

Learning at the Ends: From Hand to Tool Affordances in Humanoid Robots

Giovanni Saponaro*, Pedro Vicente*, Atabak Dehban*[†], Lorenzo Jamone^{‡*},
Alexandre Bernardino* and José Santos-Victor*

*Institute for Systems and Robotics, Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal

Email: {gsaponaro,pvicente,adehban,alex,jasv}@isr.tecnico.ulisboa.pt

[†]Chamपालimaud Centre for the Unknown, Lisbon, Portugal

[‡]ARQ (Advanced Robotics at Queen Mary),

School of Electronic Engineering and Computer Science, Queen Mary University of London, United Kingdom

Email: l.jamone@qmul.ac.uk

Abstract—One of the open challenges in designing robots that operate successfully in the unpredictable human environment is how to make them able to predict what actions they can perform on objects, and what their effects will be, i.e., the ability to perceive object affordances. Since modeling all the possible world interactions is unfeasible, learning from experience is required, posing the challenge of collecting a large amount of experiences (i.e., training data). Typically, a manipulative robot operates on external objects by using its own hands (or similar end-effectors), but in some cases the use of tools may be desirable; nevertheless, it is reasonable to assume that while a robot can collect many sensorimotor experiences using its own hands, this cannot happen for all possible human-made tools.

Therefore, in this paper we investigate the developmental transition from hand to tool affordances: what sensorimotor skills that a robot has acquired with its bare hands can be employed for tool use? By employing a visual and motor imagination mechanism to represent different hand postures compactly, we propose a probabilistic model to learn hand affordances, and we show how this model can generalize to estimate the affordances of previously unseen tools, ultimately supporting planning, decision-making and tool selection tasks in humanoid robots. We present experimental results with the iCub humanoid robot, and we publicly release the collected sensorimotor data in the form of a hand posture affordances dataset.

I. INTRODUCTION

Robotics is shifting from the domain of specialized industrial manipulators operating in a predictable and repeatable environment, towards more adaptable systems that are expected to work among people. This new generation of robots needs to operate in an unstructured and dynamic world, exposed to a multitude of objects and situations that cannot be modeled *a priori*. A crucial ability needed by these robots to succeed in such environment is to be able to predict the effects of their own actions, or to give a reasonable estimate when they interact with objects that were never seen before; this applies both to objects that are acted upon, and to those that may be used as tools to extend the capabilities of the robot body. The interplay between these elements (actions, objects, tool, effects) can be described by resorting to the concept of affordances (action possibilities),




Figure 1. The iCub humanoid robot performing motor actions with different hand postures onto a physical object. In the background screen, we show the visual routines that monitor the evolution of the environment.

a concept that was initially proposed in psychology [1], supported by neuroscience evidence, and applied to several artificial intelligence (AI) and robotics systems; indeed, a growing line of research takes inspiration from human intelligence and development in order to design robust robot algorithms for affordance learning and perception, that can, in turn, support action selection and planning in unstructured environments [2].

In humans, learning the affordances of one's own hands (i.e., what actions one can do with them, what effects one can obtain) is a long developmental process that begins in infancy [3], [4], and continues during childhood through exploration of different actions and different objects [5], [6]. Functional tool use appears later [7]–[10], and the knowledge previously acquired by babies during manual exploration of objects is likely to play a role. Definitely, one of these roles is that the increased hand dexterity acquired during development allows the child to correctly grasp, manipulate and orient a tool; however, another role may be that the child “sees” in the shapes of some tools relevant characteristics that remind the child of previously used shapes of the own hands (although no experimental evidence of this perceptual

skill has been provided in the developmental psychology literature, as far as the authors know).

