
Augmenting Kinesthetic Teaching with Keyframes

Baris Akgun
Maya Cakmak
Jae Wook Yoo
Andre L. Thomaz

Georgia Institute of Technology, Atlanta, Georgia 30332 USA

BAKGUN@GATECH.EDU
MAYA@CC.GATECH.EDU
JYOO@GATECH.EDU
ATHOMAZ@CC.GATECH.EDU

Abstract

Our work is aimed at designing Learning from Demonstration (LfD) systems for non-expert teachers. We compare two teaching methods for LfD, namely *trajectory method*, in which whole demonstrated trajectory is used, and *keyframe method*, in which a sparse set of consecutive configurations is used, with a user study. We use kinesthetic teaching as our demonstration approach in which the teacher physically guides the robot to perform a certain skill. Based on our results and observations, we then propose a hybrid method for skill learning that utilizes both trajectory and keyframe information.

1. Introduction

Robot Learning from Demonstration (LfD) deals with the challenges of enabling humans to program robot skills by simply demonstrating their successful execution (Argall et al., 2009). We are particularly interested in kinesthetic teaching in which the human teacher physically guides the robot to perform the skill, as in figure 1. Kinesthetic teaching has several advantages considering the common issues in LfD. Since the teacher is directly manipulating the robot there is no correspondence problem and demonstrations are restricted to the kinematic limits (e.g. workspace, joint limits) of the robot. Kinesthetic teaching can also pose several challenges for everyday users who do not have experience manipulating a robot arm with many degrees of freedom.

While there has been a large body of work concentrated on representations and learning algorithms, the




Figure 1. A kinesthetic teaching interaction between the robot and a human teacher.

usability of kinesthetic teaching has not been explored in depth. In many of the practical LfD applications, the teacher will not be an expert in machine learning or robotics. Thus, our research explores the ways in which Machine Learning can exploit human social learning interactions—*Socially Guided Machine Learning* (SG-ML)—and aims at identifying challenges that everyday users face when trying to interact with a robot using the common methods in the field. We then explore improvements to these methods to increase their usability.

The common approach in kinesthetic teaching, and most LfD interactions, is for each demonstration to be an entire state *trajectory*, which involves providing a continuous uninterrupted demonstration of the skill to the robot. An alternative is to provide a sparse set of *keyframes* that achieve the skill when connected together. We first present an experiment that evaluates these approaches. We find that both are

suitable to kinesthetic teaching from the user’s perspective. Keyframe demonstrations mitigate temporal alignment issues that can arise with multiple trajectory demonstrations but discard the useful velocity information present in a full trajectory. Based on these results, we propose an alternative view of keyframes, as points that highlight task and goal constraints. We suggest a new interaction scheme that we think is more natural to the user and exploits both keyframe and trajectory interaction methods to provide a novel approach to skill generalization.

2. Related Work

In kinesthetic teaching, demonstrations are often represented as arm joint trajectories and/or the end-effector path (Calinon & Billard, 2009; Hersch et al., 2008). These are referred to as the robot’s position in *joint* and *task* space respectively. Some also consider the position of the end-effector with respect to the target object of the skill (Billard et al., 2006; Gribovskaya & Billard, 2009).

Typically, start and end points of a demonstration are explicitly demarcated by the teacher. Most studies subsample the recorded data with a fixed rate (Billard et al., 2006; Amor et al., 2009). Demonstrations are often time warped such that a frame-by-frame correspondence can be established between multiple demonstrations (Hersch et al., 2008). Another recent approach is to only record keyframes and learn a manifold that models the behavior with them (Bitzer et al., 2010). In this paper we consider both trajectory and keyframe representations.

Human-robot interaction has not been a focus of prior work on kinesthetic teaching, but there are a few examples. In (Weiss et al., 2009), kinesthetic teaching is embedded within a dialog system that lets the user start/end demonstrations and trigger reproductions of the learned skill with verbal commands.

An interesting modification to the kinesthetic teaching interface is kinesthetic correction (Calinon & Billard, 2007a;b), where the teacher corrects aspects of a learned skill in an incremental learning interaction. The teacher selects joints to control manually in the next execution. The selected motors are set to a passive mode, allowing the user to control them while the robot executes the learned skill with the other joints. In another study (Lopez Infante & Kyrki, 2011), four types of force controllers that effect the response to users are evaluated for kinesthetic teaching. The study addressed human preferences on which controller was the most natural.

