
Affordance-Based Altruistic Robotic Architecture for Human-Robot Collaboration

Adaptive Behavior
XX(X):1–19
©The Author(s) 2018
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/

SAGE

Mert Imre^{1,3}, Erhan Oztop², Yukie Nagai³ and Emre Ugur¹

Abstract

This paper proposes a computational model for altruistic behavior, shows its implementation on a physical robot, and presents the results of human-robot interaction experiments conducted with the implemented system. Inspired from the sensorimotor mechanisms of the primate brain, object affordances are utilized for both intention estimation and action execution, in particular to generate altruistic behavior. At the core of the model is the notion that sensorimotor systems developed for movement generation can be used to process the visual stimuli generated by actions of the others, infer the goals behind, and take the necessary actions to help achieving these goals; potentially leading to the emergence of altruistic behavior. Therefore, we argue that altruistic behavior is not necessarily a consequence of deliberate cognitive processing but may emerge through basic sensorimotor processes such as error minimization, i.e. minimizing the difference between the observed and expected outcomes. In the model, affordances also play a key role by constraining the possible set of actions that an observed actor might be engaged in, enabling a fast and accurate intention inference. The model components are implemented on an upper-body humanoid robot. A set of experiments are conducted validating the workings of the components of the model, such as affordance extraction and task execution. Significantly, to assess how human partners interact with our altruistic model deployed robot, extensive experiments with naïve subjects are conducted. Our results indicate that the proposed computational model can explain emergent altruistic behavior in reference to its biological counterpart, and moreover engage human partners to exploit this behavior when implemented on an anthropomorphic robot.

Keywords

altruistic behavior, computational modeling, brain-inspired robotics, affordances, human-robot interaction, goal inference

1 Introduction

It has been observed that at earlier ages chimpanzees and infants tend to help fulfilling others' goals without being explicitly triggered by an external agent, even in the absence of any reward. Accomplishing such helping task seemingly requires estimating others' intentions to predict their future goals. The experiments performed by Warnaken and Tomasello (Warneken & Tomasello, 2007; Warneken & Tomasello, 2006; Warneken, Chen & Tomasello, 2006) showed that the altruistic behavior can be observed even from 14 months of age for infants in simple tasks like handing the objects that are out-of-reach for the others. By observing the stretching of the arm toward an object, infants are able to predict other's future goal of grasping it. Although many studies (Tomasello, Carpenter, Call, Behne & Moll, 2005; Provasi, Dubon & Bloch, 2001; Baldwin, Baird, Saylor & Clark, 2001; Warneken & Tomasello, 2006)

proposed that the infants can indeed help others through some kind of early empathy towards them, recent studies (Kenward & Gredebeck, 2013) suggest that such pro-social behaviors might be due to a desire to accomplish 'unfulfilled goals', evidenced by the fact that 18-month-old infants would help not only to living organisms but also to the spherical objects in accomplishing their predicted goals. Further data (Warneken & Tomasello, 2009) support the idea that the infants are willing to help without any kind of social

¹Bogazici University, Istanbul, Turkey

²Ozyegin University, Istanbul, Turkey

³National Institute of Information and Communications Technology, Osaka, Japan

Corresponding author:

Emre Ugur, Bogazici University, Department of Computer Engineering, Bebek, Istanbul, 34342, TURKEY.

Email: emre.ugur@boun.edu.tr

signal, e.g. when an experimenter whose hands are full of books tries to put them in a closed bookshelf, infants, without any reward expectation, come to open the door so that the experimenter can fulfill his goal.




Figure 1. An example scene from our experimental setup. One subject is seated on one of the long sides of the table and the robot with a Kinect attached on its head stands on the opposite long side. Both are sharing a table top environment with different daily objects.

Inferring others' actions and executing one's own actions are closely intertwined processes. Reuse of cortical circuits for both movement generation and action estimation seems to be a key principle in sensorimotor organization in primates. For complex object manipulation in an environment, an organism extracts object properties relevant for manipulation, detects object affordances, and makes motor plans based on them to fulfill a desired change in the environment. During execution of an action, for example grasping, the agent monitors the control variables (such as distance of the hand to the target) to bring them to their desired values (i.e. error reduction), and check for unexpected disturbances by comparing the predicted change with the actual change (by using internal forward models). In this description, three key notions surface: object affordances, control variable based error reduction, and forward models, which have received considerable attention in the literature. Learned object affordances were used in dexterous behavior generation and imitation in many studies (e.g. (Oztop, Imamizu, Cheng & Kawato, 2006; Ugur, Nagai, Sahin & Oztop, 2015; Ugur, Oztop & Sahin, 2011)); control variable computations and forward models developed for action generation were exploited for mental state inference in (Oztop, Wolpert & Kawato, 2005). The neuroscientific data indicates a tight connection between affordance computation and intention understanding, in particular in relation to mirror neurons as reviewed in (Thill, Caligiore, Borghi,

Ziemke & Baldassarre, 2013). Yet, to our knowledge there are no computational models that integrate affordance extraction with mental state inference mechanisms in a biological plausible manner.

In this study, we develop a biologically inspired computational model of altruistic behavior, and implement it on an anthropomorphic robot, which integrates object affordance computation and intention extraction to explain how altruistic behavior may emerge via the dual use of sensorimotor circuits for action observation and execution. Specifically, our postulation is that altruistic behavior may be triggered by basic sensorimotor processes as error minimization, rather than derived from conscious cognitive processes. The current model builds upon the previous model of Oztop et al. (Oztop, Wolpert & Kawato, 2005) and augments it with affordance extraction mechanism for more accurate goal inference. In the model, the possible set of actions that an observed actor might be performing is limited by the affordances available in the environment which in turn speeds up and refines the intention inference. Although the model components are associated with plausible brain regions, for the robotic implementation a schema based approach (Arbib, 1981; Arbib, Erdi & Szentagothai, 1998) is adopted where some modules are programmed directly, yet others learn from robot's own experience. The neural level implementation of model components is left for a later study. In particular, in this study we aimed to use this model as part of the cognitive architecture of a physical robot that can interact with the humans with acceptable processing delays. Thus, the model implementation has to comply with real-time constraints with real sensorimotor signals and drive a physical robot in contrast to driving a simulated robot with synthetic sensorimotor signals.

In the realized system the robot, under the guidance of the model, interacts with the human by observing his/her actions, and steps in for help when it decides that the action (and the predicted goal) cannot be fulfilled. By using its own sensorimotor capabilities, the robot plans and executes a set of actions, which may be different from the ones used by its human partner. So it is not necessarily a partial imitation that leads to the altruistic help. It is worth underlining that the object affordance extraction capability of the robot -learned by prior robot experience- facilitates accurate goal inference as well as correct task execution, which contribute to the naturalness of the interaction with the robot.

After the computational model has been implemented on the robot, we assessed the sensorimotor capabilities such as the affordance extraction and task execution performance of the obtained system by running validation experiments.

Then, we conducted human-robot interaction experiments with naïve subjects to evaluate the effectiveness of the system as an altruistic partner.

2 Related Work

Altruistic helping behavior was previously realized in a robotic system by (Baraglia, Nagai & Asada, 2016) who utilized mechanisms for predictive learning of sensorimotor information and goal-alignment theory, where the robot continuously updated its predictor to learn action primitives by minimizing a prediction error. The benefit of having this updated predictor is twofold: inferring the future state of the others' actions and estimating a prediction error based on observed and inferred states. As long as the others fail in completing their inferred task, the prediction error grows up to a point where the robot steps in to help them fulfill their goal.

We extend this prediction error minimization approach through exploiting affordances. Affordances have been extensively studied in the last decades for learning action-related object properties, representing the effects of the actions, encoding multi-object models and tool use, multi-step predictions for action planning and human-robot interaction and communication (Jamone, Ugur, Cangelosi, Fadiga, Bernardino, Piater & Santos-Victor, 2016). In particular, object affordances along with human manipulation actions were detected from video sequences and from RGB-D images using an open knowledge based of visual affordances of common household objects (Koppula, Gupta & Saxena, 2013). While such perception-only systems achieve good recognition performance, their performance in relation to action and object affordance cannot be assessed. As discussed in a recent detailed survey on affordances (Zech, Haller, Lakani, Ridge, Ugur & Piater, 2017), solving the correspondence problem and the need for neurally inspired models still remain as open research challenges in affordances research.

A number of studies have worked on inferring goals of others using observations of body parts and object affordances. Zanchettin and Rocco used hand trajectory for predicting the target position from a set of pre-defined positions using Bayesian inference to improve human-robot collaboration (Zanchettin & Rocco, 2017). Another recent work modeled human actions probabilistically by defining object affordances as the combination of information related to the distance and the angular position of the human body, independent of physical properties of the objects (Dutta & Zielinska, 2018). According to the knowledge of the

authors, the novelty of our work can be summarized as: (1) development of an altruistic robot partner with affordances learning and exploitation capability, (2) development of a biologically realistic model of altruistic behavior guiding the robotic implementation, (3) validation of our robot as an intuitive altruistic partner with human-robot interaction experiments.