Inspired by these observations in developmental psychology, and motivated by a clear need of autonomous robotic systems, in this paper we investigate a possible developmental *transition from hand to tool affordances*. In particular, we explore how a learned representation of hand affordances can be generalized to estimate the affordances of tools which were never seen before by the robot. We train a probabilistic model of affordances, relating visual features of (i) different robotic hand postures and (ii) different objects, with the resulting effects caused by the robot motor actions onto such objects; training data are collected during several experiments in which the iCub humanoid robot performs manual actions on objects located on a table (see Fig. 1). Our probabilistic model is implemented as a Bayesian Network [11], we publicly release a novel dataset of hand posture affordances (<http://vislab.isr.tecnico.ulisboa.pt/> → Datasets), and we test it for generalization against an available dataset of tool affordances [12].

This article is structured as follows. Sec. II overviews related work about hand and tool affordances in the contexts of psychology and robotics. Sec. III explains our proposed framework for learning the affordances of different hand postures on a humanoid robot, and how its generalization ability can prove fruitful for robot autonomy. Sec. IV reports our current experimental results. Finally, in Sec. V we give concluding remarks and discuss the significance and scope of our contribution.

II. RELATED WORK

Several researchers have investigated the role of hand actions during human intelligence development for learning to deal with the uncertainty of the real world (e.g., toddler visual attention [13]) and *tool use*. Piaget documents an observation where his daughter makes an analogy between a doll’s foot hooking her dress, and her own finger bent like a hook [14]. Tool use competence in humans emerges from explorative actions, such as those performed with the child’s bare hands in the first year [15]. A longitudinal study on infants [10] shows that, at around 16–20 months, children start to intentionally and successfully bring faraway toys closer to themselves with the help of tools such as rakes.

Lockman [7] suggests that the actions employed by toddlers on a daily basis with their everyday objects in their surroundings, likely put forward the idea that these actions initially incorporate many of the (previously learned) motor patterns that infants employ with their hands and arms for exploring and learning. Szokolszky [8] stresses how tool use is dependent and continuous with other percepts: action routines, such as reaching, grasping, focusing on an object or on a person, and eating with the hand. In [9], Lobo highlights the following points about the relationship between early self-exploration behaviors and developing object exploration behaviors: (i) infants are already actively engaging in exploratory behaviors to inform themselves about the

affordances of their own bodies, objects, and the intersect of the two in the first months of life; (ii) the emergence of reaching is an important step forward towards advanced object exploration and advanced self-exploration; (iii) the behaviors that infants use to explore their own bodies and surfaces during the first months of life may form the set of behaviors from which they later *choose*, as they begin to interact with objects. It is with these considerations in mind that, in this paper, we pursue a robotic model that transfers hand knowledge to tool knowledge, and in the experimental part we verify the applicability of our model for task-dependent *tool selection*, given previous exploratory hand knowledge.

In robotics, there is an increasingly large body of literature in affordances and computational models of affordances. A recent overview is given in [2]. The idea behind robots learning object affordances, or action possibilities, is that knowledge acquired through self-exploration and interaction with the world can be used to make inferences in new scenarios and tasks, such as prediction [16], tool use [17]–[20], and planning [21], [22]. Recently, Schoeler and Wörgötter [23] introduced a framework for analyzing tools by modeling their dynamics and providing object graphs based on their parts, to support the conjecture of a cognitive hand-to-tool transfer (e.g., to understand that a helmet and a hollow skull can both be used to transport water). However, to the best of our knowledge, our work is the first contribution, in the robot affordances field, which explicitly looks at the visuomotor possibilities offered by different hand morphologies and postures (e.g., hands with extended fingers, hands with closed fingers), and exploits this information to acquire (through self-exploration) a model that is able to generalize to novel situations for a robotic agent, including making the crucial developmental leap from hand use to tool use, as observed in babies by psychology studies.

III. PROPOSED APPROACH

In this paper, we investigate how a humanoid robot can learn the affordances of different hand postures, similarly to the developmental processes observed during the first year of life in babies. We use a probabilistic model of affordances which relates (i) visual features of the agent’s own hands, (ii) visual features of a target object located on a surface, (iii) a motor action, and (iv) the resulting effects of the action onto the object, in the sense of the physical displacement compared to the initial position. We currently use three different robot hand postures, shown in Fig. 2.

