Understanding object properties from the observation of the action partner

Alessandra Sciutti, Laura Patanè, Francesco Nori, Giulio Sandini
Robotics, Brain and Cognitive Science Department
Istituto Italiano di Tecnologia
Genoa, Italy
Email: alessandra.sciutti@iit.it

***Abstract*—Humans often collaborate with each other, achieving an astounding coordination with no need of complex verbal instructions. One key element at the basis of this ability is the mutual implicit understanding of the two partners. Just by looking at each others' action, the two collaborators can infer several action related information which otherwise should have been explicitly pointed out, reducing the fluidity of the interaction. For instance, often one can infer, just by observing the other passing an object, some properties of that object which would not be otherwise accessible, as its temperature or its weight. This understanding allows both partners to be always prepared in advance, when their turn comes to act on the object, yielding to a fluid and timely interaction. Such efficient collaboration would be extremely desirable also between humans and robots. We suggest that a fundamental step toward this aim consists in the achievement of an action-based mutual understanding between the two partners. In this work, we evaluate whether humans can read a not-visible property of a manipulated object also when the manipulator is a humanoid robot. We have therefore designed a simple set of lifting behaviors for a humanoid robot and we have measured in which conditions human subjects are able to infer object weight from robot motion. Moreover, we have assessed whether this information could also be used by humans to improve their subsequent actions on the same object. The results suggest that if robot motion is planned taking into account its understandability by a human observer, the implicit information transfer between robot and human can be as efficient as that observed between humans. These findings indicate that to achieve fluid human-robot collaboration, an effort needs to be made at the planning level, to design robot motions which are not only efficient energy-wise, but which also guarantee an automatic transfer of object-related information to the partners.**

I. Introduction

In everyday life it is usual to see humans helping each other, from the kid passing a toy to his friend who cannot reach it, to the mechanic, to whom the assistant hands over the needed tool at the right moment. In all these situations, collaboration arises naturally, often even without the need of words. This implies that the observation of the other's action is enough to communicate to the observer the actor's intention and needs. This ability is common even to young children, who naturally exhibit a proactive helping behavior [13]. Moreover, such implicit communication usually informs the action partners not only about the goals or the intentions of the cooperator, but conveys also details about the objects at hand. For instance, the complexity of holding a slippery object, or an object which is too warm or too heavy, can be inferred easily from the observation of the person who is carrying it. This implies that when the object will be passed to me, I will be prepared to handle it correctly, with no need of a precise verbal explanation of the potential danger associated to the passage. Even when the situation is not dangerous, knowing in advance objects properties (e.g., their weight) allows to improve the planning of subsequent actions on the same object, increasing the efficiency and fluidity of the interaction.

One important difficulty when interacting with a robot is that it is usually designed without taking into account this kind of implicit stream of information, which in humans is associated even to the simplest action. Indeed, often robot motion planning is designed to optimize different motion parameters as energy consumption, motion smoothness or user safety. Of course, these aspects are fundamental in the economy of the collaboration. However, disregarding the implicit communication so common in human interaction can determine severe drawbacks. In fact, when receiving a parcel from a robot, the human collaborator has to either adjust his force reactively, or be informed explicitly about object weight, which requires additional attention. Therefore, this approach works quite well in situations in which the objects to be passed are well known in advance by all the co-workers, but in an unstructured task it would often result in a not efficient cooperation.

Moreover, while in general robotic devices have a very different embodiment from humans, this difference diminishes between human and humanoids. So, if even a not experienced user can assume that industrial robots are stronger and more heat resistant that a human, this intuition can not be as easily applied to a robotic platform whose size and dimensions are not too dissimilar from those of a man or of a child. This similarity could induce subjects to attribute to the robot the same range of strength as a human, hence determining a mis-reading of the information associated to its motion (e.g., understand the rapid lifting of an object as if the object was particularly light), with negative consequences on the interaction. This could of course have a further impact also on the acceptance of the robot: an incongruence between expectations and real robot behavior determining a failure of an hand over could undermine the trust in the robot - something that, once lost, is very difficult to be re-gained [4].