While various learning methods for kinesthetic teaching have been explored, there is a lack of studies with end-users testing the effectiveness of these techniques in terms of human-robot interaction. This is the focus of our work.

3. Initial Implementation

In this section, we describe our initial implementation of skill teaching interactions by *trajectory demonstrations* and *keyframe demonstrations*. In all the interactions, the human teacher manipulates the right arm of the robot (figure 1).

3.1. Trajectory Demonstrations

3.1.1. INTERACTION

The teacher is informed that the robot will record all the movement they make with its right arm. The teacher initiates the demonstration by saying “New demonstration”, moves the arm to make the robot perform the skill and finishes by saying “End of demonstration.” This process is repeated to give as many demonstrations the user desires. After a demonstration, the teacher can use the speech command “Can you perform the skill?” to have the robot perform the current state of the learned skill and adjust his/her demonstrations to attend to any errors.

3.1.2. LEARNING

Joint angle trajectories are recorded as the teacher moves the robot’s arm to perform the skill. We use the LfD method described in (Calinon et al., 2007) to learn the skill as described below.

The data is subsampled in the time dimension to a constant length before being input to the learning algorithm. In our approach learning is done in an eight dimensional space which incorporates the joint angles (7 dimensions) and time (1 dimension). As a first step, we use the K-means algorithm with a constant k . The resulting clusters are then used to calculate initial mean vectors and covariance matrices for the expectation-maximization (EM) algorithm. The EM algorithm is run to extract a Gaussian-Mixture Model (GMM) from the data. Note that the resulting GMM has k number of sub-populations which is kept constant during our experiments. Gaussian-Mixture Regression (GMR) is used to generate a trajectory to perform the learned skill. The desired time dimension vector is given to GMR which in turn generates the joint positions. The algorithm can learn from multiple demonstrations.

3.2. Keyframe Demonstrations

3.2.1. INTERACTION

The teacher is informed that the robot will only record the arm configuration when they say “Record frame”, and it will not record any movements between these keyframes. The teacher can use the speech commands “New demonstration”, “End of demonstration” and “Can you perform the skill?” in the same way as in trajectory demonstrations.

3.2.2. LEARNING

The resulting data from this technique is a sparse trajectory of joint angles. If the teacher forgets to give keyframes for the start and/or the end position (which are assumed to be the idle position of arm down to the robot’s side), these are added automatically. Then time information is generated for each keyframe using the distance and a constant average velocity between keyframes.

Learning is slightly different than the previous case, but the space is the same. As a first step, again K-means is run but this time the number k is chosen as the maximum number of keyframes across all demonstrations provided for a skill. The GMM part is the same as the trajectory version. For generating the skill, the GMM sub-population means are traversed by splining between them. We took such an approach since the GMM sub-population means obtained from the keyframe version will be of different nature than the ones obtained from the trajectory version. The former is more likely to be a transition between two trajectory segments whereas the latter is more likely to be a mid-point of a trajectory segment (Calinon et al., 2007). Thus, we want to control the velocity ourselves at each keyframe.

4. Experiments

In this section, we describe our experiments for testing the aforementioned two methods of giving demonstrations, to answer the questions: (1) When everyday people teach the robot, what are the effects of keyframe as opposed to trajectory demonstration? (2) Is there a distinction in the learning performance when learning different types of skills?

4.1. Experimental Design

We differentiate between two types of skills. *Goal-oriented* skills are related with achieving a particular world state (e.g., finger tip on a point while avoiding obstacles.) *Means-oriented* skills, on the other hand,

are not related to a world state but are related to the style of the motion. We chose four skills from each skill type.

Goal-oriented skills are as follows. **Insert**: inserting the block in hand through the hole without touching other blocks. **Stack**: stacking the block in hand on top of another block on the table. **Touch**: touching a certain point with the tip of a finger. **Close**: close the lid of a box without moving the box.

And means-oriented skills are as follows. **Salute**: perform a soldier’s salute. **Beckon**: perform a beckoning gesture, as in asking someone to come closer. **Raise-hand**: raise the robot’s hand as if it is asking for permission. **Throw**: perform a throwing gesture with a ball (without actually releasing the ball).