3 Model

This paper proposes a computational model that is inspired from the sensorimotor processing of primate brain, and that integrates the concepts of affordance, intention estimation and altruistic behavior. The sensorimotor system developed for movement generation is placed in the center of the model to be exploited not only to predict the behavior of others but also to allow the altruistic behavior emergence. To explain the model, we first present the relation of the proposed model to the cortical organization of the brain, then provide the functional description of the model components, and finally explain the deployment of the model on the robot providing the details on how it actually operates. The components and the information flow of the proposed model is presented in Fig. 2

3.1 Relation to the Cortical Organization of the Brain

We envision the model as composed of modules that are active in both action generation and action observation. The sensorimotor processing for manipulation starts from the vision sensors, and might follow different pathways to initiate movement.

Visual Processing: In the model, we lump all the visual processing in the Visual Cortex module, which serves as the source of visual information that drives the other modules. The visual information that is relevant for manipulation and hand monitoring is extracted in the dorsal pathway and superior temporal sulcus, respectively (Nelissen, Borra, Gerbella, Rozzi, Luppino, Vanduffel, Rizzolatti & Orban, 2011). The Caudal Intraparietal Sulcus (cIPS) in the parietal pathway includes neurons that encode surface and axis orientation, and is connected with Anterior Parietal Cortex that hosts neurons that encode object shape and size (Sakata, Taira, Kusunoki, Murata & Tanaka, 1997). Finally Ventral Intra-Parietal Area neurons encode spatial location in various reference frames and project to premotor area F4 (Luppino, Murata, Govoni & Matelli, 1999). To summarize, these areas include rich object information for manipulation, i.e. affordances and the whereabouts of the objects. The Superior

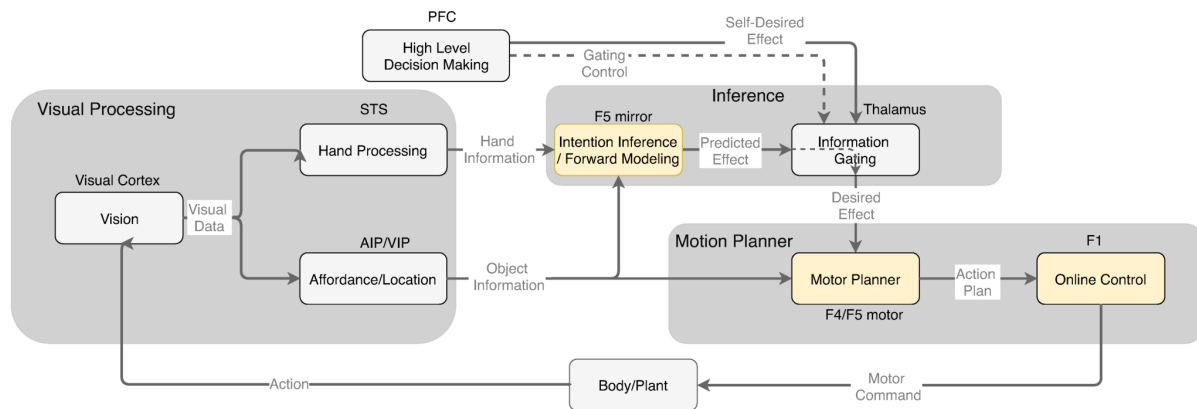


Figure 2. Proposed altruistic behavior model. This model is valid for both an agent that is observing an actor or an acting agent. The related brain parts for each of the modules are provided outside the corresponding box. The modules are grouped to present the implementation modules. Motor area/functions are colored as yellow.

Temporal Cortex represents body and its parts in terms of visual appearance. In particular, it contains neurons that represent hand-in-action regardless of the owner of the hand. STS and AIP are reciprocally connected (Nelissen, Borra, Gerbella, Rozzi, Luppino, Vanduffel, Rizzolatti & Orban, 2011) and thus may form the basis of control variable computation (such as distance or error) with probably help from other neighboring regions such as PG (Oztop & Arbib, 2002). Both STS and AIP project to ventral premotor cortex (area F5), which is a critical area that serves as the interface between the sensory and motor system for hand actions. In our model we assume that the output of AIP and VIP are used by F4/F5 motor module (a part of F5 area) to make an action plan based on the object affordances, location and intention of the organism.

Inference: An intriguing set of multimodal neurons in the area F5 are called mirror neurons. These neurons become active when the organism executes a manipulation action as well when it observes a similar action being performed by a conspecific. Although the exact function of these neurons are not known, they seem to be related with prediction, e.g. forward modeling and intention estimation (Oztop, Kawato & Arbib, 2013), or inverse modeling supporting imitation, and possibly forming the basis of self and agency concepts (Hurley, 2008; Murata & Ishida, 2007).

A forward model is a computational process that represent the mapping between the inputs to the modeled system and the elicited responses or output. An inverse model does just the opposite; thus from a motor control point of view, if the mirror neuron responses are rich enough to replicate an observed act, we might consider their function as inverse modeling rather than forward modeling. In the current model, we consider a high level forward model that

performs effect predictions about the perceived environment. Without more digression on mirror neuron function, we take the liberty to associate F5 mirror module with intention inference function with the understanding that intention understanding can be sustained by both forward and inverse models with additional circuitry (Hurley, 2008; Oztop, Kawato & Arbib, 2013; Oztop, Wolpert & Kawato, 2005). This intention inference view is in fact compatible with the recent views on mirror neuron function (Iacoboni, Molnar-Szakacs, Gallese, Buccino, Mazziotta & Rizzolatti, 2005). *Motor Planner with Gating Mechanism:* F5 motor neurons are critical for hand control for manipulation; whereas F4 is involved in reaching. The location of these areas are high in the motor hierarchy (premotor cortex), as they project to primary motor cortex. So it is plausible to envision F4 and F5 as creating a motor plan that is executed by low level motor centers. In the model, the motor planner receives its input from the Prefrontal Cortex (PFC) through Thalamus, which converts the desires of an organisms (e.g. 'want to eat') into desired sensory changes that can be used by the motor planner. The Prefrontal Cortex is an area that is known to be involved in mentally demanding tasks and working memory. So we assume that the desires or high level goals of an organism are represented by PFC. The thalamus is a subcortical structure serving as a communication relay center between brain regions (sensory, motor and cortical) with modulatory and inhibitory mechanisms (Guillery & Sherman, 2002). In particular, the mediadorsal thalamus (MD) is a thalamic relay that contributes to cognitive processes such as learning and decision making (Mitchell, 2015). Furthermore, MD and some other thalamic nuclei project strongly to the cortical motor centers including the ventral premotor cortex (i.e. F4, F5) (Fang, Stepniewska

& Kaas, 2006). Considering the information relay function of the thalamus we propose that the thalamus undertakes the gating function under the control of Prefrontal Cortex, without specifying the exact nuclei within the thalamus. So we postulate that (1) during action observation, the propagation of F5 mirror output (predicted effect) to F5 motor is suppressed, i.e. not relayed to motor centers through the thalamus due to the inhibition from PFC, and (2) during action execution, PFC monitors the ongoing action to detect possible disturbances. It is known that impulsive imitation behavior can be caused by an impairment in PFC function, which modulates thalamus relay function (Guillery & Sherman, 2002). Compatible to this, in the proposed model we argue that infant PFC is responsible for the almost automatic (i.e. non-inhibited) altruistic helping behavior by applying executive control on the thalamus relay mechanisms. Initiation of such behavior is the result of automatic adoption of the predicted effect of an observed action as belonging to one's self. This happens in infants when social norms are not developed yet to inhibit this action. In adult humans the PFC normally inhibits imitation and altruistic behavior, and only let the sensorimotor system imitate or help on a voluntary and purposeful basis, except for the pathological cases.

3.2 Functional View of the Model

To be basis for an interactive robotic system the modules of the model are designed to fulfill functional requirements with real-time constraints although we keep the connection to the possible brain regions (Fig. 2). The model is required to operate in one of the two modes: `goal-achievement-mode` or `goal-inference-mode`. In the former the model executes actions to create a desired change, i.e. fulfill a goal, whereas in the latter the agent observes an ongoing action carried out by an actor by making predictions about the aim of the actor. The following main modules constitute our model:

- *Vision Module* captures the visual processing of the brain by performing early visual processing and computing spatial relations in the environment. The input of this module is an RGB-D image captured by a Kinect * sensor.
- *Affordance/Location Module* captures the function of the anterior intra-parietal and the ventral intraparietal area (AIP/VIP) of the primate brain. The input of this module is the processed object point cloud information from the *Vision Module*. As output, it computes shape, size, orientation and location of each detected object in the environment.
- *Hand Processing Module* captures the particular functionality of the superior temporal sulcus (STS), namely hand detection and trajectory formation by using the output of the *Vision Module*. This information is then used to find a list of objects that are potentially the targets of the observed action.
- *Intention Inference/Forward Modeling Module* is inspired by the possible functional properties of mirror neurons and related circuitry. In `goal-inference-mode`, this module is used to infer the goal of the another agent taking into account the affordances of the involved objects and the hand trajectory of the acting agent. In `goal-achievement-mode`, the acting agent can use this module to continuously predict results of its own actions to detect failures.
- *High Level Decision Making Module* represents certain functions of the prefrontal cortex, in particular, voluntary action generation and decision making. In this study the robot is not operated as an agent with its own goals, so this module is not active in the robot experiments; but, we still included it in the model for completeness.
- *Information Gating Module* captures the information relay function of the thalamus under the control of prefrontal cortex for directing and transforming the predicted effect of an observed action (output of the *Intention Inference Module*) or self-desired effect into a representation that can be used by the motor planning circuit.
- *Motor Planner Module* captures the mechanisms of reach and grasp served by the premotor cortex, in particular, motor neurons of F4 and F5 areas. The input to this module is the *desired effect* that needs to be created in the environment. Object affordances are again exploited to search for the motor signals that create movements for generating the desired effect.
- *Online Control Module* represents the basic functionality of the primary motor cortex (F1 area) and thus has the responsibility of executing actions instructed by the *Motor Planner Module*. In the robotic system

*<https://msdn.microsoft.com/en-us/library/hh438998.aspx> accessed: June 2018

this also includes the low level controllers driving the robot motors.