An important question therefore becomes whether it is possible to make a robot partner able to implicitly communicate with a human as humans naturally do. This topic has been extensively investigated in the context of human-robot interaction, with particular emphasis on the role of non-verbal communication (in terms of pose, expressions, gestures, actions) in allowing the transmission, between partners, of information about the status of the robot, its future intentions, its availability as helper or its need for help [12, 1, 8]. We are interested to extend this capability to communicate through robot motion also the hidden properties of the manipulated object, to facilitate the future collaboration in the use of the same object, with no need of a dedicated communication channel.

In this paper we have addressed the feasibility of this objective, by evaluating the possibility for a humanoid robot to provide information about the weight of an unknown object, just by performing a lifting action. Humans have in fact been shown to be very good at inferring object weight from the observation of someone else's action [11, 3]. We have therefore designed a set of simple robot lifting actions that could potentially convey weight cues to the human partner and we have tested whether this information can be actually exploited by a human observer to plan his own subsequent actions.

II. Methods

The study consisted of two separate experiments, one aimed at evaluating the ability to judge weight from the observation of a lifting action (*Weight Judgment Task*, Fig. 1A) and one designed to assess whether such weight estimate is used by the observer to prospectively plan his/her own successive lifting (*Motor task*, Fig. 1B and C). The first task consisted in asking subjects to estimate the weight of a few opaque bottles after having watched movies of an actor (either human or robot) lifting them. In the second experiment, instead, participants were requested to observe a lifting action performed in front of them (again either by a human or by a robot) and then to actively lift the same object previously lifted by the actor. In both tasks, the robotic actor was the iCub robot, i.e., a humanoid robot developed as part of the EU project RobotCub, approximately $1m$ tall and with the appearance of a 3.5 years old child [9]. Robot motion was designed by defining three viapoints in joint coordinates and by manually assigning the velocity at each joint and the duration of each movement phase by using a graphical user interface. More details about the human and the robot lifting motions are reported in Fig. 2 and in Fig. 3, respectively. The motivations behind the selection of the robot kinematics are discussed in the Results section. Below we report a detailed description of the two tasks.

A. Subjects

30 healthy, right-handed subjects agreed voluntarily to participate in the experiments. 12 subjects participated to the

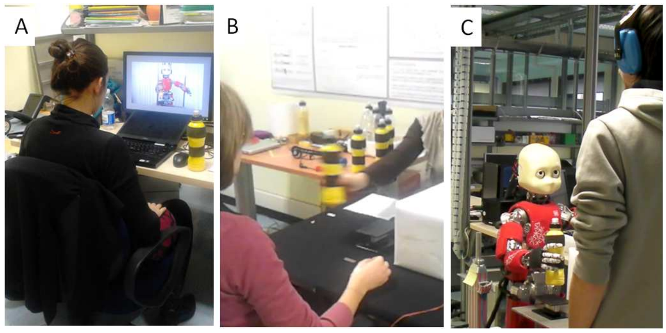

Fig. 1. Snapshots of the experimental conditions: A) *Weight Judgment Task*; B) *Motor Task* - HUMAN condition; C) *Motor Task* - ROBOT condition.

Weight Judgment Task (mean age: 32 years $\pm$ 10.4 (SD); 3 females and 9 males), while the remaining 18 subjects (mean age: 29 years $\pm$ 8 (SD); 9 females and 9 males), took part to the *Motor Task* (8 subjects in the *Human* condition and 10 subjects in the *Robot* one).