Our experiment has two conditions, *Trajectory Demonstrations (TD)* and *Keyframe Demonstrations (KD)*, which correspond to the teaching methods explained in Sec. 3. We use a within-subject study design, *i.e.* participants interact with the robot in both conditions.

We counterbalance the order of the conditions and vary the type of skill taught to the robot across the conditions. Each participant taught one *means-oriented* skill, and one *goal-oriented* skill. As practice, participants teach a *pointing* skill to the robot for familiarization with the condition prior to the other skill demonstrations. In addition, the participant is allowed to move the robot’s arm to practice each skill.

4.2. Measures

4.2.1. SUBJECTIVE EVALUATION

We asked 7-point Likert-scale questions, administered after each condition. The questions are about *Feel*, *Naturalness*, *Ease*, and *Enjoyability*. We also asked open-ended questions at the end of the experiment.

4.2.2. QUALITY OF THE LEARNED SKILLS

The performance of goal-oriented skills are scored separately by two of the authors, using three levels of success criteria (Success-PartialSuccess-Fail). The scoring is based both on recorded videos of the experiment and on skill performances recreated on the robot. If there is a conflict between the coders, they revisit the example and reach a consensus on the scoring.

Unlike the goal-oriented skills, success for *means-oriented* skill is subjective. Therefore, expert ratings of the recreated movements are used to evaluate the performance. The experts, whose specialities are in animation, are asked to answer three 7-point Likert-

scale questions. The questions are about *appropriate emphasis*, *communicating intent*, and *closeness to perfection*.

4.2.3. INTERACTION METRICS

We also measured the number of demonstrations, the number of keyframes, the time stamps for every event, and all trajectories of the joint movement during demonstrations.

5. Results

We conducted a user study with 34 participants (6 females, 26 males between the ages of 19-47). The means-oriented skills taught by participants were rated by two animation experts.¹

5.1. Comparison of trajectory and keyframe demonstrations

5.1.1. SINGLE DEMONSTRATION IS COMMON

Fig. 2 shows the number of demonstrations provided by participants in TD and KD. We see that teaching with a single demonstration was common in both modes. For goal-oriented skills, a larger portion of the participants provided a single demonstration in the TD condition than in the KD condition (19 versus 10). It was common in the KD condition to forget to provide keyframes that allow the robot to avoid obstacles while trying to achieve the goal. These frames were often provided by participants in the consequent demonstrations after observing the performed skill colliding with obstacles. An example of this is shown on Fig. 3. For means-oriented skills, more participants taught with a single demonstration in the KD condition than in the TD condition (31 versus 26).

5.1.2. TRAJECTORY MAY BE MORE SUCCESSFUL WHEN TEACHING GOAL-ORIENTED SKILLS IN A SINGLE DEMONSTRATION

Table 1 provides the distribution of participants according to the success of the goal-oriented skills that they taught². When the skill is taught with a single

¹We note that our observations reported in this section did not vary across particular skills (*i.e.* insert, close, stack, *etc.*).

²In reporting our evaluation of success of goal-oriented skills we chose to present the exact distribution of all participants and descriptive statistics as opposed to assigning a value to each success level and reporting averages and Wilcoxon signed rank test results, since such results would depend on our arbitrary choice of number of levels and values assigned to each level.

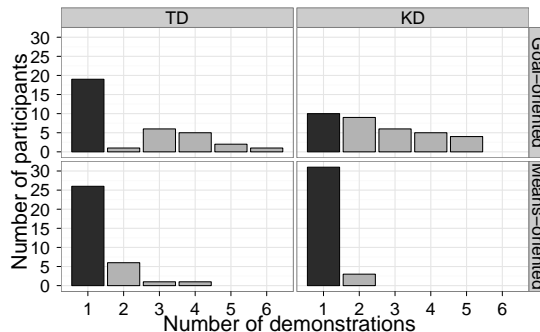


Figure 2. Histogram of number of demonstrations provided by participants in KD and TD conditions.

Table 1. Number of participants who achieved different levels of success in all conditions.

