Evaluating Intent-Expressive Robot Arm Motion

Christopher Bodden, Daniel Rakita, Bilge Mutlu, and Michael Gleicher
Department of Computer Sciences, University of Wisconsin–Madison
1210 West Dayton Street, Madison, Wisconsin 53706 USA
{cbodden, rakita, bilge, gleicher}@cs.wisc.edu

***Abstract*— Planning effective arm motions is integral to manipulation tasks. In general, motion synthesis methods have focused on functional objectives, such as minimizing time and maximizing efficiency. However, recent work in human-robot collaboration suggests that choices in motion design can influence collaboration performance and quality. Some motion designs are easier than others for human observers to interpret. In this paper, we explore the tradeoffs in robot arm movements designed to be observed by people. Through a series of human-subjects experiments, we compare collaboration performance between several motion-synthesis methods explored by prior work. We find that a number of factors, including the design of the robot arm and metric for success, affect the relative merits of different approaches.**

Fig. 1. This work considers the effects of different methods for planning robot arm motion on collaborative task performance. We compare straight-line paths (left), minimal joint-distance paths (middle), and a state-of-the-art "legible" motion type (right) for several robots with different kinematics. The white line shows the end-effector path.

I. Introduction

Creating the motions of a robot arm is an important part of planning for manipulation tasks. These motions are typically synthesized to achieve functional goals, such as being as fast as possible or providing sufficient clearance around obstacles. However, in scenarios that involve collaboration between humans and robots, robot arm motions also communicate to human observers information regarding the state, intent, or capability of the robot. For example, a robot reaching for an object not only performs the task, but its reach toward the object also communicates the robot's intent before it reaches the object. Recent work has shown that different movement-synthesis strategies communicate the robot's intent with different levels of clarity [1] and that movements can in fact be designed to improve the effectiveness of this communication [2].

In this paper, we compare the intent-expressive capability of three existing methods for synthesizing robot arm motion in a series of human-subjects experiments. Our study considers motion paths that minimize path length in joint-angle space (*simple*), paths that minimize end-effector path length to approximate straight lines (*straight*), and an optimization objective specifically designed to be easy for observers to interpret (*legible*). We compare collaboration performance resulting from these movement approaches with multiple robot-arm designs (different kinematics) and multiple measures of success. We discuss the many factors that shape the effectiveness of the different path types in communicating the robot's intent in order to inform future design of and research into robot arm motions. Our findings indicate that many factors influence the relative performance of the different path types and that there is not a single "best" approach.

II. Related Work

Two common strategies for creating arm motions are to minimize the change in joint angles (i.e., to linearly interpolate configurations in joint angle space) or to move the end-effector in as straight a line as possible (i.e., to minimize the distance traveled by the end-effector). Bortot et al. [1] observe that people view these motions differently. Dragan and Srinivasa [2] discuss how different motions can be specifically designed to be "legible," that is, to better convey their intent to human observers. In subsequent papers, they provide an optimization approach to synthesize these motions [3] and show that these motions can improve human-robot team performance in collaborative tasks [4].

Given the promise that motion design has for human-robot collaboration, a deeper understanding of the design space for robot arm motion is needed. Zhao et al. [5] provide the most comprehensive study of collaboration performance resulting from a variety of movement types, including "simple," "straight," and "curved". However, this study only approximates the legibility optimization using "curved" motions. Similarly, "straight" paths are not considered in the comparisons conducted by Dragan et al. [2], [4]. Furthermore, all studies to date consider only a single robot design and do not discuss how different success criteria might change the effectiveness of motions. For example, Zhao et al. [5] and Dragan et al. [4] have both considered absolute duration as a metric to measure when participants correctly predicted a target location. Dragan et al. [2] have also considered a relative scoring metric for correctness versus path duration. In this paper, we seek to gain a deeper understanding of the effects of different motion designs by comparing them with different robots and different success metrics.

III. Comparing Trajectories

In this work, we wish to analyze how effectively different motion synthesis methods are at expressing intent. First, we discuss our method for synthesizing robot arm trajectories. Second, we discuss two cases of minimal-energy planning that are particularly interesting when it comes to intent-expressive motions. Finally, we discuss a method for synthesizing motion designed specifically to express intent.