- Finally, *Body/Plant Module* represents simply the embodiment of the system which moves and interacts with the environment or other agents. For a human this is the body; for our implementation this the robotic system itself.

Most of the summarized modules above are utilized both during action execution and action observation. The common processing of the visual stimuli created by self and the others' actions gives rise to altruistic behavior as described in the next subsection.

3.3 Model Operation for the Altruistic Robot

The robot execution starts only after a valid unfulfilled goal is observed, therefore, we explain the operation of the goal inference first.

Goal-inference-mode Our robot observes the table environment continuously to detect a human manipulation action. Whenever a hand is detected, the robot engages in goal inference. To be concrete, the model adheres to the following execution flow: The *Vision Module* receives visual information from the environment and processes it to extract hand and object information. The object information is then processed in the *Affordance/Location Module* to retrieve action possibilities provided by all the objects based on their properties. Parallel to this, *Hand Processing Module* processes the hand trajectory and extracts the likelihood of each object to be the target of the acting human. The affordances and likelihoods of the objects are then used by the *Intention Inference/Forward Modeling Module* to compute the goal of the acting human. The inferred goal of the partner human is encoded as a predicted effect in the environment and forwarded to the *Gating Mechanism Module*. However action execution is not triggered unless this predicted effect is allowed to be relayed to the *Motor Planner Module* as a (self-) desired effect. This can be interpreted as the default behavior of the *Information Gating Module*: inhibiting the transformation of the predicted effect as the desired effect (see Fig. 2). When it is observed that the task is not being completed by the human, this inhibition is lifted and the *goal-achievement-mode* is activated. Therefore the suitable dis-inhibition mechanism automatically yields an altruistic behavior. In the current implementation, the system lifts the inhibition when it detects that the human hand is not making any progress towards completing the inferred task.

Goal-achievement-mode The robot in this mode executes actions to fulfill a desired effect (e.g. place object 1 on top of object 2) by actively interacting with the environment. The execution flow of the model in this mode is as follows: As in *goal-inference-mode*, the visual features and the object affordances are computed by the same modules. *Intention Inference/Forward Modeling Module* is used to make predictions based on object affordances and effects of robot's own actions.

The movement is planned by *Motor Planner Module* based on the inputs of desired effect from the *removalGating Mechanism Information Gating Module* and object affordances and location information from the *Affordance/Location Module*. The movement plan is sent to the *Online Control Module* that implements the plan by executing the action primitives in the right sequence. The *Online Control Module* in turn converts the primitive actions into low-level motor commands for the robotic control system which creates changes in the environment. In parallel to this, the perception modules of the model monitors the scene, and performs the functions they are responsible for completing the perception-action cycle.

4 Robotic Implementation

4.1 Robot Perception

The robot is equipped with a Kinect depth sensor, and has a built-in hand-vision coordination mapping so that the depth image from the sensor can be mapped to robot task space provided (Fig. 1). After the depth image is obtained, and the pixels outside the region of interest are filtered out, the objects are found by segmenting the remaining pixels using the Connected Component Labeling algorithm (Haralick & Shapiro, 1992). An object entering into the scene from the side is labeled as "human hand". Assuming the objects make continuous movements, the robot can track objects even when they are pushed or grasped by the human hand.

Feature Computation: For each detected object, the robot computes a number of features. These features are encoded in a 87 sized continuous vector composed of shape, dimension and local distance related properties of the object. Shape features are encoded as the distribution of local surface normal vectors on the object surface. Specifically, histograms of normal vectors along each axis, 8 bins each, are computed to form $3 \times 8 = 24$ sized feature vector. Dimension encodes the object size in 3 different axes. Local distance features encode the distribution of the local distance of all pixels to the neighboring pixels. For this purpose, for




Figure 3. An execution example from the subject experiments. The top sequence (from robot's perspective) and the bottom sequence (external camera from side) belong to the same execution. The green texts in the figure are the object numbers used to track the objects, and the affordances of the corresponding objects. *gr*, *re*, *in*, *st*, *pu*, *ro* respectively means *graspable*, *reachable*, *insertable*, *stackable*, *pushable* and *rollable*. There are five objects on top of the table and the subject aims to insert the orange toy inside the red basket. (a),(1): before task execution starts, (b): the robot detects the hand (cyan box) and the objects (red boxes), (2): the subject interacts with the first object, (c): the object affordances are predicted. (3): the human transfers the first object up to the reachable area for him, (d): the potential trajectories to the each object (green dashed arrows) and the actual trajectory, (4,5,6,e): the robot executes the inferred plan, takes the orange object from the hand and inserting it inside the basket successfully. Video url: <https://youtu.be/eEX4c14YEQI>

each pixel, distances to the neighboring pixels are computed along each 4 direction on the image. For each direction, a histogram of 15 bins with bin size of 0.5cm is created, obtaining a $4 \times 15 = 60$ sized vector.

Affordance Learning and Detection: Affordances are defined as the relations between objects, actions and effects following (Şahin, Çakmak, Doğar, Ugur & Üçoluk, 2007); and learned by training classifiers that predict the effect of actions given the object features. A separate classifier is trained for each action to predict the effect of that action on a perceived object. The 87 sized feature vector that is computed from the depth image of the corresponding object is used as the input to a Support Vector Machine (SVM) classifier. Real robot interactions are used for learning: the robot executes its actions on objects, observes the generated effects, and uses this experience to build classifiers. In this work, we assume that the effect categories are known in advance (see our previous work on how effect categories can be autonomously discovered by robot actions on single objects (Ugur, Oztop & Sahin, 2011) and on pairs of objects (Ugur, Nagai, Sahin & Oztop, 2015)).

4.2 Goal Inference

The robot uses two levels of prediction mechanisms for inferring the aim of an observed human. At the higher level, the most probable action that can be executed by the human partner is computed; whereas at the lower level, the resulting effect that would be created due to the execution of the anticipated action is predicted. We describe both mechanisms below.

4.2.1 Behavior Prediction: In order to predict the action of the acting agent, i.e. human partner, the robot searches over each (action, object) pair to find the most likely pair based on the hand trajectory observed so far and the object affordances. To do this, the system creates expected hand trajectories for each object on the table given the initial human hand position. Then these trajectories are compared against the observed actual trajectory of the hand via the Dynamic Time Warping (DTW) algorithm (Müller, 2007) to obtain the likelihood of the objects being the target of the ongoing human action. Next, actions that can be applied to the list of 'potential objects' are predicted by utilizing the object affordances. The potential object list is formed by filtering the objects by using a likelihood threshold based on the hand trajectory. This threshold is calculated according to $threshold = 1/n_{objs}$, where n_{objs} refers to the number of objects. The rest of the objects from the highest likelihood to the lowest are evaluated for the affordances available in that time instance for a reasonable interaction, that is an interaction that will cause a strongly predictable end scene in case of successful execution, by taking into account both of the objects and the hand position at that time. For example, if a hand holding a ball is moving towards an area with a basket and a lying cup, the prediction mechanism returns insertion into the basket, because putting a ball on top of a lying cup will result in an tumble effect that will cause a final scene not possible to predict, whereas, the insertion causes a more predictable final scene.

4.2.2 Effect Prediction: Once the action and the target of the action is determined by the behavior prediction mechanism described above, the robot computes the effect

(i.e. perceptual change) that would be generated by the completion of the ongoing action based on the affordances extracted for the target object. These effects are learned by the robot by applying the actions on the objects or object pairs during affordance learning phase. This information, as well as the previously experienced actions and their achieved goals, are stored in a structure for later processing. The effect prediction uses the predicted (object, action) pair and determines the likely change in the perception using object affordances. For example, when a cubic object is inferred to be put on another cubic object, stack effect is predicted to be observed. It is worth reminding that the stack effect is a label we use to describe a distributed sub-symbolic representation of the object properties.