COND.	DEMO.	SUCCESS	PART. SUCC.	FAIL
TD	SINGLE	15	4	1
	MULTIPLE	1	5	8
	Total	16	9	9
KD	SINGLE	5	5	1
	MULTIPLE	4	9	10
	Total	9	14	11

demonstration, more participants achieve full success in TD as opposed to KD (15 versus 5).

The large number of single demonstration instances is an artifact of our experimental design. The skills used in our experiments were chosen to be fairly easy to achieve and participants were allowed to practice a particular skill before providing an actual demonstration of the skill. We observed that this practice opportunity was used more in the TD condition, where people often practiced enough to be able to teach the skill in a single demonstration.

As mentioned earlier, participants often do not think of providing keyframes for obstacle avoidance in their first demonstrations. In some cases this does not effect the success of skill in terms of achieving the goal (*i.e.* partial success) and participants could be satisfied by this since they were not explicitly told to avoid collisions. We see that a large portion of the participants who provided a single demonstration in the KD condition at least achieved partial success.

5.1.3. MULTIPLE TRAJECTORY DEMONSTRATIONS CAN RESULT IN ALIGNMENT PROBLEMS

Only 1 participant out of the 14 who provided multiple demonstrations in the TD condition, was able to




Figure 3. An example of forgetting obstacle avoidance keyframes in the first demonstrations (dashed line), and providing them in the second demonstration (solid line) in the KD condition while teaching the *Touch* skill.

achieve success with the goal-oriented skill (Table 1). We observe that participants have trouble providing multiple trajectory demonstrations that have proper temporal alignment. As a result, parts of the demonstrations that are intended as different parts of the skill get averaged together across demonstrations. On the other hand, our learning approach for keyframe demonstrations handles the alignment problem. Note how the middle keyframes are automatically aligned in Fig. 3. As a result, more participants were able to achieve success with multiple demonstrations in the KD condition.

We saw that 3 participants in the TD condition achieved success after their first demonstration, however the learned skill ended up failing after being combined with subsequent refining demonstrations.

5.1.4. KEYFRAME DEMONSTRATIONS MAY RESULT IN PREFERABLE MEANS-ORIENTED SKILLS

Table 2 gives a summary of the ratings provided by experts for the means-oriented skills taught by participants. Both experts rated the means-oriented skills learned in KD condition higher in all three dimensions on average. The difference was only significant for *closeness to perfection*, and the difference is marginally significant when the three scales are averaged ($Z=2740$, $p=0.06$ on Wilcoxon signed rank test). This distinction is partly related to the difficulty of moving a 7-DOF arm smoothly in the TD condition. In addition, the temporal alignment issue mentioned earlier for goal-oriented skills, also exists for the few participants who provided multiple demonstrations for means-oriented skills.

5.1.5. PARTICIPANTS LIKE BOTH METHODS

Analyzing the Likert responses given by participants, we found that all ratings were biased towards higher values, and none of the scales showed a statistical difference between TD and KD (based on paired

Table 2. Expert ratings of means-oriented skills: Median and Coefficient of Dispersion (given in parentheses)

COND.	EXP.	EMPH.	INTENT	PERF.
TD	1	5.5 (0.27)	5 (0.33)	5 (0.35)
	2	3 (0.29)	3.5 (0.38)	4 (0.3)
KD	1	6 (0.21)	6 (0.17)	6 (0.2)
	2	4 (0.21)	4 (0.24)	5 (0.22)
COMP.		$Z=2679$, $P=0.10$	$Z=2677$, $P=0.11$	$Z=2796$, $P=0.03$

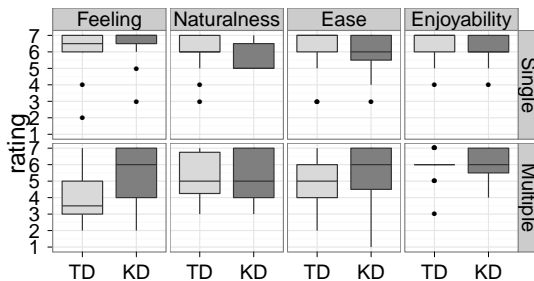


Figure 4. Subjective ratings of TD and KD conditions for goal-oriented skills separated by the number of demonstrations provided by the participant.

