An Experimental Study for Identifying Features of Legible Manipulator Paths

Min Zhao[1,2], Rahul Shome[1], Isaac Yochelson[1], Kostas Bekris[1], Eileen Kowler[2]

[1] Department of Computer Science, [2] Department of Psychology - Rutgers University
kostas.bekris@rutgers.edu - `http://www.pracsyslab.org`

Abstract. This work performs an experimental study on the legibility of paths executed by a manipulation arm available on a *Baxter* robot. In this context, legibility is defined as the ability of people to effectively predict the target of the arm's motion. Paths that are legible can improve the collaboration of robots with humans since they allow people to intuitively understand the robot's intentions. Each experimental trial in this study reproduces manipulator motions to one of many targets in front of the robot. An appropriate experimental setup was developed in order to collect the responses of people in terms of the perceived robot's target during the execution of a trajectory by Baxter. The objective of the experimental setup was to minimize the cognitive load of the human subjects during the collection of data. The extensive experimental data provide insights into the features of motion that make certain paths more legible for humans than other paths. For instance, motions where the end-effector is oriented towards the intended target appear to be better in terms of legibility than alternatives.

Keywords: Human-Robot Interaction, Legible paths, Manipulation, Co-robots

1 Introduction

The increasing availability of low-cost, compliant and human-friendly manipulators allows robots, such as Rethink Robotics' *Baxter* [1], to be placed in close proximity to human workers. Unlike traditional automation systems, which needed to be kept in cages, these compliant robots can share a common workspace with human workers. A clear benefit of this close proximity is the opportunity for cooperation between a human worker and an assistive robot.

In order for a robot to be effective in an assistive role, it is important that the human is able to easily and quickly understand the robot's intentions by observing its actions. Ideally, this understanding will come in an intuitive manner, similar to how humans are innately able to communicate with one another non-verbally when working in close quarters. Legible motion plans are an important part of making the robot understandable by human co-workers intuitively. In this context, the legibility of a motion corresponds to whether human subjects can realize early on which is the actual target of the moving arm out of many possible choices.

The goal of this work is to identify the key features of robotic motion for manipulators that contribute to their legibility. The motions are executed by two seven degrees-of-freedom manipulation arms that are mounted on a *Baxter* robot. The arms move towards grasping multiple targets, which are positioned linearly in front of the robot. As the manipulator moves, human subjects observe the robot and report their belief regarding the intended target of the arm. An appropriate experimental setup was developed in order to collect these responses, so as to minimize the cognitive load of the human subjects and achieve good accuracy.

The five different types of trajectories that were considered during this experimental study cover a variety of discriminant legibility features. Some of the features correspond to arm policies, such as the shortest path in the configuration space, and other correspond to end-effector, i.e., "hand", policies, such as the orientation of the end-effector relative to the target. The experimental results show that the legibility of different trajectories is indeed different and consistent across different targets. Motions which allow the end-effector to point towards the intended target and move along a straight line in the workspace result in enhanced legibility.

The longer term objective of identifying these legibility characteristics is the design of motion planners that incorporate these features into the planning process so as to automatically generate legible motion. The hope is that co-robots, which can generate legible motion plans, can more effectively collaborate with humans.

1.1 Related Work

Previous work has emphasized the importance of anticipatory motion [2]. By identifying a symbol, or a socially representative element of the motion, and using it as early as possible during the motion, the robot's actions are easily and quickly interpreted by observers. It has also been indicated that legible, anticipatory motion greatly assists in collaborative tasks.

Research has also focused on exploiting the repeatability of common collaborative tasks to generate anthropomorphic motions [3]. There has been work on creating metrics that can reproduce motion plans to be more human-like [4]. Another philosophy in generating motion plans has been learning by demonstration. Motions, that are demonstrated by human teachers, are used to build the policy for the robot to map its state to an appropriate motion [5, 6]. This line of work leverages anthropomorphic motions. The legibility problem, however, does not necessarily correspond to the capability of a robot to reproduce human-like motion, but how a human perceives the robot's motion.

This crucial motivation has resulted in recent important efforts in identifying aspects of and generating legible robot motion [7, 8], which have inspired and influenced the current work. In particular, these efforts have resulted in a formalization of robot motion legibility, and approaches for autonomously generating legible robotic motion plans. Further work by the authors along this line has focused on distinguishing between predictability and legibility. In the corresponding experimental process the focus was on discriminating the legibility of motion using video recordings of a robot that can potentially reach two goals in an otherwise uncluttered workspace. Familiarization [9] has been shown to improve predictability when coupled with learning.

Human beings are good at interpreting actions and relative intentions of other moving agents in their environment. This ability is developed during the first fourteen months of a person's life [10]. During daily life, there are usually two action interpretation processes [11]:

1. *Action-to-Goal* inference, in which people try to predict the result of the action based on the information accumulated during the action's execution.
2. *Goal-to-Action* inference, in which people try to predict a type of action that could achieve a determined goal

The focus of legibility is on understanding action-to-goal inference, namely how humans interpret the observed actions and then discover the underlying intention [7]. Adults, young children, and even infants could selectively focus on the key components of the behavior of others, which is relative to their intention. In psychophysical experiments the human hand was discovered to play a crucial role during interpreting and sharing actions and intentions of people by others [12, 13]. Previous psychological studies show that between nine and twelve months, infants develop a perceptual link between pointing to the target object and the targeting mechanism itself. They understand that pointing is an object-oriented action [13]. These results motivate the focus of this study on features related to the robot's end-effector.