When the predicted effect is not fulfilled for a fixed duration of time, e.g. if the hand is observed to not make any progress towards completing the task, altruistic behavior can be readily obtained by switching the robot from the `goal-inference-mode` to the `goal-achievement-mode`. After the `goal-achievement-mode` is engaged, the robot makes a plan based on its own action repertoire to complete the predicted effect on the predicted object(s).

4.3 Motion Planning

The robot aims to minimize the difference between the current and predicted situations. It achieves this by engaging its autonomous motor planning that selects and orders a set of primitive actions such as moving the arm to a given position, closing or opening the gripper, etc. In other words, there is no special *help* planning; the robot aims for full `goal-completion`, if it is allowed by the environment. An execution example from a subject experiment is provided in Fig. 3 where the robot steps in and plans its actions to achieve the inferred goal: one object being inside the other. If the full `goal-completion` is not afforded by the current context, then the robot aims to create a change in the environment so that the predicted behavior of the human becomes executable. In this example, placing the target container object to the proximity of the human hand may allow the human to act on it, and consequently help facilitate task completion.

5 Experiments

The experiments involve a dexterous dual hand manipulator robot implementing the sensorimotor mechanisms described in Section 4. The robotic platform, Baxter[†], is equipped with two 7 degrees of freedom arms and two parallel grippers.




Figure 4. The experimental setups for hand-trajectory-based target object prediction (Section 5.1.2) are shown. The hand in each experiment was instructed to move along the green dashed line starting from the green diamond mark. The possible end points are marked with coins.

A Kinect depth sensor that is placed on top of the robot serves as its means to perceive the environment (Fig. 1). The environment is composed of a black table with a set of object on it, where the Baxter and a seated human partner (not visible in the figure) perform various manipulation tasks on the objects.

The object manipulation and human-robot interaction scenarios are designed as follows. Given a set of objects on the table, the human partner is instructed to either get hold of a particular object or create a certain spatial configuration between two objects such that they become related with one of the predicates of next to, on top, or inside. The far side of the table is made unreachable to the human partners by adjusting their chair positions individually. Furthermore, the humans are not allowed to lean towards front. The robot is initially set to the `goal-inference-mode`. Once a hand is detected, it extracts object affordances and starts to continually observe the human action to infer the goal of the action. If the robot concludes that the inferred goal cannot be achieved by the human, the `goal-achievement-mode` is triggered. In this mode, the robot attempts to achieve the goal through its own actions and the learned knowledge of the object affordances.

With this interaction scenario in mind, the designed experiments can be grouped into two: *prediction evaluation experiments* (Section 5.1) and *interaction experiments* (Section 6.2). The first group evaluates the learning and prediction performance of the individual modules of the system; whereas the second group evaluates the system as a whole with subjects in human-robot interaction settings.

5.1 Prediction Evaluation Experiments

We designed experiments to evaluate the key components of our model to ensure that our system can be deployed as an interactive partner for a human. The experimental setups are detailed below, and the results of these experiments are provided in Section 6.1.

[†]<https://www.rethinkrobotics.com/baxter/> accessed: June 2018

5.1.1 Affordance learning and prediction: The aim of this experiment is to test the performance of the learning-based affordance prediction mechanism that uses the object features extracted from the point cloud data as input. For learning and test, we generated 417 scenes by using 28 different objects that were placed at various locations on the table. The objects were either plush vegetable toys, daily rigid objects such as boxes, balls, cups, or 3D printed symmetric objects such as cylinders, cubes, containers. The robot interacted with each object in each scene using its push and grasp actions, stored the effects of these interactions, and learned classifiers that predict effects of actions given object features. In order to test the performance of our affordance prediction models, and systematically analyze its generalization capability, this dataset was used to generate random n sized subsets from the full set of m objects for training, and $m - n$ sized subsets for testing. This process was repeated 100 times to assess the average generalization ability of our affordance prediction mechanism.

5.1.2 Trajectory based target prediction: This experiment is designed to evaluate the performance of the system in predicting the intended target object of the human subject by using only hand trajectory information. Identical objects that offer the same affordances were placed on the table. The task of the subject was to perform a reaching movement targeted at a pre-specified object, at a fixed position on the table. The subject was instructed to follow a pre-defined path, and then stop at one of the four end-points located 10, 20, 30, and 40 cm away from the target object (Fig. 4). The target object was placed on the table at the far left side of the subject who was instructed to move his hand in a diagonal path towards the target starting from the home position. Our target object prediction system was tested with 2, 4, and 6 objects on the far side of the table, where each object was horizontally separated by 30, 15 and 10 cm distances, respectively. For each configuration, the subject made 10 incomplete reaching movements as instructed. Therefore, $(4 + 4 + 4) \times 10 = 120$ reaching movements were performed in total.

5.1.3 Trajectory and affordance based target prediction: To investigate the efficacy of exploiting both hand trajectory and object affordance information in inferring the goal of an observed action, the first author was given the task to bring a small ball towards a large insertable but unreachable open box (target object) with the intention of placing the object inside the box. The robot was required to predict the target object by using the incomplete action trajectory of the subject together with the affordance information pertaining to the objects in the scene. The number of




Figure 5. Trajectory and affordance based prediction experiment setup from the robot's perspective. The prediction capability of the robot was tested in configurations with gradually higher level of complexities. (a) and (b) show the initial and final configurations, respectively. New objects are added at each configuration and tested 10 times from different start points. The objects are numbered with the order of their inclusion to the setup: an open box, a toy paprika, a reversed cup, a basket, a medicine box and a cup.

objects is gradually increased to generate progressively more challenging configurations. Fig. 5 (a) shows the initial configuration of the objects (a ball in reachable area and a larger box in unreachable area) from the robot's perspective prior to the action execution of the subject. The first scenario includes no distractors so it is rather straight forward for the robot to predict the goal of the human partner who grasps the ball and transfers it towards the only remaining object. The difficulty in inferring the target object was gradually increased by adding an object to the target area at a time as shown in Fig. 5(b). The prediction capabilities of the system was tested in each of these six gradually more complex configurations. For all the configurations the subject repeated his insertion attempt 10 times. In each trial, the subject picked up the ball from a different position and carried it towards the target box until his reach limits. At this point, the system made two predictions, one with the hand trajectory, and the other with the perceived object affordances together with the hand trajectory.

5.2 Human-Robot Interaction Experiments

We designed a number of experiments to evaluate the complete system that integrates mechanisms of affordance computation, target object prediction, behavior and effect prediction, planning and execution in human-robot interaction settings. The experimental setups are detailed below, and the results of these experiments are provided in Section 6.2. In the first set of experiments, the performance of the system was tested on randomly generated scenes and tasks. The first author who designed and implemented the robotic system took the subject role for this purpose. In the second set of experiments, our system was tested with naïve subjects who had no information about the model, the underlying implemented system, perception mechanisms of the robot or the reasoning behind the design of the experiments.




Figure 6. Two random generated scenes are on the left and the corresponding experimental setups are on the right.

5.2.1 HRI experiments in randomly generated scenes:

The aim of this experiment is to evaluate the performance of our integrated system in helping a subject who is familiar with the robot and the system. To this end, we used a random scene generator algorithm that selects 6 random objects out of 20 objects (selected for this experiment) and instructs us to place them on the table at randomly selected locations. This algorithm picks at least three locations that are reachable by the robot and at least two reachable by the human, and distributes the chosen objects. The algorithm then selects one object from each side, and designates them as objects to be interacted with. The subject then voluntarily decides a suitable goal using these objects, the goal being either bringing them on top of each other, or inserting one inside the other. Example scenes generated by the random scene generator and the corresponding real setups are provided in Fig. 6.

The subject attempted to clear 24 randomly generated interaction tasks in this experiment. In each configuration, at least one object was unreachable by the subject and reachable by the robot, therefore the subject was not able to complete the task by himself and required the altruistic helping behavior of the robot for task completion.

5.2.2 HRI experiments with naïve subjects: The aim of this set of experiments is to investigate the suitability of our system as an altruistic partner for naïve subjects (subjects with no prior information on our robot, model, implementation or the experimental design - See Appendix A for details) in real world experimental tasks. The subjects were informed about the aim of the experiment as to test the robot in human robot collaboration with naïve users, how their data will be kept and used including the video recording. Each subject was given 10 tasks. Before a subject started his/her experiment, the seat was positioned according to the physical specifications of the subject. Some capabilities of the robot were not shown to the subjects in order to later test if the subjects would expect help from the robot with actions that were not observed before. For this, no

insertion action was shown to the first half of the subjects, and no stacking action was shown to the second half. Before watching the video the subjects were warned to pay attention to how the robot collaborates with the humans, how the humans receive help from the robot, and how the humans act to collaborate with the robot. Instead of explicitly describing the subjects how the robot infers actions and steps in for help, we expected that the subjects could infer such information from the videos if our system and setting were natural. The object configurations that the subjects were engaged in the experiments are shown in Fig. 7.