A. Synthesizing Trajectories

To generate robot arm trajectories, we employ a flexible, optimization based technique known in animation literature as *spacetime constraints* [6]. Our technique allows us to produce trajectories that minimize an objective function subject to a set of constraints. The constraints allow us to set motion requirements, such as a target position. The objective allows us to set the properties of the motion, such as minimal-energy. Both the objective and constraints may consider the robot's joint positions and/or state configurations.

We define a robot's trajectory as a function that maps time to state configurations, $\Theta : \mathbb{R} \rightarrow \mathbb{R}^n$ where n is the number of degrees of freedom, so that the configuration along a trajectory at time t is $\Theta(t)$. We denote the kinematic function that maps from configuration space to a position, the forward kinematics, as FK, such that the end-effector position at time t is $FK(\Theta(t))$.

Optimizing the variational problem results in a trajectory:

$$\Theta^* = \operatorname{argmin} g(\Theta) \text{ subject to } c_i \diamond k_i, \tag{1}$$

for the motion range t_0 to t_f. The objective function g is a function over the *trajectory* that returns a scalar value. Each constraint, c_i, is either $\diamond \in \{=, \leq, \geq\}$ to a constant k_i.

Constraints allow us to define the requirements for the motion. In this paper, we consider motions with an initial pose constraint ($\Theta(t_0) = k_0$), a positional constraint on the end-effector at the end of the motion ($FK(\Theta(t_f)) = e_f$), and a constraint that the grasped object is vertical at the end of the motion $\hat{FK}(\Theta(t_f)) \cdot [0, 1, 0] = 1$, where $\hat{FK}(\Theta(t))$ denotes the vertical vector for the end-effector orientation.

To solve the variational optimization problem, we discretize the trajectory and approximate the objective function with finite differences sampled along time. This discretization produces a non-linear programming problem over the variables of the representation of the trajectory that can be modeled using automatic differentiation and solved using commonly available variants of Sequential Quadratic Programming (SQP) as described by Witkin and Kass [6] and Gleicher [7]. We use freely available tools from standard libraries (the Python `ad` automatic differentiation package and the SLSQP solver from `scipy`). We note that for the class of spacetime objectives we consider, we cannot use per-frame iterative approaches [8]. To execute the trajectory, we send the discretized result to the robot as a series of waypoints. While these waypoints are linearly interpolated, quadratic smoothing in the robot controller and the close spacing of the waypoints leads to robot movements that appear continuous. The robot controller executes the trajectory as fast as possible within joint velocity limits of the robot.

B. Objective Functions

Motion planning in robotics typically involves minimizing movement in the joint configuration space of the robot. This functionally simple motion is easy to plan and minimizes the joint-space energy spent by the robot. A simple objective function to achieve these types of trajectories using our technique is to minimize the length of the trajectory in configuration space (minimal joint movement):

$$\text{Simple} = \int_{t_0}^{t_f} \|\Theta(t)'\|^2 dt. \tag{2}$$

However, to signal intent using the trajectory of the robot, we must consider the human's cognitive system and not just the most functional motion for the robot. Research in neuroscience suggests that the trajectory of the motion is quite important for understanding intent [9]–[11], and that we tend to move our own manipulators in straight lines [12], [13]. This evidence gives rise to an interesting question: if we move our hands in straight lines, will straight-line motion in robots help us better understand their intent? Using our motion synthesizer, the following objective produces straight-line end-effector paths by minimizing the movement of the end-effector (as straight as possible subject to constraints):

$$\text{Straight} = \int_{t_0}^{t_f} \|FK(\Theta(t))'\|^2 dt. \tag{3}$$

In practice, we add a regularization term to the Straight objective, $\|\Theta'(t)\|^2$, to prevent joints from rapidly changing.

Other objective-based strategies, such as those proposed by Dragan and Srinivasa [3], can be used in this framework. We implement their proposed objective in this work and add their "trust region" by regularizing the total joint movement:

$$\text{Legible} = \frac{-\int P(G_r \mid \xi_{S\rightarrow Q}) f(t) dt}{\int f(t)\, dt} \tag{4}$$

where G_r is the goal; $\xi_{S\rightarrow Q}$ is the trajectory from S to Q; $P\left(G_r \mid \xi_{S\rightarrow Q}\right)$ is the probability of inferring G_r given the trajectory; and $f(t)$ is a linearly decreasing weight over time.