This work identifies the possible features of motion of a manipulator, which make it easier for its motions to be understandable by humans. The goal of the current work is to expand upon the existing experimental studies on this subject [7, 8] in two primary directions. This effort considers a workspace with many potential targets for the robot to interact with. Secondly, experiments are performed with human subjects placed in close proximity to the manipulator robot during trajectory execution and do not make use of recordings. These experiments confirm aspects of previous work, such as the contradictory nature of shortest and legible paths, and they also reveal important features of legible paths in cluttered scenes. For instance, the direction and workspace path of the end effector are shown to significantly influence a human observer's capability to realize the robot's intended target.

2 Generating Different Manipulator Paths

In order to evaluate key features of legible paths for manipulators, this experimental study considers a variety of path types for a dual arm manipulator, i.e., a Baxter robot by Rethink Robotics. There are four *arm policies* considered in this study:

1. shortest path in configuration space (i.e., minimizing change in joint angles),

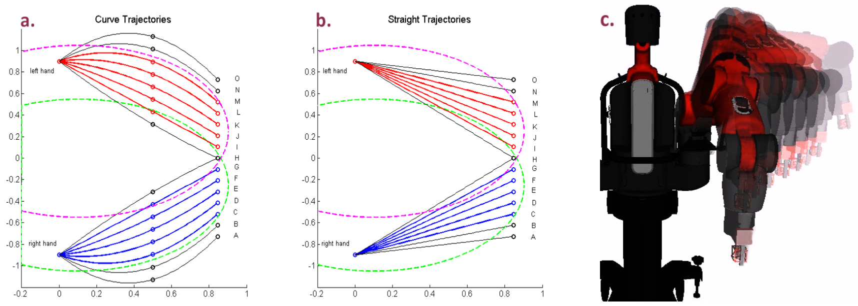

Fig. 1. *Left(a)* : "curve" and *Center(b)*: "straight" paths seen from above. The points on the left side of each plot represent the starting position for the left (red) and right (blue) end-effector. The lines show paths to reachable targets. Each hand has its own reachable region (green curve for right; purple curve for left hand). *Right(c)* one of the "overhead" paths in simulation [14]. The end-effector remains vertical and points downward.

2. overhead motion frequently appearing in "pick and place" paths,
3. shortest, straight-line path for the end-effector in workspace, and
4. "curved" path for the end-effector in the workspace to exaggerate intent (see Fig. 1a).

Likewise, this study considers two potential *hand policies*:

1. hand goes immediately to final joint position (e.g., overhead grasp) and stays there for the duration of the motion, and
2. hand points toward the goal in the workspace at all times. The pointing feature of these paths can be seen as a symbol generating anticipation of the motion [2].

Paths are generated both for the left and right arm of the robot. For each arm and for every type of trajectory, a fixed start position that is raised from *Baxter's* at-rest position is used. It helps in terms of target reachability. Different paths are generated for multiple targets. The targets are placed evenly along a line on a table in the manipulator's reachable workspace (see Fig. 1 (a and b)). By combining the above mentioned policies and pruning incompatible combinations, five different classes of path are considered in the experimental study:

1. "Shortest" path: This is the shortest path in the configuration space computed on an asymptotically near-optimal version [15] of a probabilistic roadmap [16] in the Open Motion Planning Library [17]. This class makes use of arm policy 1 (fig. 2:1) and immediately provides a path for the hand as well.
2. "Overhead Down" path: Similar to paths employed for pick-and-place tasks by Baxter robots in industrial settings, where the end-effector moves in a position over the target and points downwards throughout the motion (see Fig. 1c). This class makes use of arm policy 2 and hand policy 1 (fig. 2:2).

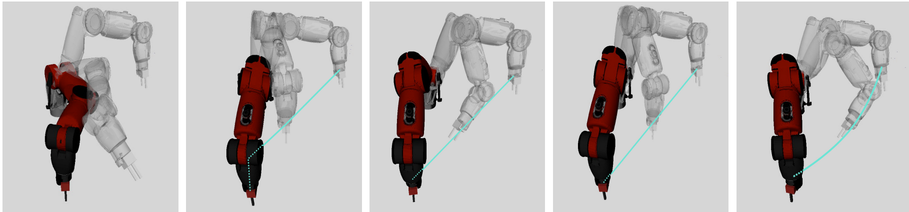

Fig. 2. *Left to right:* 1: Shortest C-space path, 2: Overhead down, 3: Straight pointing to target, 4: Straight down, 5: Exaggerated "curved" motion pointing to target