B. Weight Judgment Task

Subjects were instructed to observe several videos in which an actor (either a human or a robot) lifted and placed a set of bottles apparently identical but of different weights (respectively $100, 200, 300$ and $400g$ for the HUMAN video and $100, 250$ and $400g$ for the ROBOT video, see Fig. 1A). After the observation of the stimuli, participants had to judge the observed weight, by selecting a number between 1 and 9, corresponding to a range from $50g$ to $450g$ with $50g$ increments. When the actor was HUMAN, she was informed of object weight before action execution. This choice was made to mimic the natural object-passing scenario, where the person who manipulates the object first is usually aware of its weight (see Fig. 2 for a description of actor's movements characteristics). When the actor was ROBOT, its motion was planned as described before (and in Fig. 3). Before each experiment subjects were asked to lift nine bottles from $50g$ to $450g$ with a step of $50g$ in order to have a clear range of the weights that they would have had to recognize subsequently. The two sets of videos (HUMAN and ROBOT) consisted of 32 trials for the HUMAN set (i.e., eight lifting movements for each of the four bottles) and 30 trials for the ROBOT set (i.e., ten lifting movements for each of the three bottles) both performed in random order.

C. Motor task - weight estimation for movement planning

This task was aimed at evaluating whether the observation of human or robotic lifting actually allows the partners to prepare in advance their own next action on the same object. We asked subjects to observe a human or a robot lifting and placing an opaque bottle (as in the previous experiment) and then to lift the same bottle themselves from a custom-built analog scale (see Fig. 1B and C). For each lift, we measured participant's loading phase duration, i.e., the time spent by the subject between the beginning of the lifting and the time the bottle actually abandoned the supporting surface

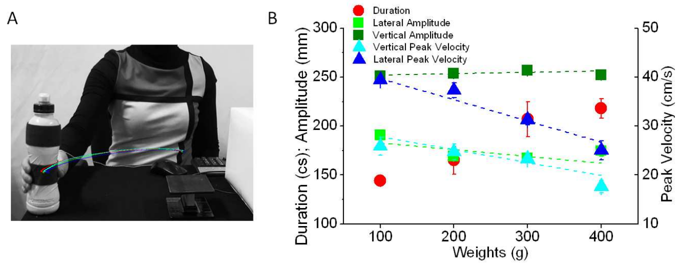

Fig. 2. Human lifting action features. A) Snapshot of the human lifting motion with superimposed the trajectories for different weights (blue - 100g, red - 200g, green - 300g and cyan - 400g. B) Kinematic features of the human actor's lifting actions as a function of object weight. Error bars represent standard errors of the average.

(see [6] for details on the method). Ideally, humans tend to maintain this time invariant with respect to the lifted weight (at least for a small range of loads), but to do so they have to know in advance the weight they will have to lift. In fact, until the complete object lift off, the subject does not receive any feedback about object weight, performing the force loading in a complete feed forward fashion [7]. A perfect knowledge of the object prior to the lifting would then imply a minimal dependence of loading phase duration on its weight, as it will be predictively compensated. When the object is unknown, instead, lighter-than-expected objects will be raised too early, and heavier-than-expected loads will take longer time to lift off (see Fig. 5 A for a representation of this trend; blue line and symbols for the known object weights and green line and symbols for the unknown ones). For each subject we computed the percentage of variation in loading phase duration between the heaviest - $400g$ - and the lightest - $100g$ - weight (Perc. Var. in the following). With this parameter we could infer the amount of information about object weight gained from someone else's (human or robot) action observation.

Both when the actor was a human or a robot, we performed additional control conditions. More precisely:

The HUMAN experiment was structured in three different phases: *Known*, in which the subject had to lift the bottle and bring it on the scale from a starting point at about $25cm$ distance. This was done to let subject directly experience the weight of the bottle. After that, he had to lift the bottle from the scale and place it on a higher support; *Unknown*, in which the subject had to perform the lifting of the bottle directly from the scale without having any previous information concerning its weight; *Observation*, in which the subject had to observe a human actor performing the lifting movement from the starting point (the same of the phase *Known*) to the scale and then perform the lifting action. Each of the three phases consisted of 32 trials in random order (i.e., eight lifting movements for each of the four bottles). Also the order of the three phases was randomized between subjects.