In the next section, we evaluate whether straight-line motions will express the robot's intent better than functional robot motion. We also consider how straight-line and functional motions compare to state-of-the-art legible motion. We generate many trajectories for these three motion types. Table I shows the average time to synthesize motions using our framework. The legible motions are those using the objective proposed by Dragan and Srinivasa [2], [3]. These motions take longer to plan because each trajectory requires calculation on every goal in a large discrete goal set.

IV. Evaluation

To compare the intent expressiveness of *straight* motion, *simple* motion, and *legible* motion as proposed by Dragan et al. [2], [3], we developed a simulated collaborative task

TABLE I

SYNTHESIS TIMES FOR MOTIONS

Robot	Motion Type	Mean (seconds)	Std. Dev. (seconds)
Mico	*Simple*	4.81	0.33
	Straight	5.37	0.16
	Legible	71.98	15.68
Reactor	*Simple*	3.15	0.21
	Straight	2.13	0.10
	Legible	40.02	5.41
UR5	*Simple*	5.10	0.25
	Straight	5.68	0.82
	Legible	57.35	6.69

(Figure 2) and conducted online human-subjects experiments using the Amazon Mechanical Turk marketplace. In these studies, we considered three objectives:

1) *Simple* (using Equation 2), called "shortest" by Zhao et al. [5] and "predictable" by Dragan et al. [2];
2) *Straight* (using Equation 3), called "straight" by Zhao et al. [5] and not considered by Dragan et al.
3) *Legible* (using Equation 4), approximated as "curved" by Zhao et al. [5] and called "legible" by Dragan et al. [2], [3].

A. Task Design

We designed a collaborative task in which the participant's performance was directly affected by their ability to interpret a simulated robot's intent. During the task, we measured how the participant's inference of the goal position changed over time. We motivated the task with a story about performing a chemistry experiment. The robot's task was to place cylinders at locations along a linear track, and the participant's task was to place a plate (controlled by his/her mouse) where he/she thought the robot was going to place the cylinder.

Participants were instructed to always use their best estimate of the robot's target as soon as they were able to make an estimate and not wait until they were sure of the goal. To encourage this behavior, they were told that the controls were unreliable and might stop working in the middle of the trial, leaving the plate at their last best guess (some unscored "contact loss" trials were randomly interspersed to remind participants of the need to update their guesses continually). To ensure that our measures started from a consistent location, participants were required to move the plate to the midpoint of the track prior to starting each trial.

To sample goals in an off-line and reproducible way, we pre-computed robot motions for a dense sampling of goal positions. This sampling approach was also necessary for producing legible trajectories, as the formulation proposed by Dragan et al. [3] is suited best for discrete sets of targets. To ensure consistency of measurement across participants and to avoid potential bias in sampling toward one side or based on distance, we divided goals into six buckets (three buckets on either side of the center at close, medium, and far distances). Each participant observed randomly sampled motions from each of the six buckets to ensure that the entire workspace was tested for each participant.