3. "Straight" path: The robot moves its end effector along a linear path from the initial position to the target object while the end effector points towards the target (see Fig. 1b). This class uses arm policy 3 and hand policy 2. (fig. 2:3)
4. "Straight Down" path: The robot moves its end effector along a linear path from the initial position to the target object while the end effector remains in a vertical orientation pointing down. This class makes use of arm policy 3 with hand policy 1. (fig. 2:4)
5. "Curved" path: The robot moves its end effector along an exaggerated curved path while pointing at the target. This class is inspired by ideas in previous work towards generating legible paths [7] (see Fig. 1a). This class makes use of arm policy 4 combined with hand policy 2. (fig. 2:5)

The above set of trajectories is designed to avoid confounding the effects of hand policies with the effects of arm policy, while keeping the total number of trajectories to a reasonable number so as to be able to extract useful conclusions. Note that there are two types of trajectories that are sharing the same arm policy (straight-line for the end effector in the workspace) but are different in terms of the employed hand policy. There are also two control classes, reflecting standard manipulation strategies ("shortest" and "overhead down" trajectories). In this way, the relative importance of these features can be discovered by comparing the time it takes for human subjects to realize the motion's target.

To ensure that for all classes there is ample time for subjects to give feedback about their belief of targets, all trajectories in this study are scaled to be performed in 8 seconds.

3 Setup for Collecting Human Responses

The experimental setup is designed to effectively record the responses of subjects' belief about the target of trajectories executed by the robot. A requirement was that both the targets and the robot were within the view of the subjects. The subjects also had a clear view of the entire motion of the robot manipulators. For studying legibility, the subject must be able to pay attention to the motion of the robot without distractions. Minimizing the cognitive load of the subject during the experiment involves minimizing distractions as well as making the

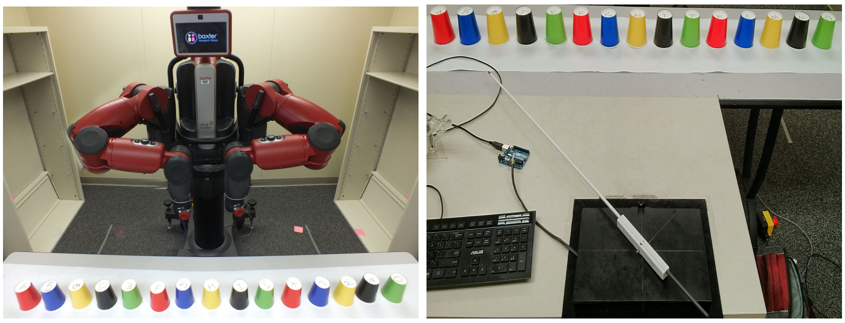

Fig. 3. (left) The start position of the trajectories on the Baxter robot during the experimental setup. (right) A view of the pointing device from the subject's perspective.

data recording interface intuitive and effortless. In order to achieve this, an efficient recording mechanism is desired, which is both accurate in recording the responses and easy to assemble. The recording interface should also be resilient enough to withstand repeated experimental trials.

As shown in Fig. 3, the experimental setup consists of a Baxter robot, a workstation, a table with 15 colored cups, and a pointing device. Only 10 of the cups can be reached by the robot from its starting position with all 5 trajectory classes. Five for the left arm and five for the right one. These 10 cups were designated as targets. During a trial, the robot follows one of the trajectory classes from its starting position to a target. A human sitting behind the pointing device uses the pointer to indicate his or her belief of the robot motion's intended target. The position of the pointing device is then recorded in a log together with the target and the class of the trajectory followed.

The trajectories are stored and played back during the trials in order to ensure that artifacts from the random sampling in the motion planning process do not cause discrepancies between trials of the same trajectory class. Moreover, the overhead of planning for the execution for the trajectories is avoided by generating the trajectories once and replaying them. For each of the workspace constrained paths, MathWorks' MATLAB [18] is used to perform linear interpolation among a series of points in the workspace. Then, the MoveIt! package [19] with a KDL kinematics solver [20] and an OMPL [17] implementation of a PRM* variant is used to plan trajectories between the interpolated points. The final trajectories can be played on the robot using the Baxter RSDK [1].

The pointing device is fixed to the spindle of a linear potentiometer. The edges of the resistive track are then connected to the 5 volt and ground pins of an Arduino device and the wiper to an analog input pin. An Arduino sketch then performs the necessary calculations to extrapolate from the wiper voltage the position along the line of targets at which the ray of the pointer will intersect. This distance is then forwarded to the Arduino's USB port.

For each human subject, three random permutations of the 50 recorded trajectories are generated using a python script. The trajectories of each permutation are then executed in order, recording a log of the trajectory filename and pointer positions, with time-stamps, captured from the Arduino during the playback of each trajectory. In this manner, it is possible to ensure even coverage among the classes and targets while minimizing the chance of subjects guessing the target through means other than visual perception of the robot's motion.

The playback of a given trajectory is preceded by an auditory alert, a bell sound for the left arm or a buzz sound for the right arm. These sounds alert the subject regarding which arm they should direct their attention toward. The robot then plays the trajectory, which has been scaled to run in 8 seconds, as the subject's feedback is recorded. Following the playback, the robot returns to a starting position, which is common to all the trajectories. Then, the subject is shown the number of trials that have been completed and is prompted to press any key on the keyboard to continue.