The ROBOT experiment, performed on a different group of subjects, was structured in two phases: *Control* and *Robot*. In the *Control* phase subjects had to perform the lifting of the bottle after having observed the robot lifting it in a invariant way, i.e., with no variation in the action kinematics as a function of the lifted weight. In the *Robot* phase the robot lifted the bottles with different action kinematics for different weights (see Fig. 3 for more details. Note that in the *Control* condition the motion features of the lifting applied to all the weights were those for the $400g$ bottle in this phase). Each of the two phases consisted of 15 trials performed in random order (i.e., five lifting trials for each of the three bottles). Also the order of the phases was randomized between subjects.

In the HUMAN *Observation* condition and in both ROBOT conditions, subjects had to provide also a verbal estimate of object weight, before the execution of the lifting action, choosing a value between $50g$ to $450g$ with $50g$ increments.

The analog scale used in this task was obtained by inserting a load sensor into a custom built structure and recorded load forces during the lifting of the bottles at a frequency of 1000 *Hz*. In the HUMAN experiment the kinematics of the action was recorded at 250 *Hz* by means of an infrared marker (Optotrak Certus® System, NDI) placed on the actor's right hand at the level of the metacarpophalangeal joint of the middle finger. The kinematics of the iCub actions was recorded by computing the end effector position from the joints state measured through the encoders of the motors in the robot arm and saving the hand coordinates on a file at a sample frequency of 100 *Hz* [10]. Before each experiment subjects were trained with a metronome to perform the lifting of the bottles from the scale at about a constant pace (each movement lasted about $1.5s$). During this training phase subjects lifted two transparent bottles (100 and $400g$ respectively) used as samples in order to provide an idea of the range of weights used in the real experiment.

III. Results

In this work we have studied whether it is possible to make a robot implicitly communicate with its motion an invisible property of the object it is manipulating, i.e., its weight.

A. Weight judgment from Human Observation

First we have investigated how good humans are at judging weight after the observation of a lift-and-place action performed by a human actor. Subjects were requested to estimate how heavy was an opaque bottle after observing a movie depicting a human actor lifting it. Participants exhibited a quite good precision in weight estimation, with an average absolute error of $80g \pm 4g$ (SE) for weights ranging from $100g$ to $400g$ in $100g$ steps (see open symbols in Fig. 4A). Considering the signed estimates, they tended to overestimate lighter bottles ($+52g$ for the $100g$, $+7g$ for the $200g$) and underestimate the heavier ones ($-69g$ for the $300g$, $-59g$ for the $400g$), but the estimated weight increased with the presented weight. In particular the average regression of estimated weight over real weight was 0.59 ± 0.04 (SE) with an average R^2 of 0.77 ± 0.07 (SE). A One-Way

Repeated Measures ANOVA on the estimates, followed by a Bonferroni post-hoc comparison, showed that the estimates for all the bottles were significantly different among each other ($p < 0.001$), with the exception of the $200g$ and $300g$ for which the difference in judgment was not significant.

The second question we were interested to was on which kinematics information subjects mostly based their judgment. The lifting actions recorded from the human actor were natural actions, in which multiple kinematics parameters varied as a function of object weight (see Fig. 2A). In particular, between lifting $100g$ and $400g$ movement duration varied of 51%, peak vertical velocity varied of -32%, peak horizontal velocity of -36%, while lateral and vertical movement amplitude varied of 0.4% and -8% respectively. To assess which of these parameters was more informative for observers' weight evaluation, we ran a multiple linear regression of the estimates with respect to the previous movement kinematics properties over the whole sample. The R^2 of the multiple regression was 0.33 and the regression was significant ($P < 0.001$). However, the only variables for which the estimated coefficients in the linear model were signifcantly different from 0 were peak vertical velocity ($p < 0.001$) and movement duration ($p < 0.05$, one-sample t-tests). Hence, we evaluated for each single subject the linear regression of weight estimate over these two main movement properties. The mean R^2 of the multiple regressions was 0.35 (range $0.07-0.48$), and the regression was significant in ten of the twelve subjects. Apparently therefore the cues particularly relevant for weight estimation on the basis of action observation are the vertical velocity of the action (as already suggested by [3]) and its duration, which in turn are strongly correlated in the natural lifting action showed in the videos ($P < 0.001$, $R^2 = 0.24$).