The subjects were given 10 tasks in total, in which they were free to choose the means, i.e. how to complete the task. The first two tasks involved grasping unreachable objects: in task 1 there was only one object on the table, whereas in task 2 the subject was free to choose to grasp one of the three objects. In tasks 3-5, the subject was required to grasp a small reachable object, and insert it in or stack it on one of the unreachable objects that offered different affordances such as insertability, stackability, and rollability. In task 3, for example, the toy tomato was required to be inserted into the box. Depending on the subject behavior, robot might act differently in the same task: If the subject brought the tomato to a reachable area for the robot, the robot would have grasped the tomato and inserted it into the box. Otherwise, if the subject held the tomato in an unreachable area for the robot, the robot would have pushed the box towards the subject in order to create an environment as close as to the inferred goal. Tasks 6-7 were sequential where the subject was required to select one of the two reachable small objects, decide to insert/stack it in/on one of the unreachable three objects that were either insertable or stackable, grasp the remaining reachable small object, and decide to insert it in or stack it on one of the unreachable objects again. Tasks 8-10 were also composed of subtasks similar to the previous one, with more objects to select from the reachable area for the subject, and with unreachable objects that offered affordances difficult to utilize. For example, there was a fluffy soft object which had a round shape, affording neither stackability nor insertability. Tasks with sequential subtasks were used to give subjects freedom to create a dynamic environmental task context. In the sequential tasks, previous actions could cause changes in affordances of some objects(see stable vs. variable affordances discussion in (Borghi & Riggio, 2015)). For example, in the particular implementation of task 9, if the subject did choose to stack the carrot on top of the box, the box would have become unstackable after this action, because the carrot on top of the box did not afford stackability.

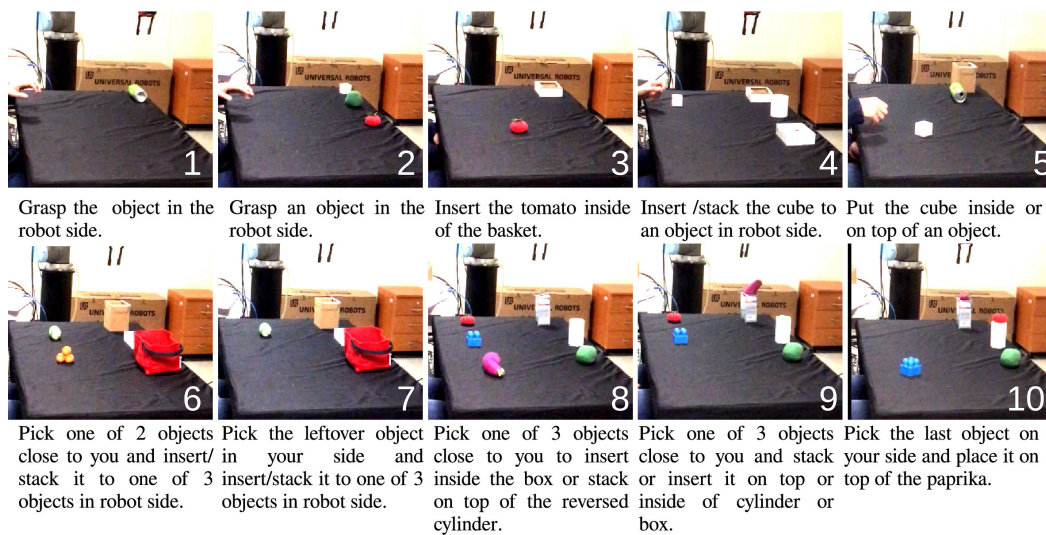


Figure 7. The provided tasks and examples of initial configurations as snapshots taken from the naïve subject experiments.

6 Results

6.1 Prediction Evaluation Experiment Results

6.1.1 Affordance learning and prediction: In this experiment, our dataset was separated into different sized training and test sets 100 times with shuffled data. In order to avoid overfitting, different scenes from the same object were not shared between the training and test sets (see Section 5.1.1 for details). Our motivation is to evaluate the performance of affordance predictors that use visual features of the objects. Fig. 8 provides the mean accuracy of each affordance predictor that used gradually increasing ratios of the data for training.

Note that the error bars indicate the 95% confidence intervals for the affordances. The accuracy is low when the training size is small as expected. When 90% objects were used in training, the mean accuracy is higher than 90%. Insertability and stackability affordances were learned with high success rates compared to the others because hollow and flat top parts were easier to detect. Furthermore, pushability and rollability affordances depend on also the material properties of the objects, which were not encoded in the feature vectors.

6.1.2 Trajectory based target prediction: In this experiment, different numbers of identical objects were placed on the target area, partial reaching attempts with different length trajectories were performed towards one of those objects, and these reach attempts were repeated 10 times for each configuration. We evaluated the performance of the target object predictors that only use hand trajectory information of the subject in the experiment setup described in Section 5.1.2. Fig. 9 provides the results where the black, dark gray and light gray bars show the results of reaching

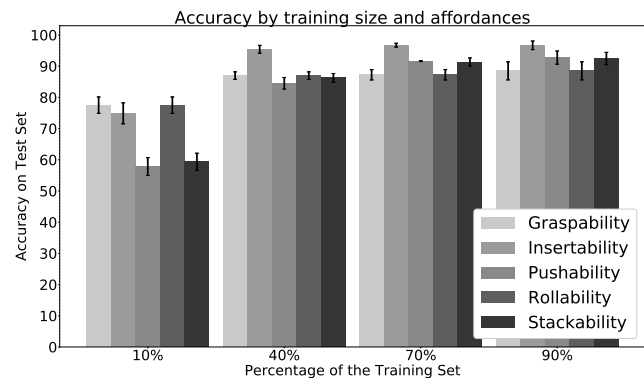


Figure 8. The accuracy of affordance predictors that use different ratios of data for training. Each training/test is repeated 100 times shuffling the dataset, therefore a distribution is presented where the bars and the lines correspond to the means and 95% confidence intervals, respectively.

attempts in settings with 2, 4 and 6 objects, respectively. In 2-object case, the target prediction was successful almost all the times unless the hand stopped 40 cm away from the target object. In 4-object case, the target prediction was successful unless the hand stopped 30 cm away. In the most cluttered setting, i.e. with 6 objects, the target prediction could be successful only when the hand came very close to the target object. These results show that, when the objects in the target area are few and separated sufficiently, the trajectory is a good indicator for target object prediction. However, when the number of objects increase and the objects start to get closer to each other, the prediction is successful only when the hand becomes close to the proximity of the target object.

These results show the limits of our hand-trajectory-based predictor: the robot cannot predict the target object when the hand is far away and the target object is closely surrounded by others.

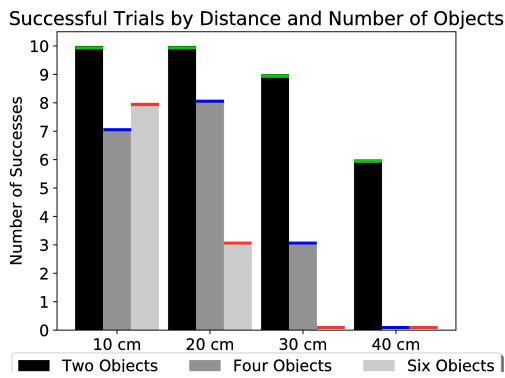


Figure 9. The performance of target object prediction only from hand trajectory. 10 trials were performed for each configuration: 2, 4 or 6 objects were placed to the target area and subject's hand halted at 10, 20, 30 and 40 cm distance from the target object.

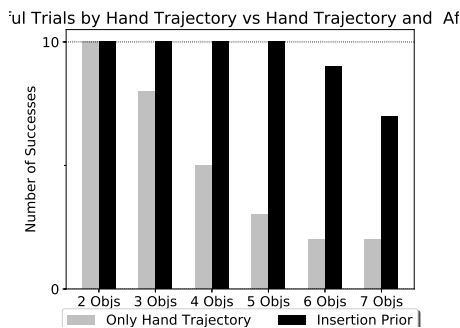


Figure 10. Gray and black bars show the success rates of the goal prediction mechanism of the system without and with exploiting affordances, respectively

6.1.3 Trajectory and affordance based target prediction:

In this experiment, we evaluated the benefit of using object affordances as an additional source of information to hand trajectory data in predicting target objects of incomplete actions in the constrained experimental setup described in Section 5.1.3. Recall that the number of the objects on the target area was gradually increased and the subject performed a partial reach action 10 times in each configuration with the aim of inserting the object held into a target container. At the end of human action, the system made two predictions, with and without perceived affordances.

Fig. 10 provides the performances of the predictors that do and do not utilize affordances. These results show that affordances were instrumental in more consistently and better detecting target objects when an affordance prior (insertion) is used: the prediction performance degraded gracefully when the system exploited affordance information in contrast to when only the hand trajectory information was available.

