Learning Skills from Human Demonstrations
Initial Attempt on a Pouring Task

*Akihiko Yamaguchi (CMU,NAIST) Christopher G. Atkeson (CMU) Scott Niekum (CMU)
Tsukasa Ogasawara (NAIST)

1. Introduction

Many robots are designed for use in domestic environments where robots will be engaged in household chores. The robots need to learn ways to do the household chores that humans are now doing. We are taking a learning from demonstration (LfD) approach to this problem [1]. In terms of the household chores, a number of tasks are developed so far; for example, bringing a beer bottle from a refrigerator to a human, making pancakes [2], and folding towels [3].

However, a key issue for robots to do household chores is how to treat different versions of each task. Consider an opening task. There are a number of ways to open a container: rotating a cap on a plastic bottle, pulling a hinge cap of a ketchup bottle, pulling a pop-tab of a beer can, tearing a bag of potato chips, and so on. In addition, when opening a tight jar, we will use a different way to open it, like holding a cap with a wet towel. We call these methods *skills*. Learning these skills is essential for robots to fully handle tasks.

In this research, we treat a pouring task to study skill learning; its purpose is to move material. Humans use many skills to pour, such as shaking a bottle to pour viscous liquid like ketchup, tapping a bottle to pour a little amount of coffee powder, squeezing a shampoo bottle, and pushing a soap pump. Thus, the pouring task is a good example for robots to learn skills.

The goal of this research is making a general pouring behavior model from human demonstrations with which the robot can pour a wide variety of materials from a wide variety of containers. This problem is decomposed into three sub-problems:

(1) Deriving a model of a skill from human demonstrations in order to reproduce the skill. Each model will have some adjustable parameters to adapt to a specific situation like bottle size and material kind.
(2) Storing skill models and pairs of a situation description and the parameters for the situation. Then, selecting a skill for a new situation with estimated parameters.
(3) Correcting the selection of a skill if necessary, and adjusting the parameters to the new situation through actual executions.

An important first step to solving these sub-problems is finding important *features* that strongly affect each performance. For example in pouring, controlling the mouth edge of a bottle is better than controlling the gripper pose, as we discuss in Section 2. Another feature is an estimation of a flow trajectory. Though a number of methods are proposed so far in LfD research, we do not see a practical solution. The issue is that we have not formulated this kind of problem yet. Therefore, in this research, we model the behaviors manually by observing human demonstrations, then improve them through experiments. In this paper, we report our modeling of pouring behavior, a shaking skill, a tapping skill, and their implementation with the PR2 robot.

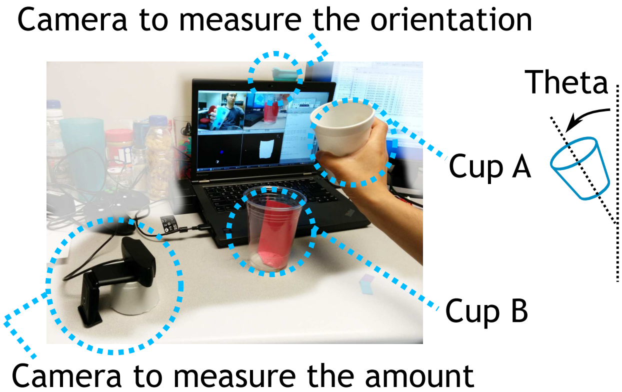

Fig. 1: Setup to measure a human demonstration.

There are several attempts for robots to learn pouring from human demonstrations [4, 5, 6]. However, they are focusing on a single pouring behavior. As far as we know, there is no method that has a capability to learn different versions of each task.

In Section 2, we model the behaviors from human demonstrations. In Section 3, we implement the skills in the PR2 robot. Section 4 concludes this paper.

2. Skill Modeling

In this section, we discuss suitable models of behaviors based on human demonstrations.

2.1 Human Demonstration of Pouring

First, we observe human demonstrations of pouring. We use the setup shown in Fig. 1 to track the human demonstrations. A human subject will pour from the cup A to the cup B where the orientation

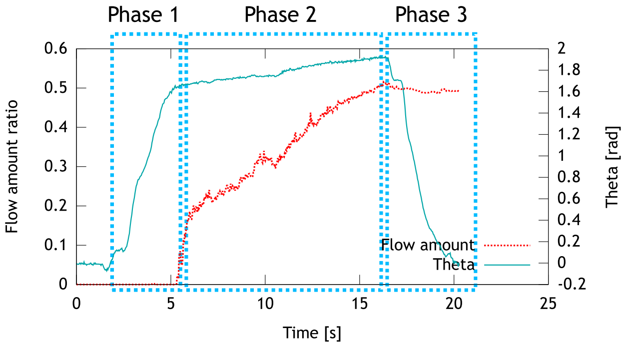

Fig. 2: Result of a human demonstration. The dotted curve shows the flow amount, and the solid curve shows the orientation (Theta) of the cup A.