After each block of 50 trials, which forms a permutation of the full set of trajectories, the subject is given a mandatory two minute break. These breaks allow the subject to rest, and to maintain attention on the perception task. Human subjects are asked to participate in this experiment only once. This is to ensure that the base legibility of the paths when first encountered, and also the rate at which learning about the paths takes place, can be accurately gauged.

4 Experimental Evaluation of Legibility

4.1 Reaction Time

In order to compare the legibility of different types of trajectories in detail, the study examined the reaction time of 30 subjects, who had the opportunity to observe 150 trials (i.e., 3 blocks of 50 unique trajectories randomly ordered). Reaction time is the time it takes a subject to converge to the correct target. This is measured as the last occurrence when the pointing device enters the target zone during a trial. Figure 4 plots three types of reaction times:

1. Time for convergence to the target itself (Figure 4a), which happened at the late part of trials;
2. Time for convergence to the range within 1 cup away from the target (Figure 4b), which happened at the middle range of trials; and
3. Time for convergence to the range within 2 cups away from the target (Figure 4c), which happened at the beginning of trials.

Data from the three blocks of 50 trajectories provided to each human subject are presented in order from left to right. In general, the straight-pointing type (red bars) was always the best. The curve-pointing (blue bars) was the second best. The shortest type (gray bars) was the worst, especially regarding the early on convergence (2 cups away, or 1 cup away from the target). One-way ANOVA tests show that there were significant differences among the different types of trajectories for all groups (Table 1, F scores and p values).

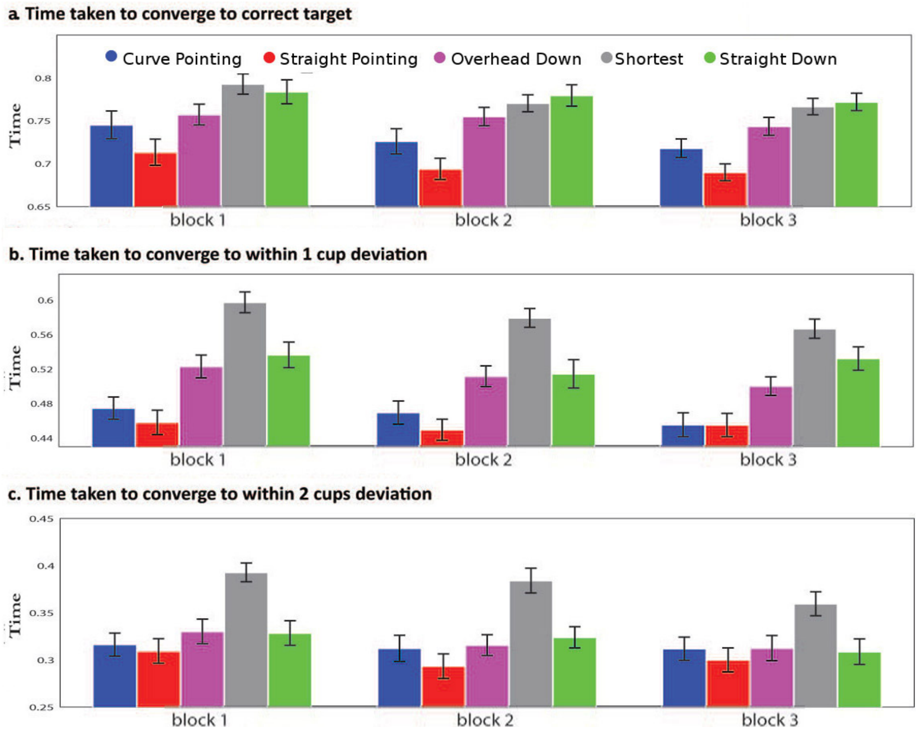

Fig. 4. Reaction time (RT): the time converging to (a) the correct target; (b) 1 cup away from the correct target; (c) 2 cups away from the correct target. There were five different types of trajectories: curve-pointing (blue), straight-pointing (red), overhead-down (pink), shortest (gray) and straight-down (green). There were three blocks of 50 trajectories in order: block 1 (left), block 2 (middle) and block 3 (right). The error bar represent -/+ 1 standard deviation error.

Block 1 (plots on the leftmost column of each figure) provides the response times when each trajectory was first presented to subjects. The performance for the shortest (grey bars) type was always the worst when converging to all types of error range. The disadvantage of the shortest type was obvious when converging into relatively large error ranges (2 cups or 1 cup away from the target). Pairwise comparison shows that it is significantly longer than the other four types (Table 1). This disadvantage decreased when approaching to the correct target finally. It suggests that the confusion of the shortest type usually appeared at the early stage of the trajectories.

