shown as dots.**

harder to interpret than label queries. Thus, they might require repetitions which adds to their cost.

As we did not commit to a particular learning method in this paper, it is not possible to make a direct comparison of the informativeness of each type of query. In the following, we try to illustrate how each type of query is informative for the learned skill. Fig. 7(a) provides a comparison of all demonstrations made by participants with the label queries made by Simon. We observe that the second label query (Fig. 6), which was chosen to be a valid way of pouring, is isolated from all the demonstrations provided by participants. We believe the reason that the queried trajectory was not covered by teacher demonstrations, was the difficulty of kinesthetic manipulation of the joints around this region. This highlights the benefits of label queries: certain ways of performing a skill might be valid even though they are unlikely to be demonstrated. Label queries can uncover such possibilities. In Fig. 7(b) we observe how the robot’s second demo query (Fig. 6) is different from the starting points of the demonstrations provided by humans. These illustrate how the different types of queries can quickly expand the applicability of a skill to unexplored regions.

## 6. DESIGN IMPLICATIONS

In designing interactions for active robot learners, the first question to be addressed is the choice of query types. Our results suggest primary usage of feature queries. These were the most common in human learning (82%) and were thought to be the smartest questions when used by the robot. On the other hand, these are the most challenging queries to produce automatically in a way that is interpretable by humans. Label queries are good alternatives that are easy to produce and are interpretable. They also have the advantage of very fast response and low effort requirements from the teacher. However, negative examples are likely to occur and are less informative in learning the skill. In addition, these negative examples can have undesirable consequences (*e.g.*, spilling cereal). One idea, inspired from the common human usage of label queries, is to do *partial* label queries. This would involve the robot performing smaller portions of a skill one at a time and asking whether the performance is correct so far after each segment.

In terms of the question form, we found that closed form

questions are desirable as they are common in human learning, and they facilitate the interpretation of the response by limiting possible answers. In addition, the fact that label and feature queries were yes/no questions was pointed out by participants as a positive property. However we observed that limiting the answers to yes/no could result in inaccurate information or missed opportunities for additional information. Thus, it is worth trying to build in the capability to interpret answers like “sometimes” or “possibly”. In addition, question forms can be used as a transparency mechanism that reveals the robot’s expectation and confidence about the anticipated answer.

Another important consideration in designing question asking behaviors is the physical grounding of the question. In particular skills are often represented with features that are hard to verbalize, but easy to instantiate (*e.g.* positions and rotations). For instance in our experiment, Simon was able to ask about the importance of “maintaining the end-effector orientation about the x-axis at about 90 degrees” without needing to explain these terms. This potential should be exploited whenever possible.

## 7. CONCLUSION

In this paper, we identify three types of queries for Active Learning, with which LfD methods can be augmented to achieve faster learning. We present two experiments; the first characterizing the use of these queries in human task learning and the second evaluating them in a human-robot interaction setting. We find that *feature queries* (explicit tests on certain aspects of a task) are the most common in human learning and are perceived as the smartest when used by the robot. In addition, we find humans use several physical actions to highlight or instantiate features that they want to make queries about. Based on these findings we provide several guidelines that can inform the design of question asking behaviors on a robot learner.

## Acknowledgment

This work is supported by the National Science Foundation, grant number IIS-0953181.

## 8. REFERENCES

- [1] B. Argall, S. Chernova, B. Browning, and M. Veloso. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.
- [2] M. Cakmak, C. Chao, and A. Thomaz. Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development*, 2(2):108–118, 2010.
- [3] S. Calinon and A. Billard. Statistical learning by imitation of competing constraints in joint and task space. *Advanced Robotics*, 23(15):2059–2076, 2009.
- [4] R. Cantrell, P. Schermerhorn, and M. Scheutz. Learning actions from human-robot dialogues. In *IEEE Intl. Symp. on Robot and Human Interactive Communication (RO-MAN)*, 2011.
- [5] C. Chao, M. Cakmak, and A. Thomaz. Transparent active learning for robots. In *ACM Intl. Conf. on Human-Robot Interaction (HRI)*, 2010.
- [6] S. Chernova and M. Veloso. Multi-thresholded approach to demonstration selection for interactive robot learning. In *ACM/IEEE Intl. Conf. on Human-Robot Interaction (HRI)*, 2008.
- [7] G. Druck, B. Settles, and A. McCallum. Active learning by labeling features. In *Empirical Methods in Natural Language Processing*, pages 81–90, 2009.
- [8] M. Gervasio and K. Myers. Question asking to inform procedure learning. In *AAAI Spring Symposium on Agents that Learn from Human Teachers*, 2009.
- [9] M. Gervasio and K. Myers. Learning to ask the right questions to help a learner learn. In *Intl. Conf. on Intelligent User Interfaces (IUI)*, 2011.
- [10] E. Gribovskaya, F. d’Halluin, and A. Billard. An active learning interface for bootstrapping robot’s generalization abilities in learning from demonstration. In *RSS Workshop Towards Closing the Loop: Active Learning for Robotics*, 2010.
- [11] D. Grollman and A. Billard. Donut as I do: Learning from failed demonstrations. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2011.
- [12] D. Grollman and O. Jenkins. Dogged learning for robots. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2007.
- [13] A. Guillory and J. Bilmes. Simultaneous learning and covering with adversarial noise. In *Intl. Conf. on Machine Learning (ICML)*, 2011.
- [14] N. Jetchev and M. Toussaint. Task space retrieval using inverse feedback control. In *Intl. Conf. on Machine Learning (ICML)*, 2011.
- [15] G. P. Kearsley. Questions and question asking in verbal discourse: A cross-disciplinary review. *Journal of Psycholinguistic Research*, 5:355–375, 1976.
- [16] R. Lomasky, C. Brodley, M. Arnecke, D. Walt, and M. Friedl. Active class selection. In *European Conf. on Machine Learning (ECML)*. 2007.
- [17] M. Lopes, F. Melo, and L. Montesano. Active learning for reward estimation in inverse reinforcement learning. In *ECML/PKDD*, 2009.
- [18] M. Muhlrig, M. Gienger, S. Hellbach, J. Steil, and C. Goerick. Automatic selection of task spaces for imitation learning. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2009.
- [19] J. Otero and A. Graesser. Preg: Elements of a model of question asking. *Cognition and Instruction*, 19(2):143–175, 2001.
- [20] H. Raghavan, O. Madani, and R. Jones. Active learning with feedback on features and instances. *J. of Machine Learning Research*, 7:1655–1686, 2006.
- [21] S. Rosenthal, A. Dey, and M. Veloso. How robots’ questions affect the accuracy of the human responses. In *IEEE RO-MAN*, 2009.
- [22] S. Rosenthal and M. Veloso. Modeling humans as observation providers using pomdps. In *IEEE RO-MAN*, 2011.
- [23] B. Settles, M. Craven, and L. Friedland. Active learning with real annotation costs. In *NIPS Workshop on Cost-Sensitive Learning*, 2008.