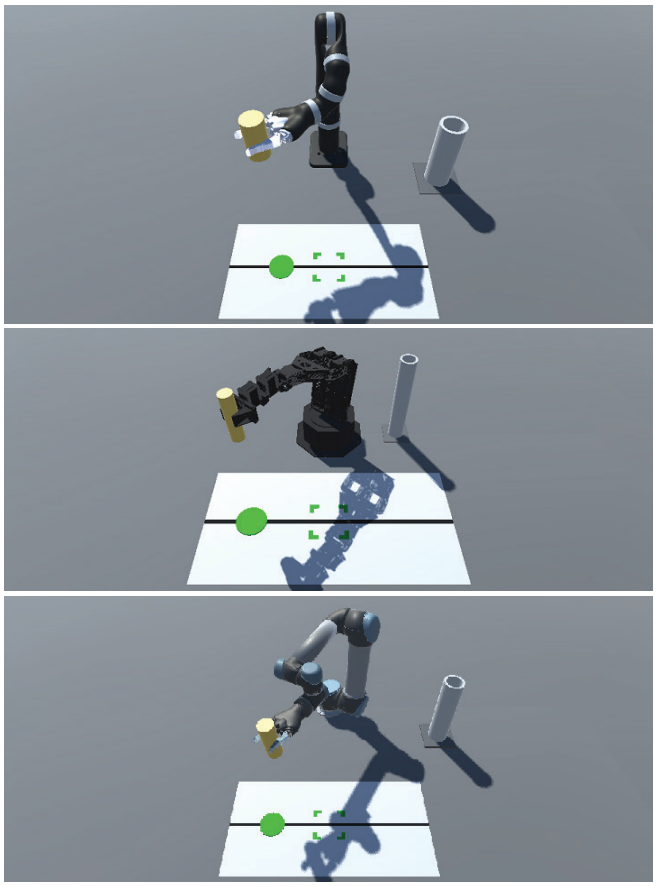

Fig. 2. The collaborative task for the simulated Mico (top), simulated Reactor (middle), and simulated UR5 (bottom). Participants were asked to place the green plate where they thought the robot arm would set the object.

We conducted the experiment using three different simulated arms to understand whether the kinematics of the robot affected the intent-expressiveness of the motions. We chose a six-DOF Kinova Robotics Mico, a six-DOF Universal Robots UR5, and a five-DOF Trossen Robotics PhantomX Reactor. The kinematics of each of these robots are very different, providing a good sample from which to test. The task setup for the study and each robot is shown in Figure 2.

B. Hypotheses

At a high level, we hypothesized that both *straight* and *legible* motion would have positive effects on collaborative task performance. We developed the following hypotheses based on *legible* motion being designed to help participants infer the goal [2] and *straight* motion most closely matching human arm motion and expectations [12], [13].

1) *Straight vs. Simple (H1)*: *Straight* motion will positively affect collaborative task performance when compared to *simple* motion. This hypothesis stems from previous work in neuroscience suggesting *straight* motions most closely match human arm motion [12], [13].
2) *Legible vs. Simple (H2)*: *Legible* motion will positively affect collaborative task performance when compared to *simple* motion. This hypothesis is a result of previous work by Dragan et al. [2], [4].

While we did not make specific hypotheses regarding the relative effects of *straight* and *legible* motion, our analysis also included an exploratory comparison of these motions.

C. Measures

To capture participant performance, particularly how quickly and accurately the robot's motion allowed participants to infer the robot's goal, we considered four objective measures described below.

- *Score* measures the remaining path percentage when the participant predicts the correct goal. This metric is important for providing a measure that is independent of path duration.
- *Time correct* measures the time (in seconds) it took the participant to predict the correct goal, which provides an important measure of absolute time.
- *Total error* measures the summation of the participant's prediction error over the entire trajectory.
- Finally, *idle time* measures how quickly the participant starts to move the plate (makes his/her first inference).

D. Experimental Design & Analysis

To test our hypotheses, we conducted three separate 3×1 between-subjects experiments (one for each of the three robots: Reactor, Mico, and UR5). In each experiment, the manipulated variable was motion type (*legible*, *straight*, *simple*). Each participant observed a total of 24 simulated motions (four random targets within each of the six buckets) for one of the three motion types.

None of the data we collected for any of the measures or robots met the normality assumptions for analysis of variance (ANOVA) after testing the goodness-of-fit for each measure using the Shapiro-Wilk W test ($p < 0.001$ for all measures and all three robots). Therefore, we analyzed our data using the Kruskal-Wallis test by ranks for motion type. Pairwise significance tests were performed using Wilcoxon rank-sum tests for each pair of motion types. Because we used non-parametric tests, we also include the median ("Mdn") value for measures in our results.