The reaction time for the straight-pointing type (red bars) was consistently shorter than the others when converging to the range within 1 cup of the target and to the target, and it is marginally shorter than the others when converging to the range within 2 cups of the target. The curve-pointing was always longer than the straight-pointing but shorter than the remaining three. This implies the legibility of the straight-pointing path. The curve-pointing path is the second most legible among the five types. The ease of understanding of the straight-

	Block 1	**Block 2**	**Block 3**
target	5.2 (<.01) 2-4; 2-5	8.34 (<.001) 1-5; 2-3; 2-4; 2-5	11.37 (<.001) 1-4; 1-5; 2-3; 2-4; 2-5
1 cup	16.74 (<.001) 1-4; 1-5; 2-3; 2-4; 2-5; 3-4; 5-4	14.38 (<.001) 1-4; 2-3; 2-4; 2-5; 3-4; 5-4	14.91 (<.001) 1-4; 1-5; 2-4; 2-5; 3-4
2 cups	7.3 (<.001) 1-4; 2-4; 3-4; 5-4	7.61 (<.001) 1-4; 2-4; 3-4; 5-4	3.3 (=.013) 2-4; 5-4

Table 1. One-Way ANOVA analysis for RT of 2 cups away, 1 cup away and pointing to the target for each block. In each cell, the values in the first row are the F-score (p-value). The second row lists all pairwise types which are significantly different from each other from post-hoc test (1- Curve-pointing; 2- Straight-pointing; 3- Overhead-down; 4- Shortest; 5- Straight-down).

pointing and the curve-pointing trajectories could be due to the fact that the end-effector was always pointing to the target. The end effector (i.e., hand) was previously reported as an important cue in understanding people's intentions [12, 13, 10]. The advantages of the end-effector pointing to the target were strongest when converging to the range of 1 cup away from the target. It suggests that the characteristics of curve-pointing and straight-pointing helps people understand the intention of the robot by converging to the smaller error range more quickly.

The "curve-pointing" did not perform as well as "straight-pointing", which was surprising given the conclusions of previous studies [7]. Nevertheless, this could be due to the difference between the "two-targets" setting in previous studies and the "crowded-target" setting in the current experiment. With multiple targets in a crowded environment, the curve path was more likely to confuse people, rather than improve legibility.

Overhead-down was the third most legible trajectory and it was better than the straight-down. This also makes sense, because whenever the overhead-down trajectories reached to the top of the cup, subjects knew the answer for sure. The straight-down trajectories were still on the way to the top of the cup at the same time point. This leads to similar reaction times between these two types of trajectories when converging to 2 cups or 1 cup away from the target. It leads into shorter reaction time for the overhead-down trajectory when converging to the correct target.

Figure 5 provides a study of the effects of learning over the duration of the trajectory. It shows that the time converging to the target was decreasing across blocks, which means subjects did learn trajectories. The learning effect is larger from block 1 to block 2, than from block 2 to block 3. This could be because subjects were already well trained when entering into block 2 and might get tired in block 3. The learning effect also varies among different types of trajectories. The shortest type shows greater learning effect than the others in all convergence cases (Figure 5 a,b,c). These results suggest that the shortest type was the hardest one to be interpreted at the early stage, but it can be learned by more training. Nevertheless, the learning does not allow it to reach the legibility level of alternatives such as the "straight-pointing" path. Additionally, the learning effect appears also as decreasing variance in later blocks (block 2 and 3).

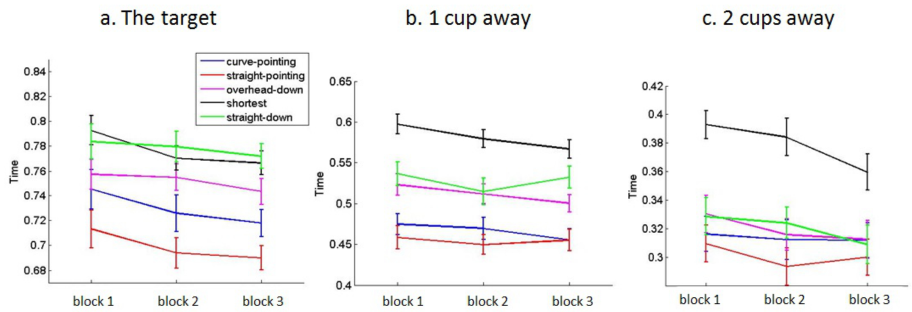

Fig. 5. Learning effects reflected in reaction time. The time converging to (a) the correct target; (b) within 1 cup deviation from the true target; (c) within 2 cups of the true target. There were five different types of trajectories: curve-pointing (blue), straight-pointing (red), overhead-down (pink), shortest (black) and straight-down (green). The error bars represent -/+ 1 standard deviation error.

4.2 Predicted Target over time

The root mean square of the distance between the pointed cup and the true target reflects the accuracy of the subject's prediction of the target over the course of the trial. It is averaged across trials for each subject, and then across all subjects. Figure 6 shows that the root mean square varied for different types of trajectories at the beginning and converged to the correct target location in the end over a normalized time scale. The convergence was fast during the middle range of trials (0.3-0.7) for all trajectories. A consistent result is that the predictions for the shortest trajectory were further away from the true target than any other type of trajectories.