A. Motor Control

The motor actions that we provide to the robot for learning the affordances of its actions performed with bare hands are: tapping an object from the left side (with the palm of the hand), tapping an object from the right side (with the back of the hand), pushing an object away from the agent, and drawing the object towards the agent (i.e., pulling). The low-level control routines to realize the motor behaviors are based

on works and software modules previously made available by other researchers in the iCub community [24], [25]. The location of the target object (i.e., the location where the robot performs an action) can be anywhere on the table, provided that it is within the reachable space of the robot end-effector, and that it satisfies geometric safety limits to avoid self-collisions. We determine this location with visual segmentation routines.

B. Internal Model and Hand Imagination

From a developmental psychology perspective, body awareness appears to be an incremental learning process that starts in early infancy [26] or probably even prenatally [27]. Such awareness is supported by a neural representation of the body that is constantly updated with multimodal sensorimotor information acquired during motor experience and that can be used to infer the limbs’ position in space and guide motor behaviors: a body schema [28]. The body schema is not necessarily restricted to proprioceptive and somatosensory perception *per se*, but can also integrate visual and perhaps auditory information. Moreover, sub-representations of the body according to their function are also proposed by several authors. The visual information is usually called body image [29].

In this work, the learning of the sensorimotor mapping of affordances is based on the joint-space representation. The body schema and the forward model are then exploited to simulate motor behaviors and imagining sensory outcomes. The same approach seems to characterize humans [30]. The hand imagination is implemented through an internal model simulator developed within a previous work of ours [31]. From a technical perspective, using the simulated robot rather than the real one to obtain the hand posture visual shape, serves to filter out noise from the image processing pipeline. Although it is not always true that we can generalize from simulation to the real robots, in this case, we adopt a graphically and geometrically precise appearance model of the robotic hand (based on the CAD model), therefore we can use the internal model simulation without losing generality or compromising the overall idea (see the real and simulated hands in Fig. 2).

C. Visual Features

We characterize hand postures and physical objects by their shape (i.e., the contour or silhouette obtained with segmentation algorithms), and we extract exemplary 2D visual features from them. The features that we extract are pre-categorical shape descriptors computed as geometric relationships between perimeter, area, convex hull and approximated shapes of the segmented silhouettes of the objects in front of the robot, and they are based on [32]. Our visual feature extraction is similar to our previous works [19], [20], [22], however it now incorporates a richer set of 13 features instead of 5: convexity, eccentricity, compactness, circularity, squareness, number of convexity defects (i.e., number of cavities along the contour, for example the “holes” between fingers in a

hand image), and seven central normalized moments. It is worth noting that we do not classify or label our shapes into classes, but we merely reason about their shape-derived features: this gives our system some flexibility, making it able to process unknown situations not seen during training.

D. Affordance Learning

Fig. 3 shows the structure of the Bayesian Network [11] that we train with robot self-exploration hand affordance data, using the hand postures of Fig. 2. This structure is similar to the one that gave us the best generalization results in our previous work [20], thanks to its dimensionality reduction which reduced the number of edges, therefore the computational complexity of training and testing, and most importantly it reduces the amount of training data required to observe the emergence of some learning effect.

All the nodes of the network, except for EffectX and EffectY, are entered as interventional variables, which means that we force a node to take a specific value (i.e., they are the prior inputs, when reasoning on the effects), thereby effectively severing its incoming arcs. The two effect nodes can be seen as the outputs (when we perform inference queries about their values). The inference performed on the networks is of the type $p(\text{Effect} \mid \text{parents}(\text{Effect}))$, which, considering the topology of our network from Fig. 3, amounts to this marginalization over our two effect nodes (horizontal and vertical displacement of the object):

$$p(\text{EffectX}, \text{EffectY} \mid M, O, A), \quad (1)$$

where M is the vector of features of the manipulator (if the manipulator is a hand we refer to its feature vector as H , if it is a tool we refer to its feature vector as T), O is the vector of features of the target object, A is the motor action identifier.