E. Participants

For each experiment, we recruited 36 participants (12 per motion type). For the Reactor, 11 females and 25 males were recruited with an average participant age of 32.3 (SD = 9.37, Max = 67, Min = 20). Participants reported moderate robot familiarity (M = 3.50, SD = 1.46) on a seven-point rating scale. Fifteen participants reported prior robot research experience. For the Mico, 20 females and 16 males were recruited with an average participant age of 33.3 (SD = 8.47, Max = 54, Min = 22). Participants reported moderate robot familiarity (M = 3.47, SD = 1.59), and nine participants reported prior robot research experience. Finally, for the UR5, 19 females and 17 males were recruited with an average participant age of 36.4 (SD = 9.67, Max = 57, Min = 21). Participants reported moderate robot familiarity (M = 3.17, SD = 1.78), and eight participants reported prior robot research experience. Each participant was paid $2.00.

TABLE II
DESCRIPTIVE STATISTICS

Robot	Measure	Simple			Straight			Legible		
		M	*Mdn*	*SD*	*M*	*Mdn*	*SD*	*M*	*Mdn*	*SD*
Reactor	*Score*	0.47	0.44	0.20	0.55	0.57	0.19	0.53	0.57	0.20
	Time Correct	1.73	1.79	0.62	1.83	1.73	0.75	1.60	1.42	0.66
	Total Error	0.27	0.26	0.12	0.29	0.28	0.14	0.26	0.24	0.15
	Idle Time	0.78	0.72	0.35	0.90	0.85	0.37	0.76	0.70	0.35
Mico	*Score*	0.28	0.20	0.22	0.44	0.45	0.22	0.41	0.42	0.23
	Time Correct	2.25	2.49	0.69	1.84	1.79	0.74	1.88	1.83	0.73
	Total Error	0.33	0.28	0.15	0.27	0.25	0.14	0.28	0.27	0.15
	Idle Time	0.96	0.90	0.50	0.88	0.75	0.43	0.92	0.85	0.52
UR5	*Score*	0.40	0.41	0.23	0.42	0.39	0.24	0.35	0.36	0.24
	Time Correct	1.61	1.58	0.61	1.99	2.13	0.83	1.80	1.80	0.66
	Total Error	0.25	0.24	0.13	0.31	0.26	0.16	0.26	0.25	0.10
	Idle Time	0.87	0.78	0.40	1.03	0.86	0.56	0.90	0.81	0.39

TABLE III
EFFECTS OF MOTION TYPE

Robot	Measure	H(2)	p
Reactor	*Score*	25.1	<0.001
	Time Correct	17.2	<0.001
	Total Error	8.06	0.018
	Idle Time	35.7	<0.001
Mico	*Score*	63.0	<0.001
	Time Correct	43.6	<0.001
	Total Error	16.2	<0.001
	Idle Time	5.93	0.052
UR5	*Score*	11.5	0.003
	Time Correct	29.8	<0.001
	Total Error	12.8	0.002
	Idle Time	8.17	0.017

TABLE IV
PAIRWISE COMPARISONS FOR EFFECTS OF MOTION TYPE

Robot	Measure	Level Difference p-value		
		Straight - Simple	*Straight - Legible*	*Simple - Legible*
Reactor	*Score*	<0.001	0.59	<0.001
	Time Correct	0.29	<0.001	<0.001
	Total Error	0.08	0.006	0.24
	Idle Time	<0.001	<0.001	0.50
Mico	*Score*	<0.001	0.21	<0.001
	Time Correct	<0.001	0.57	<0.001
	Total Error	<0.001	0.54	0.003
	Idle Time	0.01	0.19	0.31
UR5	*Score*	0.29	0.001	0.02
	Time Correct	<0.001	0.01	0.002
	Total Error	0.001	0.005	0.51
	Idle Time	0.006	<0.05	0.36

F. Results

Figure 3 summarizes the results of all three experiments. Descriptive statistics are shown in Table II, which highlights the ordering of the results as follows: green shows the top-performing motion type, followed by yellow, which is followed by orange (worst performing). If the two top- or bottom-performing motion types are not significantly different, they are both colored yellow. Inferential statistics for the effect of motion type on all measures are shown in Table III. Table IV shows pairwise comparisons for significant effects. The results for each experiment are briefly summarized in the paragraphs below.