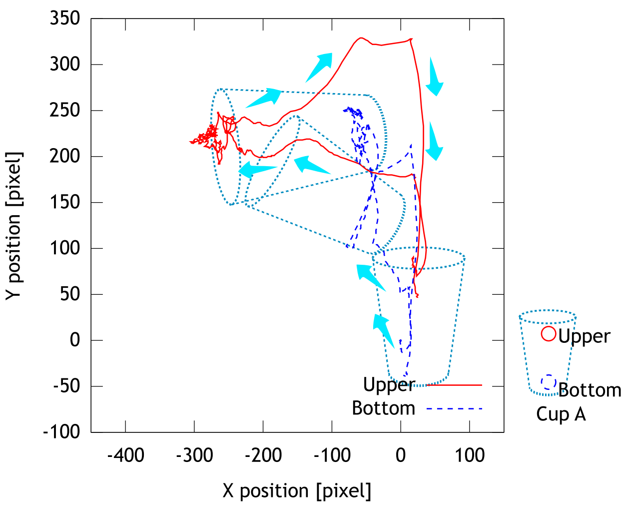

Fig. 3: Trajectories of markers on the cup A. Pouring is done where the x position of the upper marker is around -300.

of the cup A and the amount of material in the cup B are measured by RGB cameras. The material in the cup A is dried *peas* which behave like water, but are more convenient for measuring the amount and for real robot experiments. The human subject pours the material to a target amount of 0.5 which is half of the cup B.

Fig. 2 shows an obtained demonstration of pouring where the flow amount and the orientation are plotted. From this demonstration, we can see that there are three phases in pouring. Phase 1: rotating the cup A quickly until flow is observed. Phase 2: after flow starts, the human rotates the cup slowly until the amount reaches the target. We found that once the flow starts, it continues without rotating the cup so much, and thus the human was more careful. Phase 3: after reaching the target amount, the human moves the cup to the initial pose.

Fig. 3 shows the x and y trajectories of markers on the cup A. We can see that during pouring, the upper marker is moving little, on the other hand, the bottom marker is moving widely. The grasping point is around the middle of these markers, thus this point is also moving more widely than the upper marker.

2.2 Modeling Pouring

Next, we discuss how to model the pouring behavior for a robot. The whole pouring task consists of grasping a cup, moving it to the other cup, pouring the material from a cup to the other cup, moving the cup to an initial location, and releasing the cup. For simplicity, here we assume that the robot starts to pour when the robot is grasping a cup and holding it near from the other cup.

From the result of Fig. 3, we think that modeling the movement of a point on the mouth edge of the cup is easier than modeling the movement of the gripper trajectory. Thus, we assume that during pouring, the cup edge point takes a constant value, and only the orientation changes to control the flow. The cup moves mostly in a 2-dimensional plane, so the orientation is modeled by a 1-dimensional variable, θ. When the robot is grasping a cup, the position of the cup edge point is constant in the local frame of the gripper. Thus, controlling the cup edge point is achieved by a standard inverse kinematics solver.

As we mentioned above, we found three phases in the human demonstration. A simple way to model this kind of behavior is using a finite state machine. When no flow is observed, the robot increases θ (Phase 1). If flow is observed, the robot slows down the movement (Phase 2). If the target amount is achieved, the robot moves θ to the initial value (Phase 3).

Through some demonstrations, we found that if the material starts flowing, it continues to flow without increasing θ. Thus, in Phase 2, we increase θ only when the flow is not observed, and keep the same value when the flow is observed.

2.3 Shaking

Humans sometimes shake a bottle to pour when the material is jamming inside the bottle or the material is a viscous liquid. Though there are a variety of shaking behaviors, we simply model the shaking behavior as a vertical motion while holding the bottle upside down.

2.4 Tapping

Tapping is used to pour material accurately; for example, pouring coffee powder. Reproducing such a motion with a robot is a bit difficult unless the robot can move the gripper rapidly. Our modeling is touching the free gripper to the cup held by the other gripper, then vibrating the gripper.