IV. RESULTS

In this section, we present the results obtained from our hand affordance network which is trained with robot experience on objects using different hand postures and motor actions, and we assess its performance.

A. Dataset Inspection

Our experimental data is obtained by making manipulation experiments on an iCub humanoid robot, in a setup like the one shown in Fig. 1, using its left arm for data collection. We consider 4 motor actions A (tapFromRight, tapFromLeft, draw, push), 2 objects O (lego piece, pear), 3 hand postures H (straight fingers, bent fingers, arched fingers). We extract the visual features from both O and H (before performing the actions). The dataset is publicly available at <http://vislab.isr.tecnico.ulisboa.pt/> → Datasets.

In Fig. 4 we show the distributions of the motion effects onto target objects caused by the robot influence when it touches objects with its manipulator. In particular, Fig. 4a shows the effects of using the different hand postures. For comparison, Fig. 4b depicts the effect of using the elongated




Figure 2. The three hand postures adopted in this work. Left column: straight hand; center column: bent hand; right column: arched hand. The first two rows are real robot postures seen from different viewpoints; the last row shows the simulated body schema CAD model of the top viewpoint. From the latter simulated view, we obtain the segmented silhouette contour of the hand and its shape features (see Sec. III-C).

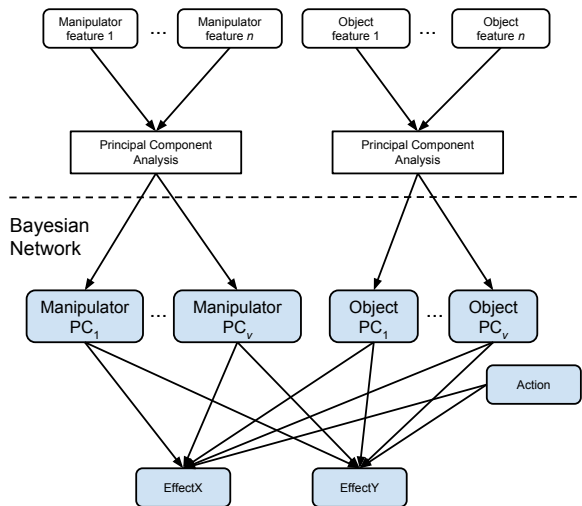


Figure 3. Structure of the Bayesian Network that we use to model the affordance relationships between robot manipulator, target object, motor action, and resulting effects, similarly to [20]. The raw visual features of manipulators and objects are real-valued and normalized between 0 and 1 (see Sec. III-C). In the network, Principal Components of manipulators and objects are discretized into 2 bins, the action node is a discrete index (see Sec. III-A), and the effect nodes are discretized into 5 bins, according to the hyper-parameters used during training, listed in Table I.

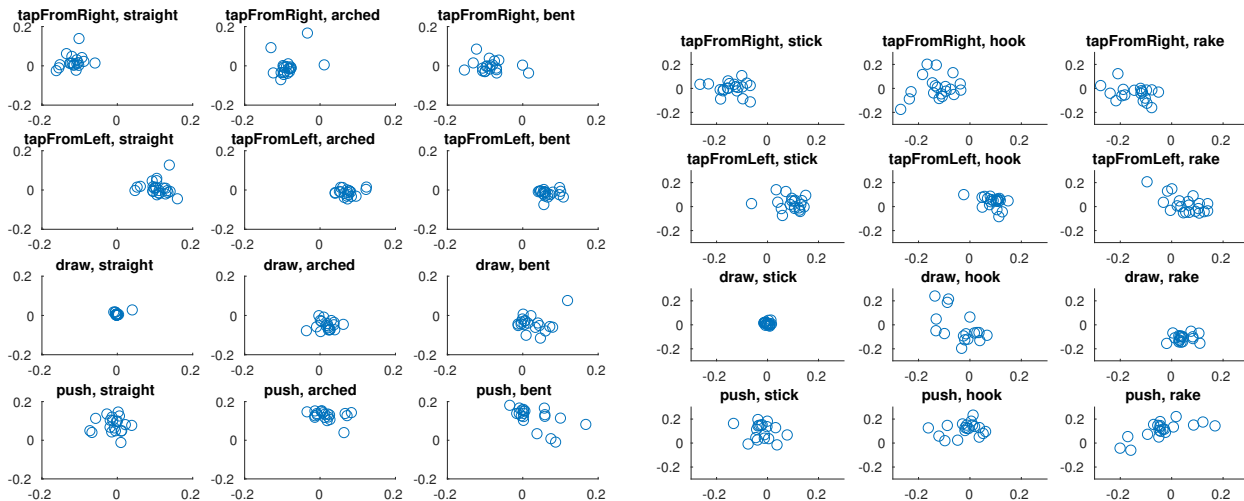
tools (Fig. 5) on the same objects. Visual inspection reveals the similarities in the effect of using tools or hands, for example, tapping from left usually results in the object moving to the right. Another prominent similarity is that drawing with a stick or with straight hand posture causes only minimal movement.