B. Weight judgment from Robot Observation

We have therefore designed a set of robotic lifting actions in which vertical velocity and hence movement duration varied as a function of object weight (see Fig. 3 for details) and we have tested, again with movies, whether this simple modification in robot kinematics was sufficient to allow subjects to discriminate the weight of an opaque bottle.

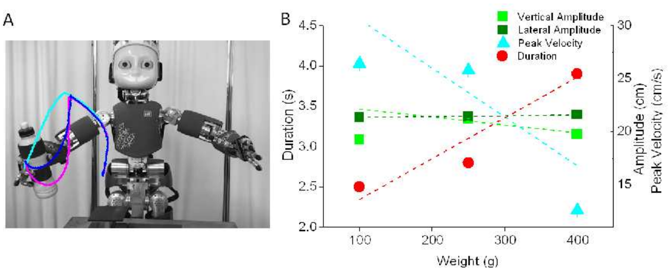

Fig. 3. Robot lifting action features. A) Snapshot of the robot lifting motion with superimposed the trajectories for different weights (blue - 100g, magenta - 250g and cyan - 400g. B) Kinematic features of the lifting actions as a function of object weight. In the *Control* condition the robot always moved its arm with the kinematic features associated in this graph to the $400g$ weight.

Subjects showed a remarkable ability in weight estimation also in presence of the robotic actor, with an average absolute error comparable to that exhibited in the human task ($72g \pm 6g$ (SE) for the discrimination among three bottles of $100g$, $250g$ and $400g$, see also Fig. 4). Also in this case, they tended to overestimate the lightest bottle (on average of about $+64g$) and underestimate the heaviest one ($-19g$), but the estimate increased significantly with the presented weight ($P < 0.001$ in a One-Way Repeated Measures ANOVA). Bonferroni post-hoc showed that subjects could discriminate between the $400g$ bottle from both the $100g$ and the $250g$ ones ($Ps < 0.001$), while the difference in the estimates between the two lighter objects was not significant. The average regression of estimated weight over real weight was 0.72 ± 0.04 (SE) with an average R^2 of 0.62 ± 0.1 (SE).

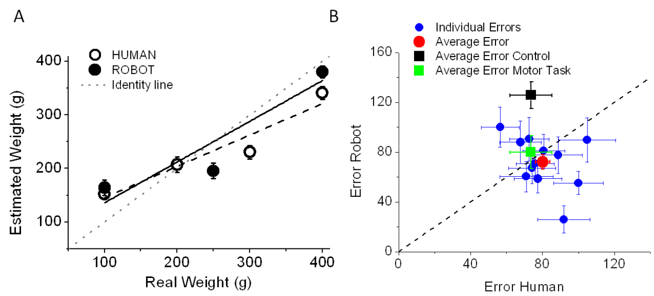

Fig. 4. *Weight Judgment Task* results. A) Average weight estimates in the HUMAN and the ROBOT condition. B) Individual subject's absolute error of the estimate in the ROBOT condition over his/her own performance in the HUMAN condition. Most data lie on the identity line, indicating a similar performance in the two tasks. The red dot represent the sample average. The squares indicate the average error in the weight estimate for the subjects in the *Motor task* (see section III-C). The black square refers to the weight judgment associated to the observation of the robot moving invariantly with respect to object weight (*Control* condition), while the green square represents the average judgment when robot motion varied as a function of object weight, as in the movies used in the *Weight Judgment Task*. Note how also in the *Motor task*, weight judgment is similar after the observation of a human or a robot actor (i.e., the green square is on the identity line), as far as the robot moves appropriately to the lifted weight. When this is not the case, the errors in the robot condition are much higher than in the human one (black square lying over the identity line).