Wilcoxon signed rank tests).

We observe that participants’ ratings are correlated with their success in teaching the goal-oriented skills ($r=.31$, $p<.001$ in Spearman’s rank correlation test, assuming Fail:1, Partial:2 and Success:3). As a result, when the participants are grouped into ones that provide a single demonstration and ones that provide multiple demonstrations, we find that participants who provided multiple demonstrations felt more comfortable with keyframe demonstrations and thought that they were easier to use (Fig. 4). We do not observe a difference in the preference of participants who provided single demonstrations in these dimensions.

5.2. Comparison of goal-oriented and means-oriented skills

5.2.1. DIFFERENT OBJECTIVE FUNCTIONS ARE USED FOR EACH SKILL TYPE

As observed in Fig 2, a much larger fraction of participants provide a single demonstration for teaching means-oriented skills both in TD and KD conditions.

Across both conditions, the average number of demonstrations provided for goal-oriented skills (2.37, $SD=1.45$) is significantly larger than the number

of demonstrations provided for means-oriented skills (1.22, SD=0.53) ($t(84)=6.18$, $p<0.001$ on t -test). This highlights a fundamental difference between goal-oriented and means-oriented skills: while goal-oriented skills have a well defined objective function, means-oriented skills are subjective and underspecified.

5.2.2. CHARACTERISTICS OF PROVIDED KEYFRAMES ARE DIFFERENT FOR EACH SKILL TYPE

The average distance between keyframes in the 7DOF joint space for goal-oriented skills is much smaller (approx. 47% smaller) than the average distance for means-oriented skills ($t(38)=-3.94$, $p<.001$ on t -test). We hypothesize that participants are providing different types of keyframes within a single demonstration. For goal-oriented skills we see a distinction between keyframes that achieve the instrumental goal of the skill, and the keyframes that let the robot avoid obstacles. Similarly in means-oriented skills we see a distinction between keyframes that actually give the skill its meaning and make it recognizable, whereas other keyframes are waypoints for getting to these frames. Participants provide a large number of frames that are close to one another around the target of goal-oriented skills. For means-oriented skills, on the other hand, they provide less frames that are separated by a larger distance. For both types of skills the waypoint keyframes or obstacle avoidance keyframes tend to be further apart. We observe that the average number of keyframes for goal-oriented skills (6.75, SD=1.89) is not statistically different from that for means-oriented skills (6.21, SD=2.17) ($t(65)=1.11$, $p=.27$ on t -test).

6. Discussion

6.1. Advantages of keyframe-based LfD

Our results indicate that trajectory demonstrations were preferable for goal-oriented skills when the skill is taught with a single demonstration while they were likely to suffer from temporal alignment issues when teaching with multiple demonstrations. While this issue can be reduced with temporal alignment algorithms or with practice by the teacher, it cannot be fully eliminated for non-expert users. Furthermore, teaching with a single demonstration is unrealistic when learning needs to generalize across different goals (*e.g.* learning to insert an object through a hole at any location or orientation). The learning approach used in this paper for trajectory demonstrations was originally evaluated with multiple demonstrations provided by its developers (Calinon & Billard, 2009). In our experiment the goals were kept constant to limit the total duration of the user-studies given the number of

other factors that were varied. Our results showed that keyframe demonstrations were good at handling multiple demonstrations, even when some of the demonstrations were failing (Fig. 3).

Another advantage of keyframe demonstrations is that they temporally segment the demonstration. The segmentation of a demonstration into keyframes could be leveraged by considering different *types of keyframes*. We saw that such types emerged in our experiments, *e.g. goal keyframes* and *waypoint keyframes* in goal-oriented skills. Such labels for keyframes could easily be obtained from the teacher and used in different ways. For instance, the precision on the goal could be improved by using a smaller clustering threshold for *goal keyframes* since they are expected to be closer to one another.

In addition to these, the keyframe-based approach (i) allows a more compact representation of demonstrations, (ii) can easily be combined with planning, (iii) eliminates jerky movements of trajectories provided by humans, and (iv) reduces the amount of practice before demonstrations and mental workload during demonstrations required from the teacher.