6.2 Human-Robot Interaction Experiment Results

The human-robot interactions were recorded by two cameras and stored along with the execution logs during experiments. The intended goal of the human partner could be obtained from the execution logs for the experiment that used randomly generated scenes (since the goal was instructed to the human), or extracted from the questionnaires filled by the subjects in the naïve subject experiments. Object selection, affordance prediction and plan generation related errors were identified by inspecting the generated plans stored in the execution logs. Finally, the videos were inspected to assess whether each execution was successful or not, and the cause of the problem, in case of a failure.

6.2.1 HRI experiments in randomly generated scenes:

In these experiments, we evaluated the performance of the integrated system in randomly generated configurations and two object manipulation tasks. 24 random scenes were generated with 6 objects and the first author attempted to achieve randomly generated tasks that can only be achieved with the help of the robot. These experiments required both inference of the goal of the subject and physical help of the robot. The results indicate that although the goal inference was correct in about 96% of the trials, the requested task could be completed by the subjects with the help of the robot in about 71% of the trials due to execution related errors (see Table 1). In these experimental trials, as the goal was beyond the subject's capabilities in all the tasks, the robot either finished the incomplete goal directly (e.g. by grasping an object from subjects' hand and releasing it inside a box), or assisted the subject for the completion of the goal (e.g. by pushing the box towards the subject so that the subject can drop the object directly inside the box). Fig. 11 shows snapshots from the experiments where the robot directly completed the goal Fig. 11(a) or assisted the subject for goal completion Fig. 11(b).

With this finding, we conclude that the proposed model and its implementation on our robot can be used in real-time, even though a near-perfect performance could not be obtained. Therefore, we moved to experiments with naïve subjects to assess the suitability of our system as an altruistic partner in natural settings.

6.2.2 HRI experiments with naïve subjects: We evaluated the performance of the integrated system using 10 naïve subjects in 10 tasks described in Section 5.2.2. Each subject was assigned the same set of tasks that required robot's help in their attempts of reaching objects, inserting objects inside others, or stacking objects on top each other.

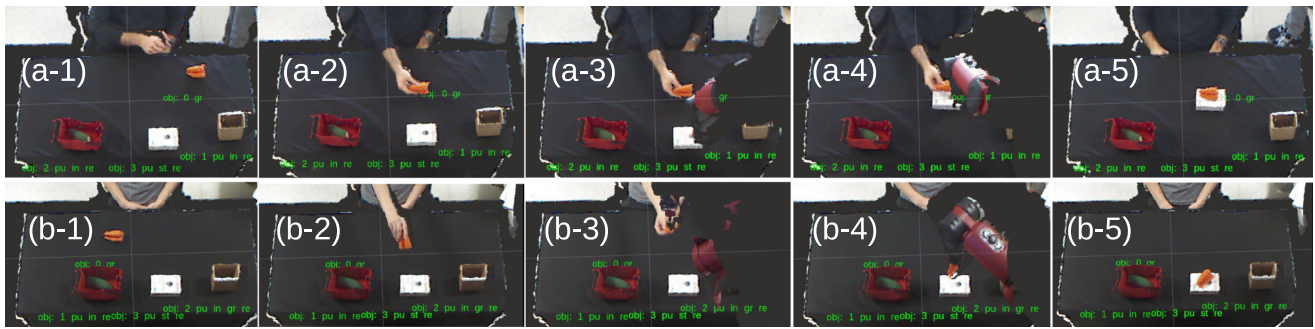


Figure 11. There are two execution examples in the figure taken from different naïve subject experiments (a-1 to a-5 from subject 5 and b-1 to b-5 subject 9 both in round 7). The aim of both subjects are to stack the carrot on top of the object in the middle. In the upper case since the robot detected the carrot as unreachable, it planned to push the object in the middle to the subject for enabling him to complete the task. Whereas, in the lower case the robot detected the carrot as reachable and stacked it on top of the object by itself.

Table 1. The results of the HRI experiments in randomly generated configurations

Result	Rates (randomly generated scenes)
Successful execution	70.84%
Failure in affordance computation	4.16%
Failure in predicting the 1st obj	0.00%
Failure in predicting the 2nd obj	0.00%
Failure in planning	0.00%
Failure in execution	25.0%

After each task was attempted, the subjects were given a form that includes 3 questions about their intended task, what type of actions the robot took, and their grade from 1 to 5 on the success of the robot to help them complete their intended task (see Appendix B, Q1.1-1.3). The grades given by the subjects to the third question are summarized in Table 2. In 55% of all tasks, the subjects graded the robot with highest two grades (4 or above), therefore, it can be argued that the subjects were satisfied with the altruistic helping behavior of the robot to complete their intended task in most of the time during the experiments. 15% of the tasks, the subjects did choose an average grade of 3. In the rest of the tasks, the subjects graded the system with low grades, i.e. with 2 or 1. The subjects were asked to fill another form (see Appendix B, Q2.1-2.10) that was provided after they completed all the experiments, which probed for their overall experience with the robot. Two questions related to the success of the robot in inferring and helping the human (Q2.1 and Q2.3) received an average score of 3.90 ± 0.91 , and two questions related to the realism of the system in speed and behavior (Q2.4 and Q2.5) received an average score of 2.75 ± 0.91 . The subjects found the system successful, yet averagely realistic.

The ratings of the naïve subjects were subjective and dependent on their level of understanding and engagement in the tasks. To have a less subjective assessment, the main author also evaluated the results of the naïve subject experiments by reviewing the questionnaires, the execution logs and the videos. The results indicate that 68 of 100 trials

Table 2. The grades given by the subjects for their evaluation on the success of the altruistic helping behavior of the robot to complete their aim after each trial are provided. T1 and S1 correspond to task 1 and subject 1, TM and TV correspond to mean and standard deviation of the grades given to the corresponding task, and SM and SV refer to the mean and standard deviation of the grades given by the corresponding subject.

	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	TM	TV
T1	5	5	5	5	4	5	3	4	5	4	4.5	0.71
T2	4	5	5	4	1	1	4	3	5	5	3.7	1.57
T3	4	4	5	5	5	4	4	5	3	1	4.0	1.25
T4	2	5	1	5	1	1	2	5	5	1	2.8	1.93
T5	3	2	5	5	1	5	4	2	3	2	3.2	1.48
T6	1	5	4	5	4	3	2	3	4	3	3.4	1.26
T7	5	4	2	3	5	3	2	4	5	1	3.4	1.43
T8	3	1	1	2	1	3	1	4	5	1	2.2	1.48
T9	2	5	5	5	4.5	5	3	2	5	4	4.05	1.26
T10	4	1	5	4	1	1	4	2	3	3	2.8	1.48
SM	3.3	3.7	3.8	4.3	2.7	3.1	2.9	3.4	4.3	2.5	3.4	
SV	1.3	1.7	1.7	1.0	1.8	1.6	1.1	1.1	0.9	1.5		

were successfully completed. According to the evaluation, the robot never failed to infer the first target object that the naïve subject intended to grasp. Given the inferred goals, the robot never made an incorrect plan as well. Evaluations have shown that the most common cause of the task completion failure was due to the object manipulation problems (19%). Furthermore, in 8% of the interactions, the execution failed because of a collaboration failure between the robot and the human in handing an object. Push execution was another problematic case with 5% of the interactions, some thin objects when pushed along their major axis escaped from the gripper and some light objects toppled over the table; and did not move as expected. While the former case can be improved with better close-loop controllers, the latter case is challenging as the weight of the object is difficult to perceive from vision.

Table 3. The results of the HRI experiments with the naïve subjects experiments in the pre-defined configurations.

Result	Rates (volunteer naïve subjects)
Successful execution	68.0%
Failure in affordance computation	1.00%
Failure in predicting the 1st obj	0.00%
Failure in predicting the 2nd obj	4.00%
Failure in planning	0.00%
Failure in execution	19.0%
Failure in human collaboration	8.0%

6.2.3 Discussion We observed that the success criteria of the naïve subjects were diverse and sometimes counter-intuitive especially when we consider that their answers should have reflected the robot’s capability in completing the tasks. For example, in a scenario where the subject tried to insert a grasped object into an unreachable box, she expected the robot to pick up the object from her hand that was stretched towards the robot even if her hand was unreachable for the robot. In this case, the robot pushed the box towards the subject, which enabled the subject to drop the object into the box; however the subject still annotated this behavior as a failure. This example suggests that humans might see the robots failing even if the robot successfully helps them in clearing the given tasks, probably because they are not aware of the capabilities of the robots: reachability region of the robot in this particular example. Yet, in other examples, although the robot was incapable of completing the tasks, the subjects gave a high grade if the action plan of the robot was correct from their perspectives. In a setting where the subject was required to stack an object over an unreachable one, the robot was rated as successful when it picked up the object from her hand and released it over the target, independent of the result of the stacking action. In other words, even when the objects tumbled of the stack as the result of the release action of the robot and the task was not accomplished as specified, the subject gave a high grade to the robot.