A comparison between the performance in weight judgment when the robot or the human was the lifter did not show any significant difference in terms of the slopes of the regression of estimated weight over real weight nor in terms of average absolute errors in the estimate ($P > 0.05$ in the corresponding paired-sample t-tests). This can be seen in Fig. 4 both for population averages (panel A) and at the individual level (panel B). Hence, the simple modification in the robot lifting actions proposed succeeds in allowing the human partner to correctly estimate the lifted weight, as much as a human lifting action would.

C. Weight estimation for action planning (Motor task)

In this task we have evaluated whether the observation of a human lifting allows the partner not only to correctly judge the object weight, but also to prepare in advance his next action on the same object. To this aim, we assessed participants' loading phase durations for the lifting of different weights (see [6] and Methods section for details) as a function of their a priori knowledge of the object to be lifted. First, we compared a condition where subjects had no information about objects weight before the lift ("Unknown"), with a condition in which they had the possibility to lift it before the measure was performed ("Known") or they could infer its weight by observing a human lifting it ("Observed"). The results confirmed that knowing object weight in advance allows for maintaining much more constant the loading phase duration than lifting an unknown weight (compare a Perc. Var. of $31\% \pm 8\%$ in the "Known" condition with the $100\% \pm 11\%$ of the "Unknown" and the blue and the green line and symbols in Fig. 5A, respectively, for a representative subject). Furthermore the observation of a lifting action provides enough information to significantly increase such constancy (Perc. Var. $63\% \pm 8\%$, $P < 0.01$ for all comparisons, in a One-Way Repeated Measures ANOVA on Perc. Var. followed by Bonferroni Post-hoc, see red line and symbols in in Fig. 5A).

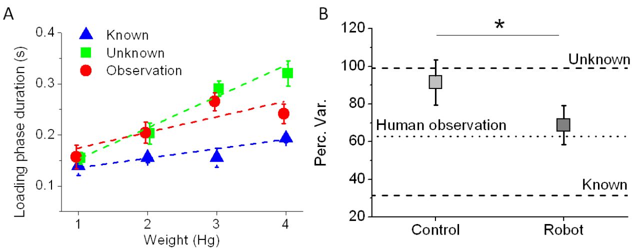

Fig. 5. Lifting action results. A) Loading phase duration as a function of object weight in different conditions for a representative subject. B) Average percentage of variation in loading phase duration between heavy (400g) and light (100g) bottles for Robot observation across all subjects. The horizontal lines indicate the corresponding percentage of variations measured during the human condition. The star indicates significant difference evaluated with a pair sample t-test.

After that, we evaluated whether the observation of the robotic lifting previously designed could be as efficient in facilitating human action planning. Moreover, we wanted to assess whether the proposed kinematic modification was actually necessary to improve human performance in a collaboration, or if a standard robot motion could achieve a similar result. Therefore, we evaluated human stability in loading phase duration also when the robot was lifting the opaque bottle with a motion not explicitly planned to vary with object weight, but just aimed at the maximization of motion smoothness. The results clearly showed that the weight cues independent of robot kinematics (e.g., the slip among robot fingers of the heavier bottles) are not enough to allow for a correct planning of the subsequent action on the bottle (Perc. Var. on average $91\% \pm 12\%$ see light gray square in Fig. 5 B). On the contrary, the proposed robot lifting is very efficient in allowing the human partners to anticipatorily plan their subsequent actions, significantly reducing the dependency of the loading phase duration on object weight (see dark gray square in Fig. 5 B; Perc. Var. on average $69\% \pm 10\%$, $P < 0.05$ in a pair-sample t-test).