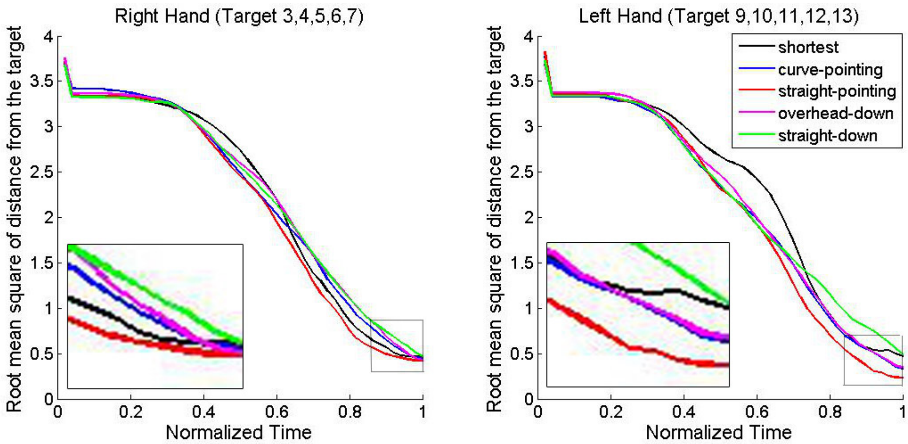

Fig. 6. Root mean square of distance from the target along the normalized time scale for five types of trajectories: shortest (black), curve-pointing (blue), straight-pointing (red), overhead-down (pink) and straight-down (green).

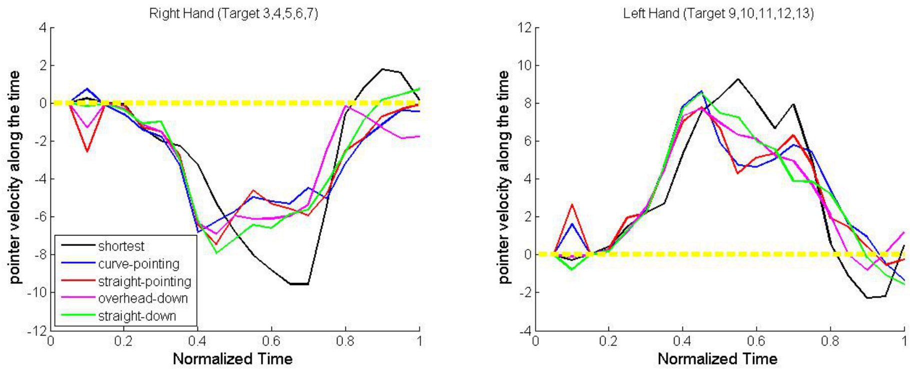

Fig. 7. Pointer velocity along the normalized time scale for five types of trajectories: shortest (black), curve-pointing (blue), straight-pointing (red), overhead-down (pink) and straight-down (green).

The pointer velocities (Figure 7), which was how fast subjects moved the pointer, were peaking during the middle range of trials (.4-.6). The graph shows that the shortest trajectory type resulted in different behavior than the other four. These results suggest that the subjects might not have been able to predict the target during the early parts of the shortest trajectories as well as they were able to do so for the other types. Frequently during the shortest paths, the end-effector was overshooting the target and then returning back to it, which complicated the interpretations of the motion even close to the completion of the path.

4.3 Performance by cup

The understanding of different types of trajectories was also related to the location of the target. In order to better analyze different trajectories, it is useful to further examine performance (mean position and pointer velocity over time) for different cups. Figure 8 shows the mean distance from the target for each cup and Figure 9 shows the pointer velocity for each cup.

By comparison to the cups located on the edge of the target set (i.e., 3, 4, 12, 13), for the cups located near the center (i.e., 6, 7, 9, 10), the mean deviation from the target was more likely to cross the 0 level (Figure 8). As the subjects typically begin with the pointing device centered, this suggests that subjects were more likely to overshoot the target. The starting position of the robot is nearer the edge cups than the center cups. The overshooting could indicate that they are following the arm rather than predicting the target accurately. Subjects were more likely to overshoot the target for the shortest type. The traces in Fig. 8 illustrate the subjects' reactions. The trajectories with the lowest reaction times also demonstrate the least overshooting.

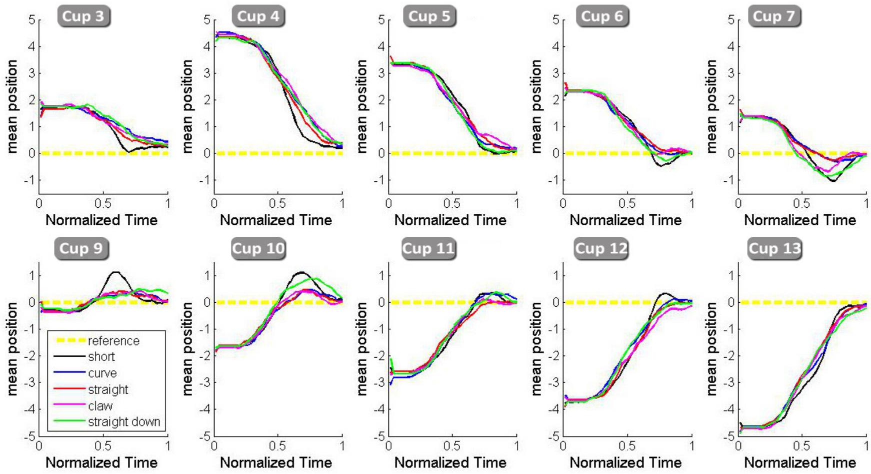

Fig. 8. Mean distance from target along normalized time scale for each cup (cup No. labeled on the top of each plot). Five types of trajectories: shortest (black), curve-pointing (blue), straight-pointing (red), overhead-down (pink) and straight-down (green).