The lack of a good understanding of robot capabilities not only resulted in dissatisfaction for the subjects, but also significantly affected the physical interaction performance. Related to affordance detection capability, our system computes object affordances from the visual perception of the immediate environment. For example, the robot perceives a lying cylinder on the table as a rollable, non-pushable and non-stackable object; and reasons over these affordances throughout the interaction trial. From a human perspective, on the other hand, the same object can be exploited as a stackable and pushable object when a simple ‘rotate’ action is applied. Therefore, collaboration attempts fail when humans expect the robot to infer or execute such ‘obvious’ actions that were not included in the action repertoire of the

robot and that were not considered to change to affordances by the robot. The low grades obtained in Task 8 are due to the mismatch between insertion affordance perception of the robot and the subjects. From the perspective of the robot, the robot learned that the objects with large surface areas are not insertable to smaller gaps. However, from the perspective of the human subjects, the soft toy eggplant can be picked up, rotated and squeezed into the box. As the robot decides object affordances from their perception on the table and the robot does not have rotate or squeeze primitives in its action repertoire, it failed to help to the subjects when the subject chooses insertion with the eggplant. In Task 9, the subjects generally selected an action that involves the toy tomato, where the system obtained high grades by stacking the tomato on the cylinder or the box. The object detection capabilities of the robot are limited currently with the available point cloud information in our system. Therefore, an accurate position perception of the objects that are in the hand of people is not possible; and the robot requires some help from human while picking up the object from the hand. Subject 7, for example, was either not aware of this fact or found this unacceptable, and therefore rejected to move the object in her hand towards the robot gripper while the robot tried to pick up the object. A final observation was that even with the prior exposure of the video recoding, a subject (subject 10) failed to understand/remember how two object interaction tasks were achieved by the robot. The subject 10 could not achieve the relatively simple third and fourth tasks, and failed in the rest of the tasks as well.

A final remark on naïve subject experiments is that, when required, the subjects were able to predict that the robot could engage with actions that were not previously shown to them in the introduction videos. For example, even if the first group of subjects were not shown insert action of the robot, they could infer that the robot could help them if the task requires insertion of the objects.

To sum up, we provided video recordings of human-robot interactions to communicate the capabilities of the robot to naïve subjects with minimal explicit prior information about the inference and execution mechanisms of the robot and about the aim of our experiments. Still, our observations suggest that longer physical interaction experience with robots is required in such HRI settings as otherwise naïve users cannot develop a theory of mind for the autonomous robots, and cannot engage in effective collaboration with them.

7 Conclusion

In this paper a biologically inspired model is proposed to explain the mechanism behind the altruistic behavior observed in young infants. The proposal suggests that the altruistic behavior is not necessarily a result of complicated cognitive processes but can emerge through basic sensory processing. To leverage the plausibility of the model, it was implemented on a physical robot and evaluated in five different experimental setups, evaluating the components of the system one by one, and verifying the integrated system in human-robot interaction scenarios. The experiments in which the first author acted as the subject were used for verification and testing of the technical aspects of our proposed system. While some of the experiments (behavior prediction with/without affordance) checked whether particular modules had significant effect on system performance, the others (behavior prediction in densely cluttered environments or robot action selection in randomly generated scenes) were used as stress tests to assess the performance of the robot in difficult configurations. These experiments were designed to be significantly more difficult than the naïve subject experiments, and shows the limits of our system rather than how humans act or perform with it.

It was found that the affordances that involve actions with single and paired objects (e.g. graspability and insertability) could be predicted with high accuracy respectively based on their visual features but the inability of detecting material properties of the objects did not allow close-to-perfect affordance prediction. This confirmed that some affordances depend on the material properties, and therefore perception systems that do not perceive such properties are expected to make mistakes in detecting affordances. While advanced machine learning methods and large datasets can be used to discriminate tactile properties from haptic and visual data (Gao, Hendricks, Kuchenbecker & Darrell, 2016), integration of such detailed and specific methods is out of the scope of this paper.

The learned object affordances, in turn, was shown to increase the performance of inferring the goals of observed actions of others especially if the target objects of those actions are located in cluttered environments. If the hand trajectory is used without affordance information, the goal inference performed well only if the scene is not cluttered or the hand is in the vicinity of the target object.

Our HRI experiments in randomly generated challenging environments and tasks verified the integrated altruistic robotic model in real world. We analyzed the reasons of failures systematically and concluded that while our system

performed well in affordance detection, goal prediction and action planning for help, the robot's helping capability was significantly affected from the difficulties in physical interaction with such a high variety of objects without any opportunity of taking haptic information into account during the control of the simple parallel gripper. Next, we designed HRI experiments with naïve subjects who were not provided any information about robot's behavior and underlying mechanisms except a very short video recording from subject - robot experiments. The naïve subjects were given 10 different tasks which required robot's assistance in different forms, and we observed that the robot could successfully infer the goals of the naïve subjects and make correct plans to complete their incomplete action executions, although it suffered from problems in actually execution of actions largely due to manipulation problems again and in some cases due to subjects' inability to understand how to physically collaborate with the robot. The system should be improved with advanced closed-loop manipulation controllers that exploit haptic information and feedback for better handling of objects. In particular, the adoption of anthropomorphic dexterous hands for the robot end-effector would eliminate grasping errors, widen the objects that can be worked with, and allow a more intuitive collaboration. Another issue is that the speed limitation of Baxter was reported as an hindrance for the naturalness of the system, and thus should be thought upon. Never the less, our general conclusion from HRI experiments is that our altruistic robotic model could successfully infer the intended goals of actions and engage in interaction for completing the unfulfilled goals with correct actions. Brain-inspired dual-use of sensorimotor components in inferring others' actions and achieving self-goals were shown to be effective in the real world especially when learned affordances were exploited. Relied on the error minimization mechanism, altruistic behavior was shown to be the result of basic sensorimotor processes rather than deliberate cognitive processes in our robotic model and implementation. However towards a more natural human-robot collaboration and higher performing altruistic behaviors, the action representation should be improved with advanced closed-loop controllers and with haptic feedback allowing faster and more reliable action execution.

Our system can complete the unfulfilled goals of human manipulation actions; however this is only one type of human-robot collaboration. For example, a human might require robots collaboration in tasks that are beyond his/her capabilities or might benefit from the help of the robot to speed-up goal achievement. The former case can be achieved

by our framework for certain manipulation tasks, however reasoning about tasks that require for example physical collaboration such as carrying heavy objects is beyond the capabilities of our system as predicting effects of joint action executions is not addressed in our work. For speeding goal achievement, the robot would need to predict higher level goals that might require execution of sequence of actions on several objects. For instance, in a table tidying up scenario, the robot would need to infer the general goal by combining its observation of human actions that are of the form of pick-up and place. In our current design, the goals are associated to particular objects that are involved in human actions. Therefore, the goal is required to be detached from the particular manipulated objects, and represented in a more abstract level to achieve abstract goal satisfaction. We believe that our model is suitable for such extension, however deciding the right level of abstraction or the right goal from a set of possibilities is an open problem on its own (Argall, Chernova, Veloso & Browning, 2009).

As a future work, we also would like to work on more realistic human-robot collaboration settings that are continuous rather than episodic and that include inferring and planning multi-step actions with chained effect predictions in object affordances. Finally, the more biologically realistic implementation of model components should be considered. This way, the results may be used to generate predictions at finer granularity for developmental neuroscience and psychology.

Acknowledgements

This work was supported by the European Unions Horizon 2020 research and innovation programme under grant agreement no. 731761, IMAGINE; JST CREST “Cognitive Mirroring: Assisting people with developmental disorders by means of self-understanding and social sharing of cognitive processes” [grant number: JPMJCR16E2]; Bogazici Research Fund (BAP) project IMAGINE-COG++ [grant number: 18A01P5].

References

- Arbib, M. A. (1981). Perceptual structures and distributed motor control. *Comprehensive Physiology*.
- Arbib, M. A., Erdi, P. & Szentagothai, J. (1998). *Neural organization: Structure, function, and dynamics*. MIT press.
- Argall, B. D., Chernova, S., Veloso, M. & Browning, B. (2009). A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5), 469–483.
- Baldwin, D. A., Baird, J. A., Saylor, M. M. & Clark, M. A. (2001). Infants parse dynamic action. *Child Development*, 72(3), 708–717.
- Baraglia, J., Nagai, Y. & Asada, M. (2016). Emergence of altruistic behavior through the minimization of prediction error. *IEEE Transactions on Cognitive and Developmental Systems*, 8(3), 141–151.
- Borghi, A. M. & Riggio, L. (2015). Stable and variable affordances are both automatic and flexible. *Frontiers in human neuroscience*, 9, 351.
- Şahin, E., Çakmak, M., Doğar, M. R., Uğur, E. & Üçoluk, G. (2007). To afford or not to afford: A new formalization of affordances toward affordance-based robot control. *Adaptive Behavior*, 15(4), 447–472.
- Dutta, V. & Zielinska, T. (2018). Predicting human actions taking into account object affordances. *Journal of Intelligent & Robotic Systems*.
- Gao, Y., Hendricks, L. A., Kuchenbecker, K. J. & Darrell, T. (2016). Deep learning for tactile understanding from visual and haptic data. In *2016 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 536–543).
- Guillery, R. W. & Sherman, S. M. (2002). Thalamic relay functions and their role in corticocortical communication: generalizations from the visual system. *Neuron*, 33(2), 163–75.
- Haralick, R. M. & Shapiro, L. G. (1992). *Computer and Robot Vision, Volume I*. Addison-Wesley.
- Hurley, S. (2008). The shared circuits model (scm): how control, mirroring, and simulation can enable imitation, deliberation, and mindreading. *Behav Brain Sci*, 31(1), 1–22; discussion 22–58.
- Iacoboni, M., Molnar-Szakacs, I., Gallese, V., Buccino, G., Mazziotta, J. C. & Rizzolatti, G. (2005). Grasping the intentions of others with one’s own mirror neuron system. *PLoS Biology*, 3(3), e79.
- Jamone, L., Uğur, E., Cangelosi, A., Fadiga, L., Bernardino, A., Piater, J. & Santos-Victor, J. (2016). Affordances in psychology, neuroscience and robotics: a survey. *IEEE Transactions on Cognitive and Developmental Systems*.
- Kenward, B. & Gredebek, G. (2013). Correction: Infants help a non-human agent. *PLOS ONE*, 8(11).
- Koppula, H. S., Gupta, R. & Saxena, A. (2013). Learning human activities and object affordances from RGB-D videos. *International Journal of Robotics Research (IJRR)*, 32(8), 951–970.
- Luppino, G., Murata, A., Govoni, P. & Matelli, M. (1999). Largely segregated parietofrontal connections linking rostral intraparietal cortex (areas aip and vip) and the