IV. Discussion

Humans are extremely good at collaborating. This happens effortlessly in every day life, when we prepare together with a friend something in the kitchen or when a worker cooperates with his co-workers in building a new piece of furniture. Interestingly, the fluidity in the interaction seems to be connected to the ability to continuously exchange with the partner implicit information about what we are going to do next or which tool we are going to grasp. This stream of information is so rich that it even allows the partners to infer object properties which would be not perceivable at first sight. So, if someone is passing me a heavy suitcase or a hot dish, a glance to his movement is enough to tell me that I will have to be careful and tune my own action appropriately.

In this paper we have investigated this particular case of implicit information transfer, i.e., the communication of the weight of an unknown object, mediated by the kinematics of the action with which this object is passed. Such information is not only easily extracted by humans during the observation of their human partners, but is also exploited during the subsequent actions on the passed object in a predictive fashion. On the contrary, at least in a small range of weights as the one used in this experiment ($100g - 400g$), the change in weight alone is not sufficient to modify the way a humanoid robot lift the objects. Therefore, a robot motion not explicitly designed to somehow communicate weight to the human partner does not allow the collaborator to be prepared to the weight to be lifted. However, a simple modification which recalls the dependency of lifting velocity to object weight common to human behavior, is enough to make the implicit weight understanding in the robot lifting as effective as during human-human collaboration both for weight judgments and for the preparation to subsequent manipulations of the object.

We point out that achieving such implicit communication requires an extra effort in the planning of robot motion, aimed at considering the observers' intuitive comprehension of the observed robot behavior. This approach is strongly connected to the idea that seamless collaboration in human-robot interaction depends on the understandability and the intuitiveness of the robotic behavior [5, 2, 8]. In particular we propose that, in addition to make the goal of the robotic action, its uncertainty about the task or its availability as a helper transparent to the human partner, it would be important to communicate implicitly also the properties of the object at hand, to facilitate his/her future interactions with it.

A possible confound in our results relies in the fact that the lifting actions performed by the human and by the robotic actor were different, with the robot action spanning more the vertical direction and the human action being more horizontally oriented. This could have in theory introduced a difference in the difficulty of the task, potentially facilitating the interpretation of the weight in the robotic condition. We cannot exclude this hypothesis, although the facilitation - if present - did not make the task trivial, as subjects' performance was comparable to that measured in the task with the human actor, never reaching negligible errors. Moreover, the comparison between the two robotic conditions, in which robotic motions were similar, confirms that when the robot behavior is not appropriately designed to foster weight understanding, subjects are not able to infer this object property, even in this possibly simpler situation. Further studies are now planned to design a more precise comparison between human and robotic actions and to assess also the relevance of the adoption of a human-like versus a robotic-like motion kinematics. In particular, future research will be aimed at assessing the impact of this implicit object properties communication in more interactive scenarios, as for instance, object handover, where a prospective estimate of the object to be received could be determinant in making the object transfer successful and effortless.

V. Conclusion

Achieving the same fluidity and timing observed usually between humans in human robot collaboration is a long term aim of robotic research. A key element to this ideal efficiency relies in the implicit action-based communication common to humans, which allows both partners to be continuously informed about others' intentions and even about invisible properties of the object at hand. In this work we have shown that it is possible to achieve an efficient implicit communication also between a humanoid and a human, as far as the robot passing action is planned by taking into account also its understandability from the point of view of the human partner. Hence, humans can infer which weight a robot is lifting as efficiently as they do for a human lifting and use this acquired knowledge of weight to plan in advance their own lifting action. This approach can be applied to several other object properties (e.g., temperature, weight distribution, slipperiness) or action properties (e.g. future use of the manipulated object), which can be inferred by the collaborator just by observing the actor's motion, with no need of an explicit explanation. We believe that allowing the inference of similar object or action properties from robot motion could be a promising path to foster more natural and proficient human-robot collaborations.

