

Parental Action Modification Highlighting the Goal versus the Means

Yukie Nagai and Katharina J. Rohlfing

Research Institute for Cognition and Robotics, Bielefeld University

33594 Bielefeld, Germany

Email: {yukie, rohlfing}@techfak.uni-bielefeld.de

Abstract—Parents significantly alter their infant-directed actions compared to adult-directed ones, which is assumed to assist the infants’ processing of the actions. This paper discusses differences in parental action modification depending on whether the goal or the means is more crucial. When demonstrating a task to an infant, parents try to emphasize the important aspects of the task by suppressing or adding their movement. Our hypothesis is that in a goal-crucial task, the initial and final states of the task should be highlighted by parental actions, whereas in a means-crucial task the movement is underlined. Our analysis using a saliency-based attention model partially verified it: When focusing on the goal, parents tended to emphasize the initial and final states of the objects used in the task by taking a long pause before/after they started/fulfilled the task. When focusing on the means, parents shook the object to highlight it, which consequently made its state invisible. We discuss our findings regarding the uniqueness and commonality of the parental action modification. We also describe our contribution to the development of robots capable of imitating human actions.

I. INTRODUCTION

The social context of learning has increasingly gained attention not only in developmental psychology and cognitive science but also in robotics [1]. The focus of it is the way in which a system can develop social skills through imitation. This focus has led to enormous progress in robotics research in designing robots that are able to reproduce actions presented by humans. However, copying actions does not imply the understanding of the action organizations in terms of the goal and the underlying shared representation. Call and Carpenter [2] differentiate between the forms of imitation depending on the understanding of the goal and the environmental result. Accordingly, *mimicry* captures an imitative behavior, in which an agent copies the movement without understanding the goal. The term of *imitation*, in contrast, refers to a reproduced action, in which an agent shows the goal understanding. Yet another form is *emulation*, which is the case where an agent achieves the goal by using another movement, i.e., an agent adopts its own action repertoire to produce the same result in the environment. The abilities to imitate and to emulate are discussed as evolutionary crucial since they allow agents to be engaged in collaborative interactions [3]. Interacting with others enables the agents to share knowledge about objects and events, and to learn new skills, which is targeted in social robotics. While cross-species developmental research has shown that emulation can be found in children [4], infants as young as 18 months [5], and even chimpanzees [4], [6],

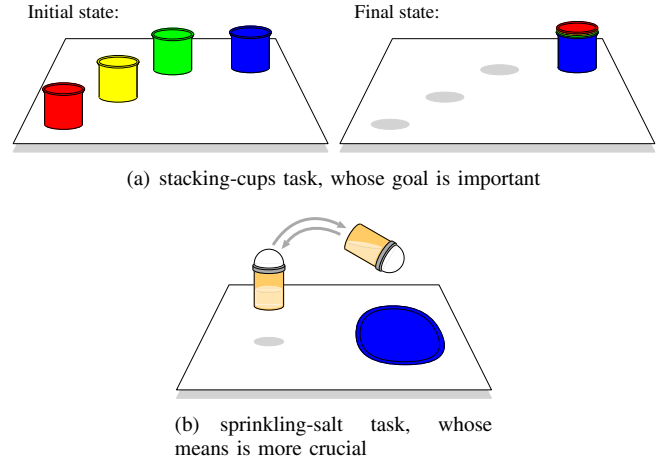


Fig. 1. What is important in a task, ‘the goal’ or ‘the means’?

the modeling of goal emulation has barely been touched in robotics research. This is due to the problem of ‘what to imitate’ [3], i.e., how to detect the relevant features of demonstrated actions, structure them, and recognize the goal of them [7]–[9].

A useful approach to this problem is to analyze how parents structure events for children [10]–[13]. Actions performed by adults would differ in terms of highlighting the goal of the actions or the means to achieve it. In studies concerning prelinguistic foundations for action components, it has been shown that infants as young as 7 months of age have the ability to notice changes in both the goal and the means [14]. However, it seems that encoding the goal is an easier task for 10-month-olds than encoding the means, which infants achieve around 13 months of age [15]. Together with research on social learning, this evidence suggests that (1) a task which is concerned with the means might be more difficult for infants to perceive and follow and that (2) parents might highlight the structure of their actions in a different way. Thus, pursuing the idea that parents try to educate infants’ attention, we investigated the differences and similarities in parental actions when they provided structures focusing either on the goal or the means.