- ventral premotor cortex (areas f5 and f4). *Exp Brain Res*, 128(1-2), 181–7.
- Mitchell, A. S. (2015). The mediodorsal thalamus as a higher order thalamic relay nucleus important for learning and decision-making. *Neurosci Biobehav Rev*, 54, 76–88.
- Müller, M. (2007). Dynamic time warping. *Information retrieval for music and motion* (pp. 69–84).
- Murata, A. & Ishida, H. (2007). *Representation of bodily self in the multimodal parieto-premotor network* (pp. 151–176). Springer.
- Nelissen, K., Borra, E., Gerbella, M., Rozzi, S., Luppino, G., Vanduffel, W., Rizzolatti, G. & Orban, G. A. (2011). Action observation circuits in the macaque monkey cortex. *J Neurosci*, 31(10), 3743–56.
- Oztop, E. & Arbib, M. A. (2002). Schema design and implementation of the grasp-related mirror neuron system. *Biological Cybernetics*, 87(2), 116–140.
- Oztop, E., Imamizu, H., Cheng, G. & Kawato, M. (2006). A computational model of anterior intraparietal (AIP) neurons. *Neurocomputing*, 69, 1354–1361.
- Oztop, E., Kawato, M. & Arbib, M. A. (2013). Mirror neurons: Functions, mechanisms and models. *Neuroscience Letters*, 540, 43–55.
- Oztop, E., Wolpert, D. & Kawato, M. (2005). Mental state inference using visual control parameters. *Cognitive Brain Research*, 22(2), 129–151.
- Provasi, J., Dubon, C. D. & Bloch, H. (2001). Do 9- and 12-month-olds learn means-ends relation by observing? *Infant Behavior and Development*, 24(2), 195 – 213.
- Sakata, H., Taira, M., Kusunoki, M., Murata, A. & Tanaka, Y. (1997). The tins lecture - the parietal association cortex in depth perception and visual control of hand action. *Trends in Neurosciences*, 20(8), 350–357.
- Thill, S., Caligiore, D., Borghi, A. M., Ziemke, T. & Baldassarre, G. (2013). Theories and computational models of affordance and mirror systems: An integrative review. *Neuroscience and Biobehavioral Reviews*, 37, 491–521.
- Tomasello, M., Carpenter, M., Call, J., Behne, T. & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28(5), 675–691.
- Ugur, E., Nagai, Y., Sahin, E. & Oztop, E. (2015). Staged development of robot skills: Behavior formation, affordance learning and imitation with motionese. *Ieee Transactions on Autonomous Mental Development*, 7(2), 119–139.
- Ugur, E., Oztop, E. & Sahin, E. (2011). Goal emulation and planning in perceptual space using learned affordances. *Robotics and Autonomous Systems*, 59(7-8), 580–595.
- Warneken, F., Chen, F. & Tomasello, M. (2006). Cooperative activities in young children and chimpanzees. *Child development*, 77(3), 640–663.
- Warneken, F. & Tomasello, M. (2006). Altruistic helping in human infants and young chimpanzees. *Science*, 311(5765), 1301–1303.
- Warneken, F. & Tomasello, M. (2007). Helping and cooperation at 14 months of age. *Infancy*, 11(3), 271–294.
- Warneken, F. & Tomasello, M. (2009). The roots of human altruism. *British Journal of Psychology*, 100(3), 455–471.
- Zanchettin, A. M. & Rocco, P. (2017). Probabilistic inference of human arm reaching target for effective human-robot collaboration. In *Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on* (pp. 6595–6600).
- Zech, P., Haller, S., Lakani, S. R., Ridge, B., Ugur, E. & Piater, J. (2017). Computational models of affordance in robotics: a taxonomy and systematic classification. *Adaptive Behavior*, 25(5), 235–271.

A Naïve subject details

Volunteer subjects were recruited through various university mailing-lists. The students who study robotics and/or have information on this particular research topic were not allowed to participate. The subjects were composed of 1 female and 9 males aged 24–28 (one with Medicine, six with Engineering, and three with Social Sciences backgrounds). None of the subjects had any prior first-hand experience with a manipulator robot.

B Questionnaires for the naïve subjects

The questions given to the naïve subjects after each of their attempted task are as follows (1-lowest, 5-highest):

- Q1.1. What was your intended action?
- Q1.2. What was the action of the robot?
- Q1.3. How successful was the robot in helping your task?

The questions given to the naïve subjects at the end of the experiment are as follows (1-lowest, 5-highest):

- Q2.1. How well did the helping behavior of the robot match with your expectations?
- Q2.2. How much were you surprised with the solutions of the helping robot?

- Q2.3. How well the robot predicted the aim of your actions?
- Q2.4. How realistic was the helping behavior of the robot?
- Q2.5. Please grade the speed of the reaction of the robot.
- Q2.6. What type of situations did you have difficulties?
- Q2.7. What type of situations did you like most?
- Q2.8. Would you like to have the help of such a robot in your daily life ?
- Q2.9. What can be done on the robot to make it more helpful and more realistic?
- Q2.10. Any other additional comments?

About the Authors



Mert Imre obtained a Bachelor Degree in Computer Engineering from Bogazici University and is pursuing Master's Degree in the same department. He is an Active Member of the Cognition, Learning and Robotics Research Group in Bogazici University. His research is focused on Humanoid Robots, Cognitive and Developmental Robotics, Artificial Intelligence, Human-Robot Interaction, Robot Learning.



Erhan Oztop earned his Ph.D. at the University of Southern California in 2002. In the same year, he joined the Computational Neuroscience Laboratories at the Advanced Telecommunications Research Institute International, (ATR) in Japan. There he worked as a researcher and later a senior research and group leader where he also served as vice department head for two research groups and held visiting associate professor position at Osaka University (2002-2011). Currently, he is a faculty member and the chair of the Computer Science Department at Ozyegin University, Istanbul. His research involves computational study of intelligent behavior, human-in-the loop systems, computational neuroscience, machine learning, cognitive and developmental robotics.



Yukie Nagai received the masters degree from Aoyama Gakuin University, Tokyo, Japan, in 1999, and the Ph.D. degree from Osaka University, Osaka, Japan, in 2004, both in engineering. She was a Post-Doctoral Researcher with the National Institute of Information and Communications Technology (NICT), Kyoto, Japan, from 2004 to 2006, and with Bielefeld University, Bielefeld, Germany, from 2006 to 2009. She was then a Specially Appointed Associate Professor with Osaka University from 2009 to 2017, and became a Senior Researcher with NICT in May 2017. Her research interests include computational modeling of human cognitive functions such as self-other cognition, imitation, and joint attention, and design of assistant systems for developmental disorders. She was the Project Leader of MEXT Grant-in-Aid for Scientific Research on Innovative Areas “Computational modeling of social cognitive development and design of assistance systems for developmental disorders” from 2012 to 2017, and is the Project Leader of JST CREST “Cognitive Mirroring: Assisting people with developmental disorders by means of self understanding and social sharing of cognitive processes” since 2016.



Emre Ugur is an assistant professor in Department of Computer Engineering Department, Bogazici University, Turkey. After receiving his PhD in Computer Engineering from Middle East, he worked at ATR Japan as a researcher (2009-2013), at University of Innsbruck as a senior researcher (2013-2016), and at Osaka University as specially appointed assistant professor (2015, 2016). He participated in several EU funded projects, including Xperience, ROSSI, MACS and Swarm-bots, and is currently PI of IMAGINE project supported by European Commission, Horizon 2020 Programme. He is interested in developmental and cognitive robotics, and intelligent and adaptive manipulation.