6.2. Limitations of keyframe-based LfD

An important limitation of keyframe-based demonstrations used in this study is that the timing or speed of movements could not be controlled by the teacher. We observed that some participants tried to achieve slower movements or stops by providing a large number of very close or overlapping keyframes. In addition, several participants mentioned adding speed related commands in their suggestions.

A common problem with keyframe demonstrations was that participants could give demonstrations that are potentially harmful to the robot (*e.g.* colliding with the table). They did not predict that the robot would move straight to the goal, because they moved it to the goal through a non-colliding path. This shows that keyframe-based demonstrations might not be very intuitive at first. However, participants were able to recover from these colliding demonstrations, and they did not repeat the same mistake on later demonstrations or skills.

7. A Follow-up Approach

From the aforementioned study, we found that kinesthetic teaching is a viable method of interaction with a humanoid robot for everyday people. Additionally we find that both keyframe and trajectory demonstrations have their advantages.

Many skills require a combination of getting from one place to another with a certain configuration in the end state, and then executing a particular trajectory. An example is scooping something from a bowl. The spoon gets above the bowl and assumes a certain orientation, and “scoop” is executed. It is unnecessary to tie a particular trajectory to getting above the bowl thus a keyframe would fit perfectly. However, scooping requires a nonlinear motion and it would be best described by a trajectory.

Combining our results and observations, we suggest a new interaction scheme for LfD which merges trajectory and keyframe demonstrations in a single interactive demonstration.

7.1. Interaction

We want to allow people to teach skills in an intuitive way and think that the ability to provide both keyframe and trajectory information in the context of a single demonstration will be useful for a variety of skills and even combination of skills (e.g. scooping and then serving).

Towards this end, we propose a *hybrid* interaction scheme in which the teacher is able to give keyframes and trajectory segments in a single demonstration. After the demonstration is initiated, the teacher can provide a keyframe by moving the arm to a desired position and saying “Go Here”. This can be repeated multiple times. At any point, the user can say “Like this” to initiate a trajectory demonstration and say “That is it” to stop it. The teacher can combine these in any order.

Note that with this approach, the teacher can give pure keyframe, single trajectory, multiple (or segmented) trajectory and full hybrid demonstrations.

7.2. Learning and Representation

Since goal-oriented skills are performed with respect to an object, we choose to represent our trajectories in the object-centric coordinates and time (7 DOF). This has the potential to lead to better skill generalization.

With hybrid interactions, resulting demonstrations have multiple segments. We assume the end points of these segments to be correlated for different demonstrations.³

³For example, if a demonstration is trajectory-trajectory-trajectory (T-T-T), we expect the following to be identical or at least have the end points correlated if the number of segments are different (e.g. if the next demonstration is T-T, we expect that one of the Ts be a combination of corresponding two Ts in the previous demonstra-

One of the challenges of this approach is to find an alignment between these segments. To do this, we employ keyframe learning (as mentioned previously) to keyframes and the start and end points of the trajectory segments. The resulting cluster memberships denote the alignment of the segments. Then we perform trajectory learning to trajectory segments and keyframe learning to keyframe segments.

7.3. A Pilot Study

A pilot study was performed with three non-expert users (who are robotics students but are not familiar with the robot and the algorithm) on the robot Simon. The subjects demonstrated the **Close** and the **Insert** skills in pure keyframe, pure trajectory and segmented-trajectory modes. Then they demonstrated **Scoop and Serve** skill in hybrid mode which includes scooping M&M’s from a bowl with a spoon and pouring into another one. All the skills were done for three goal configurations and the subjects demonstrated each skill twice per goal configuration. There was no playback of the learned trajectory. The users were told that robot would move linearly between keyframes. The users were not notified about the learning algorithm’s assumption. The aim of this pilot study was to test the usability of the approach and to analyze the habits of non-experts. A follow-up interview was done after the experiment.

We observed that users were able to quickly adapt to the interaction scheme (the study took around an hour). The users tend to segment the skills into many portions during their first demonstrations, but decrease this number as they get used to the method. One dramatic example is that a user gave 11 segments on the first demonstration and gave 5 on the third.

For some of the demonstrations of a skill, users gave different corresponding segments (e.g. keyframe and trajectory portions overlap). When asked, two of the users said they did this on purpose, to show the robot various ways of achieving the skill. This hints that a mechanism might be needed to split up demonstrations to learn multiple models of a single skill.