B. Data Augmentation

Similar to previous research [12], [19], [20], we assume that the affordance of an object and of a robot manipulator is viewpoint-invariant. By exploiting this notion, it is possible to artificially augment the trials data using *multiple views* of manipulators and objects. In all of the following experiments, we have used at least 10 viewpoints of each object and manipulator, effectively multiplying the number of available samples by more than 100 times.

C. Effect Prediction

One way to assess the quality of the learned Bayesian Network of Fig. 3 is to predict the effect distribution, given the descriptors of manipulator, object, and action, i.e., the direct application of (1). To this end, we have empirically divided the effect distribution along each axis into five bins (a



(a) Motion caused with the robot *hands* when using different actions and hand postures, as observed when interacting with 2 objects multiple times in the new experiments from this paper.

(b) Motion caused with *tools* when using different actions and tool types, from [12]. Here we show only the interactions with 2 objects, to be consistent with Fig. 4a.

Figure 4. Motion caused by different robotic manipulators when using different actions and manipulator morphologies: in Fig. 4a we use different hand postures, whereas in Fig. 4b we vary tool types for comparison. Each plot displays the geometrical displacement along horizontal and vertical direction (in meters, measured from the object initial position) from the point of view of the robot (the robot is at the 0 in the x-axis marker). For example, tapping an object from the right (tapFromRight action) usually results in making the object shift to the left direction; drawing (i.e., pulling) an object closer only works if the manipulator morphology is appropriate.




Figure 5. The three tools used in [12]: (a) stick, (b) rake and (c) hook. They provide different affordances when grasped by the hand of the robot and employed for performing motor actions onto objects. In this paper, we consider these tool affordances for comparison with our novel hand affordances.

list of the hyper-parameters that we used for training our network is reported in Table I for reproducibility). If the network predicts the correct effect bin out of five, it gets a point. Since there exist two axis directions, a random predicting machine would be correct 1/25 of the time.

Using the same network parameters but training with different data, the following accuracy scores were obtained:

- train 80% hand, test 20% hand: accuracy 72%;
- train 80% tool, test 20% tool: accuracy 58%;
- train 100% hand, test 100% tool: accuracy 53%.

To explain these scores, we note that motor control on the iCub is noisy, and actions on this platform are not deterministic or repeatable (e.g., when commanding the robot twice starting from an initial position, the same motor command can produce two slightly different configurations). Even so, in the accuracies and in Fig. 4 we see that tool effects are more varied than hand effects, making tools less reliable (i.e., more noisy) than hands. Nevertheless, by only training on the hand data, we obtain an accuracy that is comparable with the case where the network is trained on tool data, demonstrating the generalization of our proposed method.

D. Tool Selection

One question that this paper tries to embark upon is the following: if an agent gains the knowledge of how her hand postures can affect the environment, can she generalize this knowledge to other tools which look similar to her hands? To answer this question in the scope of the presented scenario, we conduct the following experiment. We suppose that an agent has defined a goal, for example to pull an object towards herself. She knows that the correct action for this task will be to *draw*, however the object is out of hand’s reach and one of the presented tools (see Fig. 5) must be selected for the task.