Mico: The results from the Mico experiment show both *straight* (*H1*) and *legible* (*H2*) outperform *simple* in all metrics except idle time. For idle time only *straight* (*H1*) outperforms *simple*. Comparison of *legible* and *straight*

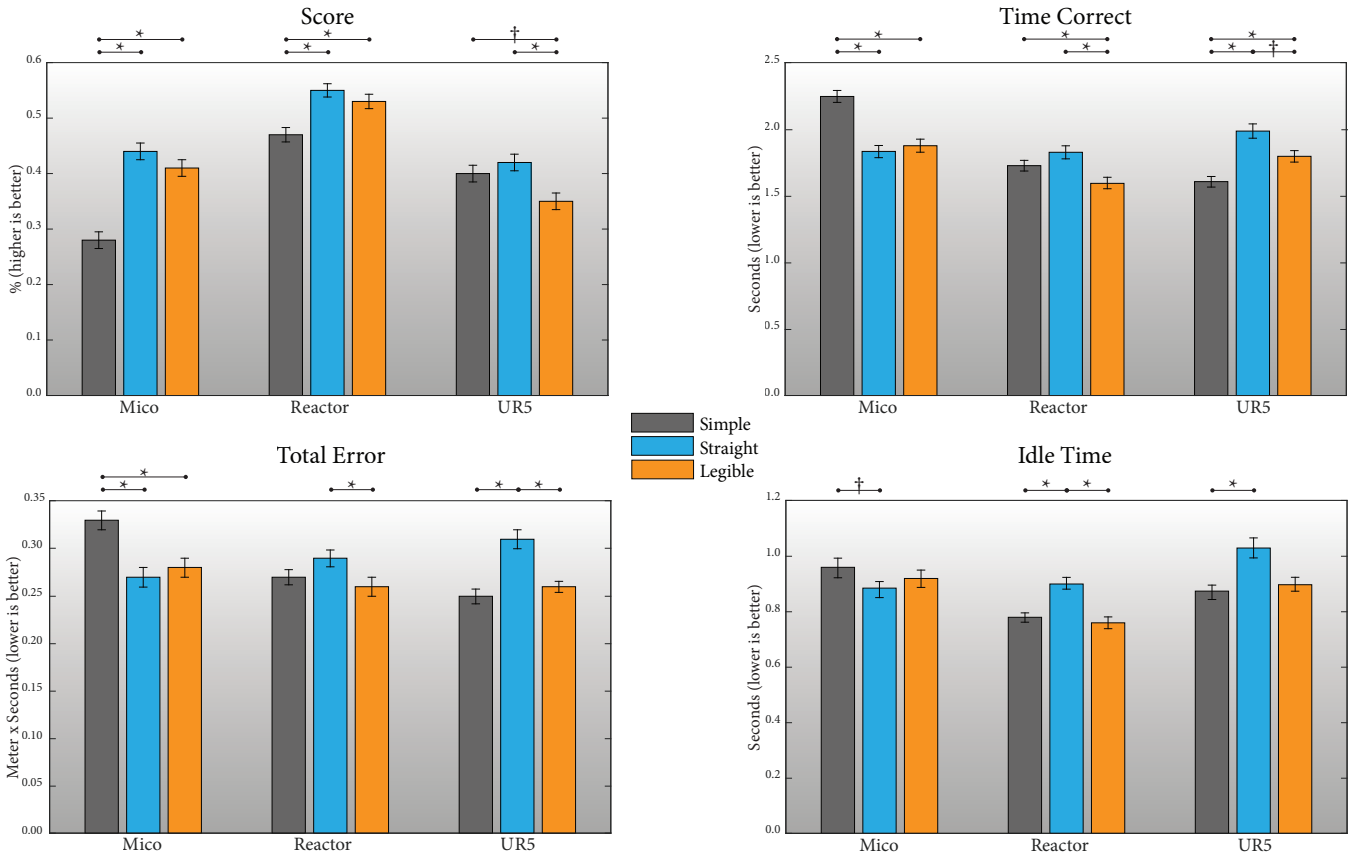

Fig. 3. The result from the three online experiments with all three robots (*Kinova Mico*, *PhantomX Reactor*, and *UR5*) to compare *legible*, *straight*, and *simple* motions. Graphs are shown for each objective measure: *Score* (top-left), *Time Correct* (top-right), *Total Error* (bottom-left), and *Idle Time* (bottom-right). *Error bars are standard error.* $*$ *denotes* $p \leq 0.01$, $\dagger$ *denotes* $p \leq 0.05$

showed no significant differences for any measures. This result supports findings by Dragan et al. [2], [4] that show *legible* motion outperforming *simple* motion.