Acknowledgments

The authors would like to thank Marco Jacono, for his help in setting up the experiment; Gabriel Baud-Bovy, for his advice on force measurements and bibliography; Laura Taverna for pictures and movies preparation; Vadim Tikhanoff and Francesco Rea for their support in the use of iCub. Parts of this work have been submitted to IROS 2013 and ICDL 2013.

References

[1] C. Breazeal, CD. Kidd, AL. Thomaz, G. Hoffman, and M. Berlin. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *Intelligent Robots and Systems (IROS 2005)*, 2005.

[2] Crystal Chao, Maya Cakmak, and Andrea L. Thomaz. Transparent active learning for robots. pages 317–324, 2010. URL http://dl.acm.org/citation.cfm?id=1734454.1734562.

[3] A. F de C Hamilton, D. W. Joyce, J. R. Flanagan, C. D. Frith, and D. M. Wolpert. Kinematic cues in perceptual weight judgement and their origins in box lifting. *Psychol Res*, 71(1):13–21, Jan 2007. doi: 10.1007/s00426-005-0032-4. URL http://dx.doi.org/10.1007/s00426-005-0032-4.

[4] M. Desai, P. Kaniarasu, M. Medvedev, A. Steinfeld, and H. Yanco. Impact of robot failures and feedback on real-time trust. pages 251–258, 2013. ISSN 2167-2121. doi: 10.1109/HRI.2013.6483596.

[5] Anca D. Dragan, Kenton C.T. Lee, and Siddhartha S. Srinivasa. Legibility and predictability of robot motion. pages 301–308, 2013. URL http://dl.acm.org/citation.cfm?id=2447556.2447672.

[6] A. M. Gordon, G. Westling, K. J. Cole, and R. S. Johansson. Memory representations underlying motor commands used during manipulation of common and novel objects. *J Neurophysiol*, 69(6):1789–1796, Jun 1993.

[7] R. S. Johansson and G. Westling. Coordinated isometric muscle commands adequately and erroneously programmed for the weight during lifting task with precision grip. *Exp Brain Res*, 71(1):59–71, 1988.

[8] W. Ju and L. Leifer. The design of implicit interactions: Making interactive systems less obnoxious. *Design Issues*, 24(3):72–84, 2008.

[9] Giorgio Metta, Lorenzo Natale, Francesco Nori, Giulio Sandini, David Vernon, Luciano Fadiga, Claes von Hofsten, Kerstin Rosander, Manuel Lopes, Jos Santos-Victor, Alexandre Bernardino, and Luis Montesano. The icub humanoid robot: an open-systems platform for research in cognitive development. *Neural Netw*, 23(8-9):1125–1134, 2010. doi: 10.1016/j.neunet.2010.08.010. URL http://dx.doi.org/10.1016/j.neunet.2010.08.010.

[10] Ugo Pattacini, Francesco Nori, Lorenzo Natale, Giorgio Metta, and Giulio Sandini. An experimental evaluation of a novel minimum-jerk cartesian controller for humanoid robots. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010.

[11] S. Runeson and G. Frykholm. Visual percepuon of lifted weight. *Journal of Experimental Psychology Human Perception and Performance*, 7:733–740, 1981.

[12] L. Takayama, D. Dooley, and W. Ju. Expressing thought: improving robot readability with animation principles. In *Proceedings of the 6th international conference on Human-robot interaction (HRI 2011)*, 2011.

[13] Felix Warneken and Michael Tomasello. Altruistic helping in human infants and young chimpanzees. *Science*, 311(5765):1301–1303, Mar 2006. doi: 10.1126/science.1121448. URL http://dx.doi.org/10.1126/science.1121448.