In this scenario, an agent looks at the available tools and the target object and performs a mental simulation of the known action. A tool is selected if it is expected to cause a movement of the target object *along the desired direction*, and it is rejected if no movement is predicted, or if the object is predicted to move against the desired direction. Because in this work we divide the direction into five bins (see Sec. IV-C), we compare the sum of the predictions in the two desired-movement bins against the sum of the predictions in the remaining bins. Since there was no interaction with the tool, it is necessary to generalize from the knowledge of previous hand explorations to tools in a zero-shot manner.

As an example of a successful generalization, the agent should predict that drawing an object with a stick is pointless because she has already experimented drawing objects with a straight hand posture, and the visual descriptors of straight hands are similar to those of a stick. Table II shows the result of this inquiry. We have also implemented a baseline in which the agent has already experienced with the tools and is asked

Table I
HYPER-PARAMETERS USED TO TRAIN THE BAYESIAN NETWORK OF FIG. 3 FOR PREDICTING THE DISTRIBUTION OF THE EFFECTS.

parameter	value (and comment)
number of Principal Component Analysis (PCA) blocks	one for manipulator, one for object
number of components of each PCA block	2
number of discretization values (bins) of each PCA component	2
number of discretization values (bins) of each Effect node	5
intervals (in meters) of the Effect bins	$]-\infty, -0.06]$, $]-0.06, -0.025]$, $]-0.025, 0.025]$, $]0.025, 0.06]$, $]0.06, \infty[$

Table II
TOOL SELECTION RESULTS OBTAINED FROM OUR “HAND TO TOOL” (HT) NETWORK, COMPARED TO ONES OBTAINED FROM THE BASELINE “TOOL TO TOOL” (TT) NETWORK [12].

action	stick	hook	rake
tapFromRight	HT: 1.0 (TT: 1.0)	HT: 1.0 (TT: 1.0)	HT: 1.0 (TT: 1.0)
tapFromLeft	HT: 1.0 (TT: 1.0)	HT: 1.0 (TT: 1.0)	HT: 1.0 (TT: 1.0)
draw	HT: 0.5385 (TT: 0.1538)	HT: 0.6154 (TT: 0.1538)	HT: 1.0 (TT: 0.4615)
push	HT: 1.0 (TT: 1.0)	HT: 1.0 (TT: 1.0)	HT: 1.0 (TT: 1.0)

to select the correct tool. As expected, all the tools can be used for the desired effects, and it is only the draw action which requires a tool with a specific shape. The numbers are normalized, as they correspond to different views of the tool and object, and they reflect the percentage of the cases where that specific tool was selected.

V. CONCLUSIONS

We introduce a computational model of hand affordances: a Bayesian Network that relates robot actions, visual features of hand postures, visual features of objects and produced effects, allowing a humanoid robot to predict the effects of different manual actions. Interestingly, we show how this knowledge, acquired by the robot through autonomous exploration of different actions, hand postures and objects, can be generalized to tool use, and employed to estimate the most appropriate tool to obtain a desired effects on an object, among a set of tools that were never seen before.

A few important comments have to be made to better focus the scope of our contribution. Our results show that, in some specific cases, it is indeed possible to generalize what was learned about hand affordances to tools that were never seen before. Clearly, this is limited to a subset of all the possible human-made tools that a humanoid robot could see and possibly use; however, the previous knowledge about hand affordances can give the robot the possibility to make a good initial estimate of how a tool could be used. It would be very interesting to investigate how further sensorimotor experience with tools can be integrated in the learned model, and possibly permit better predictions. Also, it may be interesting to explore the opposite direction: can

the knowledge acquired with a specific tool be re-used to estimate the effects of manual actions without the tool, or to shape the robot hand in the best posture to achieve some effects (i.e., from tool to hand affordances)?