4.4 Conclusions

Overall, the straight-pointing path was the easiest to understand, followed by the curve-pointing, and the shortest type was the least legible. The advantage of the straight-pointing path and the curve-pointing is likely due to the end-effector's orientation towards the target. Reaction times are slower with the straight-down than the straight-pointing, though both of them followed the same end-effector paths in the workspace.

The disadvantage of the shortest type appears as slow convergence; it takes more time to approach a certain error range than the other types. The overhead-down was better than straight-down when converging to the target. An obfuscating effect is that subjects could know the target for sure when the hand moved to the top of the cup, while at the same time point, the straight-down motion was still away from the target.

Even though the learning effect exists, it varies for different types of trajectories. The general learning benefit appears with less variance in block 2 and 3. Three types - shortest, curve-pointing and straight-pointing, could be learned across blocks, while the overhead-down and the straight-down show least learning. Even though the shortest type can be learned, it still doesn't become more legible than the other four trajectories. This implies, that it optimizing a C-space metric may not be the best strategy for providing legible trajectories.

5 Discussion

This study partially supports previous findings regarding the legibility of robot motion [7], i.e., different types of paths can have highly divergent levels legibility.

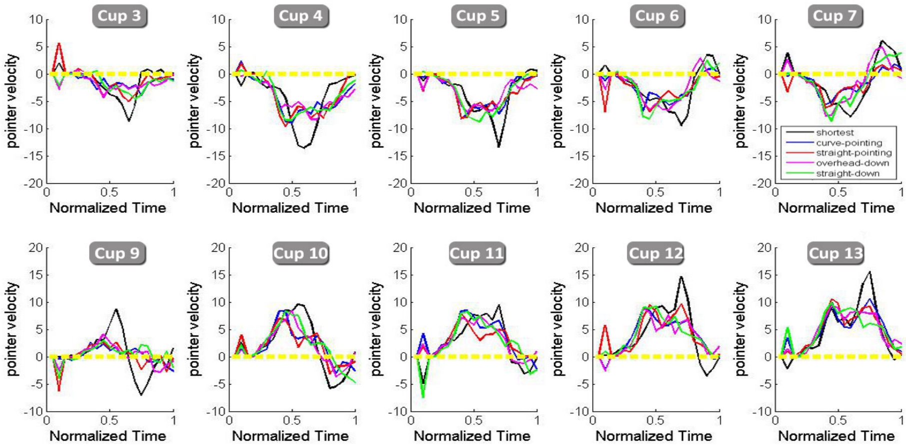

Fig. 9. Pointer velocity along the normalized time scale for each cup (cup No. labeled on the top of each plot). Five types of trajectories: shortest (black), curve-pointing (blue), straight-pointing (red), overhead-down (pink) and straight-down (green).

Shortest C-space paths, which are frequently the focus of motion planning methods, can be poor choices in terms of legibility. Similarly, paths that are currently used for pick-and-place tasks in industrial setups (e.g., "overhead") also appear not to be intuitively interpreted. Paths that focus on the orientation of the end-effector seem to be advantageous in terms of legibility, since they exhibited the best performance in estimation (high accuracy and less convergence time).

The fact that "straight-pointing" paths were more legible relative to "curve-pointing" paths corresponds to a difference from previous findings. The idea behind the "curve-pointing" paths is that legibility may increase by exaggerating the arms' motion so that it moves away from unintended targets (Fig. 1a). The difference seems to be due to the presence of multiple targets in the current work. When only two targets are present, exaggerating the motion in one direction can significantly assist in identifying the target but can be confusing in the case of multiple targets, or in cluttered workspaces. This study is intended to inform legibility in such cluttered environments.

A significant observation corresponded to the importance of the end-effector's orientation relative to the target. It was hypothesized that humans might pay particular attention to the pose and orientation of a robotic end-effector, similar to the way they respond to a human's hands. The experimental results confirmed this hypothesis. It would be worthwhile to incorporate the maintenance of such end effector orientations into the cost functions of motion planners in the future.

A question that needs to be answered is whether it is worthwhile to consider legibility of robotic motion planning paths, as opposed to relying on learning to take place in human observers. There is a learning effect when the subjects observed the same trajectories. The benefit primarily appears as reduced variance during repetitions of the same trajectories and varies across types of trajectories. Three types of paths, shortest, curve-pointing and straight-pointing, could be learned across blocks, while the overhead-down and the straight-down did not exhibit significant learning behavior. Although the learning effect existed, the benefit of the learning might not be able to override the advantage of the legible information, which was supported by the fact that the performance of the shortest type was improved in later blocks, but it was still not as good as the performance of other types.

Note that in the experiments the subjects witnessed the same path to a target in 3 blocks times rather than 3 variations of the same type of path to the same target. Certain planners, such as sampling-based ones, can vary in the repeatability of their solutions. It is not necessarily the case that similar degrees of learning would occur for the general case of repeated exposure to motions plans generated from such motion planners. Furthermore, in this study the initial condition was always the same. When a robot needs to plan on the fly and transition from one task to another, the human subject will not be exposed to the same exact trajectories repeatedly. It is interesting to consider the effects of legibility in the context of trajectories that have different initial conditions.