The responses of the users to the interview show that the interaction scheme is usable and easy to get used to. Moreover, the users explicitly mentioned to use keyframes when how to get to a point was not important or a linear motion was fine whereas they used trajectories when the skill required a complex/nonlinear trajectory. An interesting response was that one of the users chose to move from one bowl to the other with

a trajectory to ensure that the M&M's do not spill on the way.

8. Conclusions

While trajectory demonstrations are commonly used for kinesthetic teaching and for LFD in general, in this work we also consider keyframe demonstrations. We have presented a user-study with non-expert human teachers that evaluates these two kinesthetic teaching methods for various skills. We found that trajectory demonstrations are indeed intuitive for non-experts but have the potential to run into problems such as time alignment. Although we have not done user-studies, we expect that teaching with multiple goal configurations would impact trajectory demonstrations negatively. We found that using keyframe demonstrations is a potential remedy to these issues but they are not intuitive and have their own disadvantages.

Following this we have propose a framework to combine both keyframe and trajectory methods into a single demonstration interaction. We argue that this hybrid approach combines the advantages of both methods and mitigates their disadvantages. For example, using keyframes as navigation nodes provide the potential of incorporating motion planning for obstacle avoidance. Learning keyframe positions with respect to objects decreases the space that the trajectory learning needs to cover. This also helps with the generalization problem. Having the ability to combine keyframes with trajectories and to give multiple trajectory segments results in shorter trajectories which in turn helps learning and the time alignment issues.

We are in the process of further analyzing our approach quantitatively with expert demonstrations to further back our claims. Additional future work includes a user study that would test the usability of our approach for non-expert users.

References

- Amor, Heni Ben, Berger, Erik, Vogt, David, and Jun, Bernhard. Kinesthetic bootstrapping: Teaching motor skills to humanoid robots through physical interaction. *Lecture Notes in Computer Science: Advances in Artificial Intelligence*, 58(3):492–499, 2009.
- Argall, B., Chernova, S., Browning, B., and Veloso, M. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.
- Billard, A., Calinon, S., and Guenter, F. Discriminative and adaptive imitation in uni-manual and bi-manual tasks. *Robotics and Autonomous System*, 54(5):370–384, 2006.
- Bitzer, S., Howard, M., and Vijayakumar, S. Using dimensionality reduction to exploit constraints in reinforcement learning. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pp. 3219–3225, 2010.
- Calinon, S. and Billard, A. Incremental learning of gestures by imitation in a humanoid robot. In *In Proc. of the ACM/IEEE International Conference on Human-Robot Interaction*, pp. 255–262, 2007a.
- Calinon, S. and Billard, A. What is the teacher's role in robot programming by demonstration? - Toward benchmarks for improved learning. *Interaction Studies. Special Issue on Psychological Benchmarks in Human-Robot Interaction*, 8(3), 2007b.
- Calinon, S. and Billard, A. Statistical learning by imitation of competing constraints in joint space and task space. *Advanced Robotics*, 23(15):2059–2076, 2009.
- Calinon, S., Guenter, F., and Billard, A. On learning, representing and generalizing a task in a humanoid robot. *IEEE Transactions on Systems, Man and Cybernetics, Part B. Special issue on robot learning by observation, demonstration and imitation*, 37(2): 286–298, 2007.
- Gribovskaya, E. and Billard, A. Learning nonlinear multi-variate motion dynamics for real-time position and orientation control of robotic manipulators. In *Proceedings of IEEE-RAS International Conference on Humanoid Robots*, 2009.
- Hersch, M., Guenter, F., Calinon, S., and Billard, A. Dynamical system modulation for robot learning via kinesthetic demonstrations. *IEEE Transactions on Robotics*, 2008.
- Lopez Infante, Marta and Kyrki, Ville. Usability of force-based controllers in physical human-robot interaction. In *Proceedings of the 6th international conference on Human-robot interaction, HRI '11*, pp. 355–362, 2011.
- Weiss, A., Igelsboeck, J., Calinon, S., Billard, A., and Tscheligi, M. Teaching a humanoid: A user study on learning by demonstration with hoap-3. In *Proceedings of the IEEE Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 147–152, 2009.