As we made clear in the introduction, to the best of our knowledge, there is no clear experimental evidence that human children exploit visual information about their hand postures to estimate the possible affordances of an external object (i.e., a tool); neither is there evidence that visual information extracted from successfully used tools is employed to control hand postures. Also, although the early manipulation attempts of the child are probably a necessary prerequisite for effective tool use, a great deal of knowledge is indeed acquired during direct sensorimotor exploration using the tools, and during observations of tool actions performed by others (e.g., by the parents). However, our results can serve a twofold purpose for the developmental robotics community: (i) we propose a robot learning framework that presents practical advantages for robot autonomy, at least in the limited number of situations we experimented, since it permits to generate meaningful predictions about a non-finite set (i.e., tools) from experiences in a finite set (i.e., hand postures); and (ii) we lay the basis for future discussion and research in developmental psychology and cognitive science about the possible existence of similar phenomena in humans.

ACKNOWLEDGMENTS

This work was partially supported by the FCT project UID/EEA/50009/2013 and grant PD/BD/135115/2017. We thank the anonymous reviewers for their helpful suggestions.

REFERENCES

- [1] J. J. Gibson, *The Ecological Approach to Visual Perception: Classic Edition*. Psychology Press, 2014, Originally published in 1979 by Houghton Mifflin Harcourt.
- [2] L. Jamone, E. Ugur, A. Cangelosi, L. Fadiga, A. Bernardino, J. Piater, and J. Santos-Victor, “Affordances in psychology, neuroscience and robotics: a survey”, *IEEE Transactions on Cognitive and Developmental Systems*, 2016.
- [3] E. J. Gibson, *An Odyssey in Learning and Perception*. MIT Press, 1994.
- [4] D. K. James, “Fetal learning: A critical review”, *Infant and Child Development*, vol. 19, pp. 45–54, 2010.
- [5] D. Rosenblatt, “Developmental trends in infant play”, *Biology of Play*, pp. 33–44, 1977.