During preliminary experiments, there was a transition from a web-based UI in the pilot trials, to the physical pointer feedback device used in gathering the data included here. This change decreased the cognitive load placed on human subjects by the data collection interface and resulted in a reduction between the best-performing and worst-performing path classes relative to the pilot study. A human co-worker in a collaborative setting is likely to have additional mental demands beyond the robot interaction. While minimizing the cognitive load might clarify the effects of legible features, such distractions might exaggerate the legibility of robot motions. An interesting line of future research is to analyze the effect of cognitive load on legibility.

Initial pilot trials also used trajectories which varied in duration. Increasing duration of trajectory execution gives the subject more time to recognize the legible features of the motion. However, it is not clear whether the effect persists if the trajectory duration keeps on increasing. Unnaturally slow trajectories might obfuscate the features that contribute to legibility. A scope for future work would be to understand the effects of the duration and speed of trajectories on legibility consistent among different types of trajectory.

Future experiments could involve the random placement of targets over a two-dimensional subspace, the presence of obstacles, as well as stopping the motion of the arm half-way towards the target and asking the user to guess the intended target. The longer-term objective is the definition of appropriate motion planners that generate highly-legible paths. It appears that such planners and accompanying cost metrics need to be reasoning for the orientation of the end-effector and its workspace path. This line of work can eventually lead to robots that use time-efficient paths when they operate in a dark factory floor and automatically switch to humanly-legible but less efficient paths when people enter their workspace and collaborate.

References

1. Rethink Robotics: Baxter Research Robot Software Development Kit (SDK) Version 0.7.0 (2013)
2. Gielniak, M., Thomaz, A.: Generating Anticipation in Robot Motion. In: RO-MAN, IEEE (July 2011) 449–454
3. Beetz, M., Stulp, F., Esden-Tempski, P., Fedrizzi, A., Klank, U., Kresse, I., Maldonado, A., Ruiz, F.: Generality and legibility in mobile manipulation. Autonomous Robots **28**(1) (2010) 21–44
4. Gielniak, M.J.: Spatiotemporal correspondence as a metric for human-like robot motion. In: In ACM/IEEE HRI. (2011)
5. Argall, B.D., Chernova, S., Veloso, M., Browning, B.: A survey of robot learning from demonstration. Robotics and Autonomous Systems **57**(5) (2009) 469 – 483
6. Billard, A., Calinon, S., Dillmann, R., Schaal, S.: Robot programming by demonstration. In Siciliano, B., Khatib, O., eds.: Springer Handbook of Robotics. Springer Berlin Heidelberg (2008) 1371–1394
7. Dragan, A.D., Lee, K.T., Srinivasa, S.S.: Legibility and Predictability of Robot Motion. In: International Conference on Human-Robot Interaction (HRI). (2013)
8. Dragan, A., Srinivasa, S.S.: Generating Legible Motion. In: Proceedings of Robotics: Science and Systems, Berlin, Germany (June 2013)
9. Dragan, A., Srinivasa, S.S.: Familiarization to Robot Motion. International Conference on Human-Robot Interaction (HRI) (March 2014)
10. Momasello, M., Carpenter, M., Call, J., Behne, T., Moll, H.: Understanding and sharing intentions: the origins of cultural cognition. In: Behavioral and brain science. (2005)
11. Csibra, G., Gergely, G.: Obsessed with goals: Functions and mechanisims of teleological interpretation of action in humans. Acta Psychologica **1**(124) (2007) 60–78
12. Woodward, A.L.: Infants selectively encode the goal object of an actor's reach. In: Cognition. (1998)
13. Woodward, A.L., Guajardo, J.J.: Infants' understanding of the point gesture as an object-directed action. In: Cognitive Development. (2002)
14. Kimmel, A., Dobson, A., Littlefield, Z., Krontiris, A., Marble, J., Bekris, K.E.: Pracsys: An extensible architecture for composing motion controllers and planners. In: Simulation, Modeling and Programming for Autonomous Robots (SIMPAR), Tsukuba, Japan (2012)
15. Marble, J., Bekris, K.E.: Asymptotically Near-Optimal is Good Enough for Motion Planning. In: Proc. of the 15th International Symposium on Robotics Research (ISRR-11), Flagstaff, AZ (28. Aug. - 1 Sep 2011)
16. Kavraki, L.E., Svestka, P., Latombe, J.C., Overmars, M.: Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces. IEEE TRA **12**(4) (1996) 566–580
17. Şucan, I.A., Moll, M., Kavraki, L.E.: The Open Motion Planning Library. IEEE Robotics & Automation Magazine **19**(4) (December 2012) 72–82 `http://ompl.kavrakilab.org`.
18. MATLAB: version 8.2.0.701 (R2013b). The MathWorks Inc., Natick, MA (2013)
19. Sucan, I.A., Chitta, S.: Moveit! Online
20. Smits, R.: KDL: Kinematics and Dynamics Library. `http://www.orocos.org/kdl`