Reactor: The results from the Reactor experiment tell a different story. While *straight* and *legible* both outperform *simple* when considering score (relative path remaining), *legible* significantly outperforms both *straight* and *simple* in time correct (absolute time). The difference in time correct is not significant between *straight* and *simple*. The performance of *legible* and *simple* are not significantly different for idle time and total error, but both outperform *straight* motion for idle time. Finally, the difference between *simple* and *straight* is not significant for total error. *H1* is supported by the score metric, but contradicted by the idle time metric. *H2* is supported by the score and time correct metrics.

UR5: Finally, results from the experiment involving UR5 tell a third story; *straight* and *simple* both outperform *legible* when considering score. The difference in score between *straight* and *simple* is not significant. *Simple* outperforms both *straight* and *legible* in time correct, and *legible* outperforms *straight* in time correct. The performance of *legible* and *simple* are not significantly different for idle time and total error, but both outperform *straight* motion for total error. Finally, the difference between *legible* and *straight* is not significant for idle time. *H1* is contradicted by the time correct, total error, and idle time metrics. *H2* is contradicted by the score and time correct metrics.

G. Discussion

Our findings, including further post-hoc analyses conducted on our data, point toward three key factors that shape observer performance: relative vs. absolute time, viewpoint, and robot kinematics.

1) Relative vs. Absolute Duration: Score and time correct both measure of how quickly a participant correctly determined the goal of a motion. However, score is a measure of time *relative* to the total duration of the motion. Since different motion objectives produce motions that are not necessarily equal in duration for the same target, it is possible to score well relative to the duration while not achieving the best absolute time. Absolute time is therefore a more objective metric if motion durations within a single target are not equal across motions types (as in our experiments). Relative durations can be an effective metric if path durations are equal for a specific target across motion types, however. The use of relative durations allows more effective comparison across all targets because of the normalized scale.

2) Viewpoint Dependence: A variable that our motion objectives do not consider is the viewpoint from which they will be seen. We visually inspected our synthesized

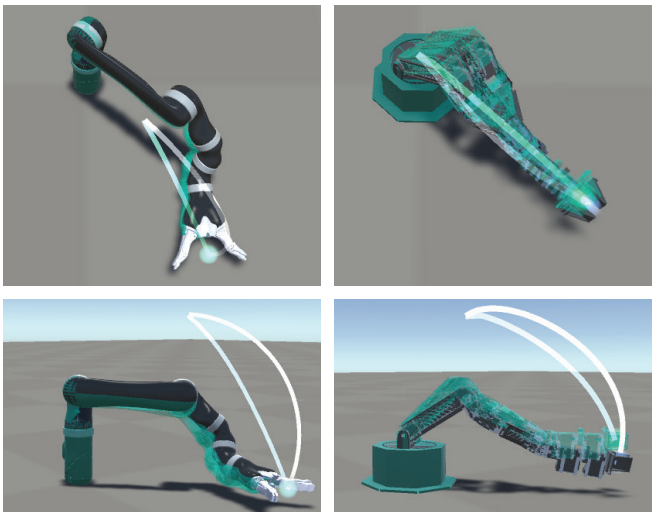

Fig. 4. Illustrations of how viewpoint can affect participant perceptions of arm motion. The top row shows *straight* (blue) and *simple* (white) motions seen from the participant's perspective for both a Mico (left) and Reactor (right). Notice how the motions look very similar for the Reactor from this viewpoint. The bottom row shows from a different viewpoint that these motions are actually very different.

motions from our participant's perspective and observed that while the motions may actually be different, they can look very similar from the participant's perspective. We hypothesize that viewpoint may be important for planning intent-expressive motions. In animation, viewpoint can be thought of as the principle of "staging" [14]. Figure 4 shows the issues that a participant's vantage point can create.