- [6] D. A. Rosenbaum, *Human Motor Control*. London, UK: Academic Press, 2009.
- [7] J. J. Lockman, “A Perception–Action Perspective on Tool Use Development”, *Child Development*, vol. 71, no. 1, pp. 137–144, 2000.
- [8] Á. Szokolszky and É. Devánszky, “The Development of Spoon-Use in the Daily Routine of Infants: A Naturalistic Observation Study”, in *Studies in Perception and Action IX: Fourteenth International Conference on Perception and Action*, 2010, pp. 74–78.
- [9] M. A. Lobo and J. C. Galloway, “The onset of reaching significantly impacts how infants explore both objects and their bodies”, *Infant Behavior and Development*, vol. 36, no. 1, pp. 14–24, 2013.
- [10] J. Fagard, L. Rat-Fischer, and J. K. O’Regan, “The emergence of use of a rake-like tool: A longitudinal study in human infants”, *Frontiers in Psychology*, vol. 5, 2014.
- [11] K. Murphy, *Machine Learning: A Probabilistic Perspective*. MIT Press, 2012.
- [12] A. Dehban, L. Jamone, A. R. Kampff, and J. Santos-Victor, “A Moderately Large Size Dataset to Learn Visual Affordances of Objects and Tools Using iCub Humanoid Robot”, in *European Conference on Computer Vision*, ser. Workshop on Action and Anticipation for Visual Learning, 2016.
- [13] C. Yu, L. B. Smith, H. Shen, A. F. Pereira, and T. Smith, “Active Information Selection: Visual Attention Through the Hands”, *IEEE Transactions on Autonomous Mental Development*, vol. 1, no. 2, pp. 141–151, Aug. 2009.
- [14] J. Piaget, *Play, Dreams & Imitation in Childhood*. W. W. Norton & Company, 1962.
- [15] L. B. Smith, S. Street, S. S. Jones, and K. H. James, “Using the axis of elongation to align shapes: Developmental changes between 18 and 24 months of age”, *Journal of Experimental Child Psychology*, vol. 123, pp. 15–35, 2014.
- [16] L. Montesano, M. Lopes, A. Bernardino, and J. Santos-Victor, “Learning Object Affordances: From Sensory–Motor Maps to Imitation”, *IEEE Transactions on Robotics*, vol. 24, no. 1, pp. 15–26, 2008.
- [17] A. Stoytchev, “Learning the Affordances of Tools using a Behavior-Grounded Approach”, in *Towards Affordance-Based Robot Control*, Springer, 2008, pp. 140–158.
- [18] V. Tikhonoff, U. Pattacini, L. Natale, and G. Metta, “Exploring affordances and tool use on the iCub”, in *IEEE-RAS International Conference on Humanoid Robots*, 2013, pp. 130–137.
- [19] A. Gonçalves, G. Saponaro, L. Jamone, and A. Bernardino, “Learning Visual Affordances of Objects and Tools through Autonomous Robot Exploration”, in *IEEE International Conference on Autonomous Robot Systems and Competitions*, 2014, pp. 128–133.
- [20] A. Gonçalves, J. Abrantes, G. Saponaro, L. Jamone, and A. Bernardino, “Learning Intermediate Object Affordances: Towards the Development of a Tool Concept”, in *IEEE International Conference on Developmental and Learning and on Epigenetic Robotics*, 2014, pp. 482–488.
- [21] N. Krüger, C. Geib, J. Piater, R. Petrick, M. Steedman, F. Wörgötter, A. Ude, T. Asfour, D. Kraft, D. Omrčen, A. Agostini, and R. Dillmann, “Object–Action Complexes: Grounded Abstractions of Sensory-Motor Processes”, *Robotics and Autonomous Systems*, vol. 59, no. 10, 2011.
- [22] A. Antunes, L. Jamone, G. Saponaro, A. Bernardino, and R. Ventura, “From Human Instructions to Robot Actions: Formulation of Goals, Affordances and Probabilistic Planning”, in *IEEE International Conference on Robotics and Automation*, 2016, pp. 5449–5454.
- [23] M. Schoeler and F. Wörgötter, “Bootstrapping the Semantics of Tools: Affordance analysis of real world objects on a per-part basis”, *IEEE Transactions on Cognitive and Developmental Systems*, vol. 8, no. 2, pp. 84–98, 2016.
- [24] U. Pattacini, F. Nori, L. Natale, G. Metta, and G. Sandini, “An Experimental Evaluation of a Novel Minimum-Jerk Cartesian Controller for Humanoid Robots”, in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct. 2010, pp. 1668–1674.
- [25] A. Roncone, U. Pattacini, G. Metta, and L. Natale, “A Cartesian 6-DoF Gaze Controller for Humanoid Robots”, in *Robotics: Science and Systems*, 2016.
- [26] C. von Hofsten, “An action perspective on motor development”, *Trends in Cognitive Sciences*, vol. 8, pp. 266–272, 6 2004.
- [27] R. Joseph, “Fetal brain behavior and cognitive development”, *Developmental Review*, vol. 20, pp. 81–98, 1 2000.
- [28] G. Berlucchi and S. Aglioti, “The body in the brain: Neural bases of corporeal awareness”, *Trends in Neurosciences*, vol. 20, pp. 560–564, 12 1997.
- [29] G. Riva, “Out of my real body: Cognitive neuroscience meets eating disorders”, *Frontiers in Human Neuroscience*, vol. 8, p. 236, 2014.
- [30] R. C. Miall and D. M. Wolpert, “Forward Models for Physiological Motor Control”, *Neural Networks*, vol. 9, pp. 1265–1279, 8 1996.
- [31] P. Vicente, L. Jamone, and A. Bernardino, “Robotic Hand Pose Estimation Based on Stereo Vision and GPU-Enabled Internal Graphical Simulation”, *Journal of Intelligent & Robotic Systems*, vol. 83, no. 3-4, pp. 339–358, 2016.
- [32] D. Zhang and G. Lu, “Review of Shape Representation and Description Techniques”, *Pattern Recognition*, vol. 37, no. 1, pp. 1–19, Jan. 2004.