3. Experiments

We implement the pouring skills modeled in the previous section on a robot, PR2. The PR2 has two

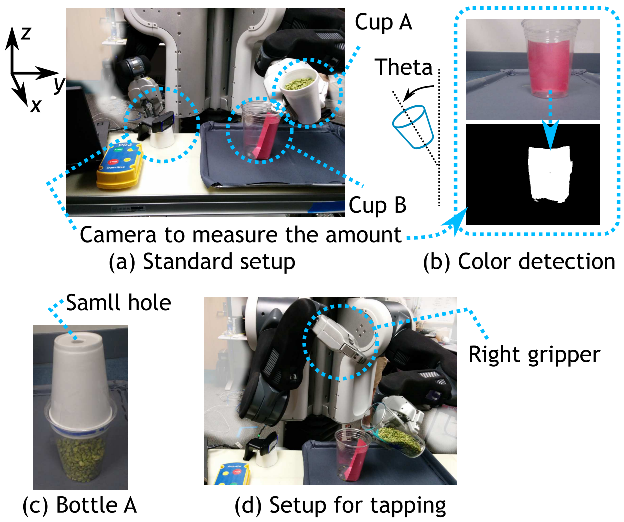

Fig. 4: Setup of the experiments.

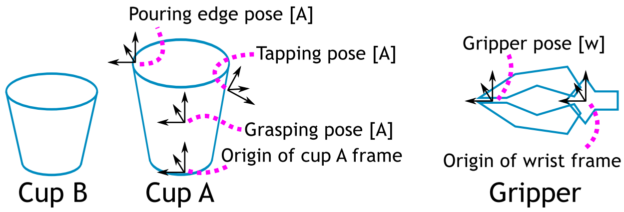

Fig. 5: Object specific vectors. A pose means a xyz position and a quaternion. [A] and [w] denote a vector defined in the cup A and the wrist frames respectively.

7-degrees of freedom arms with grippers. Fig. 4(a) shows the setup of the robot and the cups. We use dried peas to avoid hardware damage by liquid. In order to measure the amount of poured material, we use an RGB camera and detect specific colors as shown in Fig. 4(b). The ratio of colored areas is used as the amount.

We use ROS packages for the PR2 to implement low-level control and inverse kinematics solver of the grippers. As mentioned in Section 2.2, in order to control the pose of the pouring edge point with the 1-dimensional variable θ, we define several object specific vectors as illustrated in Fig. 5. The pouring edge pose [A] and the grasping pose [A] are constant vectors in the cup A frame; these are defined for each cup and bottle from which the robot pours the material. Since during grasping, the gripper pose corresponds to the grasping pose, we can compute the pouring edge pose in the wrist frame. An inverse kinematics solver for the wrist link is implemented in a ROS package, thus we can control the pouring edge pose. From the initial pose of the cup A, we can estimate the rotation axis.

Fig. 6 shows the result of pouring where the target amount is 0.5. In this graph, only Phase 1 and 2 are

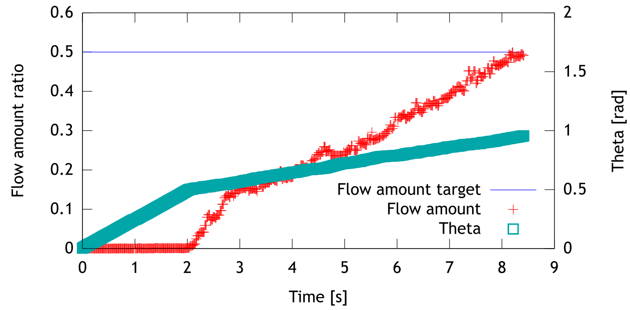

Fig. 6: Result of pouring.