3) Robot Kinematics: The three robots tested each give different results. This may be explained by their differing kinematic designs. The Reactor has only 5 degrees of freedom, which limits its ability to achieve straight paths. In contrast, the UR5 is very capable of achieving the straight objective. However, given the UR5's kinematics, simple paths have nearly straight end-effector paths, albeit with much shorter total arm movement. Simple paths are very different for the Mico robot, which has a series of non-orthogonal "wrists." The Mico's non-anthropmorphic design may mean that simple paths are less familiar to viewers, which may explain the under-performance of the simple condition for the Mico.

V. Conclusion

In this paper, we argue that the design and analysis of intent-expressive robot-arm motions must consider many factors; the choice of robot, viewpoint of the user, and metrics used for analysis will impact the choice of how to synthesize motion to express intent. We selected three very different robots in both degrees of freedom and levels of anthropomorphism to evaluate. The different motion types performed differently for each robot from a consistent vantage point. The metrics used also tell different stories within each robot's analysis based on whether they consider relative or absolute metrics. Since the relative metric is sensitive to path duration, it has extremely limited use.

Across all of the studies, we have shown that both simple and straight-line motions can be as expressive as state-of-the-art legible motion synthesis. While other methods for planning legible motions that could perform best for all robots could very well exist, the evidence in this work suggests that the effects of current methods are predictable on a robot-by-robot basis and under certain conditions, such as viewing angle.

Key areas of future work include identifying how features of a robot's kinematics affect synthesizing intent-expressive motion and investigating the effects of observer viewpoint in motion synthesis. Finally, research must continue exploring a one-size-fits-all legible motion-synthesis method.

Acknowledgments

This work was supported by National Science Foundation National Robotics Initiative Award #1208632.

References

[1] D. Bortot, M. Born, and K. Bengler, "Directly or on detours?: how should industrial robots approximate humans?" in *ACM/IEEE HRI*, March 2013, pp. 89–90.
[2] A. Dragan, K. Lee, and S. Srinivasa, "Legibility and predictability of robot motion," in *ACM/IEEE HRI*, Mar. 2013, pp. 301–308.
[3] A. Dragan and S. Srinivasa, "Generating Legible Motion," in *RSS*, June 2013.
[4] A. Dragan, S. Bauman, J. Forlizzi, and S. Srinivasa, "Effects of Robot Motion on Human-Robot Collaboration," in *ACM/IEEE HRI*, 2015, pp. 51–58.
[5] M. Zhao, R. Shome, I. Yochelson, K. E. Bekris, and E. Kowler, "An Experimental Study for Identifying Features of Legible Manipulator Paths," in *ISER*, 2014.
[6] A. Witkin and M. Kass, "Spacetime Constraints," *ACM SIGGRAPH Computer Graphics*, vol. 22, no. 4, pp. 159–168, Aug. 1988.
[7] M. Gleicher, "Retargeting motion to new characters," in *ACM SIGGRAPH*, July 1998, pp. 33–42.
[8] M. Gleicher, "Comparing constraint-based motion editing methods," *Graphical Models*, vol. 63, no. 2, pp. 107–134, Mar. 2001.
[9] S. Blakemore and J. Decety, "From the perception of action to the understanding of intention," *Nature Reviews Neuroscience*, vol. 2, no. 8, pp. 561–567, Aug. 2001.
[10] G. Gergely, Z. Ndasdy, G. Csibra, and S. Br, "Taking the intentional stance at 12 months of age," *Cognition*, vol. 56, no. 2, pp. 165–193, Aug. 1995.
[11] K. Kamewari, M. Kato, T. Kanda, H. Ishiguro, and K. Hiraki, "Six-and-a-half-month-old children positively attribute goals to human action and to humanoid-robot motion," *Cognitive Development*, vol. 20, no. 2, pp. 303–320, Apr. 2005.
[12] C. G. Atkeson and J. M. Hollerbach, "Kinematic features of unrestrained vertical arm movements," *Journal of Neuroscience*, vol. 5, no. 9, pp. 2318–2330, Sept. 1985.
[13] P. Morasso, "Spatial control of arm movements," *Experimental Brain Research*, vol. 42, no. 2, pp. 223–227, Apr. 1981.
[14] F. Thomas, O. Johnston, and F. Thomas, *The illusion of life: Disney animation*. Hyperion, 1995.