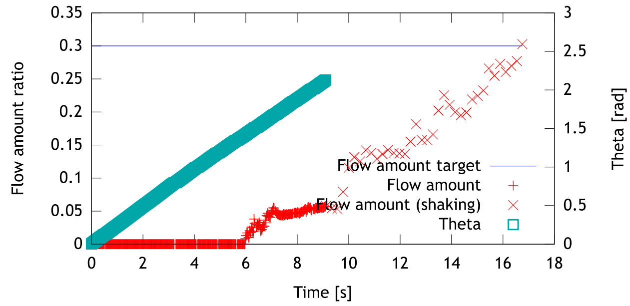

Fig. 7: Result of the shaking.

plotted; after achieving the target amount, the robot moves θ to the initial value. The flow started around 2 [s], then the robot slows down to move θ. Comparing to the human demonstration (Fig. 2), the robot behavior has similar structure. However, there are some small differences. For example, the human slows down the angular velocity before the flow starts. One possible reason is that the human tries to keep the initial flow small. Estimating the orientation where the flow starts is necessary to reproduce this behavior; humans are using visual information and/or force information. We are ignoring this behavior, but it seems to be working practically.

Shaking is useful when the material is jamming, so we use a bottle with a small hole as shown in Fig. 4(c). The peas are contained in the bottle. The robot grasps the white part, and the initial pose is similar to the standard pouring case. Fig. 7 shows the result of pouring with shaking. The orientation θ is plotted only while the robot rotates the bottle to turn it upside down. After this phase, the shaking motion starts. We can see a little flow during this phase (around 6 [s]), but the flow stops due to the jammed material, so standard pouring does not work any more. During shaking, we can see the amount is increasing. Thus, shaking is also a possible way for the robot to solve jamming.

Next, we investigate the performance of tapping. This skill uses the free gripper to tap, so we start from the setup shown in Fig. 4(d). The robot grasps a cup with the left gripper while the right gripper stays above the cups. In order to touch the right

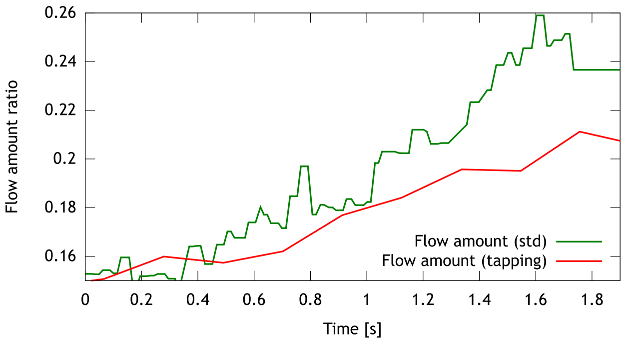

Fig. 8: Comparison of the standard pouring and tapping.

gripper to the cup, we define a tapping pose as illustrated in Fig. 5, which is a constant vector in the cup frame. Fig. 8 shows a part of trajectories of the standard pouring and the tapping. From this comparison, we can see that using tapping, the robot can pour the material slowly; in addition, the noise of amount estimation from the camera is reduced. Thus, tapping enables the robot to pour accurately. However, we have not obtained a successful result yet; since the amount of initial flow is large, the accuracy of pouring with tapping was not good.

4. Conclusion

In this paper, we investigated a way for robots to learn pouring from human demonstrations. Pouring, shaking, and tapping were modeled with finite state machines, and implemented on the PR2 robot. The real robot experiments were mostly successful. Future work is expanding the coverage of skill and making a knowledge database of skill and situations in order to enable a robot to pour a wide variety of materials from a wide range of containers.

References

[1] A. Billard and D. Grollman, "Robot learning by demonstration," *Scholarpedia*, vol. 8, no. 12, p. 3824, 2013.

[2] P. Kormushev, S. Calinon, and D. G. Caldwell, "Robot motor skill coordination with EM-based reinforcement learning," in *the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'10)*, pp. 3232–3237, 2010.

[3] J. Maitin-Shepard, M. Cusumano-Towner, J. Lei, and P. Abbeel, "Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding," in *the IEEE International Conference on Robotics and Automation (ICRA'10)*, pp. 2308–2315, 2010.

[4] M. Mühlig, M. Gienger, S. Hellbach, J. J. Steil, and C. Goerick, "Task-level imitation learning using variance-based movement optimization," in *the IEEE International Conference on Robotics and Automation (ICRA'09)*, pp. 1177–1184, 2009.

[5] M. Tamosiunaite, B. Nemec, A. Ude, and F. Wörgötter, "Learning to pour with a robot arm combining goal and shape learning for dynamic movement primitives," *Robotics and Autonomous Systems*, vol. 59, no. 11, pp. 910–922, 2011.

[6] L. Rozo, P. Jiménez, and C. Torras, "Force-based robot learning of pouring skills using parametric hidden markov models," in *the IEEE-RAS International Workshop on Robot Motion and Control (RoMoCo)*, 2013.