Sections II and III present our analytical experiment, which compared the parental action demonstration of a goal-crucial task and a means-crucial one. We examined how parents



(a) camera image focusing on parent (b) camera image focusing on infant

Fig. 2. Parent-infant interaction in stacking-cups task

differently modified their actions in order to convey the significant information of the actions to infants. Section IV discusses the results toward designing the mechanisms for robot imitation. Our contribution to the issue ‘what to imitate’ in robot imitation is described. Section V concludes the paper with future issues.

II. ANALYSIS OF PARENTAL ACTION DEMONSTRATION

We conducted a comparative experiment, in which parents demonstrated two types of tasks where either the goal or the means was important. Parental actions to infants versus adult partners were analyzed using a computational attention model.

A. Participants

Twenty-seven parents (12 fathers and 15 mothers) participated in the experiment. Their infants were 8 to 11 month old ($M = 10.12$, $SD = 1.14$) when they joined the experiment.

B. Design and Procedure

The parents were asked to demonstrate two tasks: a stacking-cups task (see Fig. 1 (a)) and a sprinkling-salt task (see Fig. 1 (b)). The stacking-cups task, whose goal is important, was achieved by picking up a red, a yellow, and a green cup, and then putting them down into a blue one. The instruction to the parents was to produce the final state of the cups, and the means action (e.g., how to move the cups or which cup to move first) was *not* instructed. No restrictions on their speech or gestures were made. The sprinkling-salt task, whose means is more crucial, was fulfilled by taking a salt dispenser, tilting (and tapping) it, and then dropping salt onto the blue tray. The instruction to the parents was the same as in the stacking-cups task, that is, they were *not* asked for more than that of achieving the goal (i.e., to get salt on the tray). The salt dispenser was made of orange clear plastic so that the salt in it was visible from the outside.

In each task, the parents interacted first with their infant and then with their spouse. The parents’ actions as well as the partners’ responses were videotaped with a camera positioned slightly behind and above the partners. Fig. 2 shows a sample scene from the stacking-cups task experiment. The infant (or the adult partner) was sitting across a table, on which the task was demonstrated, from the parent. The objects used in the tasks were introduced at the beginning of each task.



(a) attended locations indicated (b) saliency map, linearly combining (c) to (g)




Fig. 3. Attended locations and saliency map in sprinkling-salt task

C. Analysis

We analyzed the parents’ actions using an attention model based on visual saliency. The model, proposed by Itti et al. [16], [17], is inspired by the neuronal and behavioral mechanism of primates, and has been demonstrated to simulate the bottom-up visual attention of humans. The saliency is here defined as the difference between a focused location and the surroundings in terms of primitive features. For instance, if a black dot is drawn in a white background, the dot is detected as salient because of the high contrast in the intensity. Of particular note on the model is the potential to imitate infants’ attention [18]. Because infants are supposed to rely more on bottom-up features (e.g., color and motion) than top-down information (e.g., context knowledge), the model can reveal where infants would look and why they do so.

Fig. 3 shows a result from the sprinkling-salt task experiment: (a) an input image to the saliency model, where the locations attended to by the model are indicated, (b) the corresponding saliency map, and (c)-(g) the conspicuity maps with respect to the color, the intensity, the orientation, the flicker (i.e., the change in the intensity), and the motion (i.e., the optical flow), respectively. The color and the intensity maps shown in Figs. 3 (c) and (d) present higher saliency for the father’s face, his hands, the salt dispenser, and the white tray because of their distinguishable colors and/or intensity. The orientation map shown in (e), by contrast, indicates only the contours of them. The flicker and the motion maps shown in (f) and (g), both of which are concerned with the movement, present strong saliency for the father’s right hand

with the salt dispenser because they were being shaken in the scene. As the result, the saliency map (see Fig. 3 (b)), which linearly combines the five conspicuity maps, reveals three most outstanding locations: two on the father's right hand and one on the salt dispenser. In our experiment, image locations for which saliency was higher than $0.9 \times$ the maximum in the current frame were chosen as the attended locations. Note that the five conspicuity maps were weighted equally when being combined because our interest was only in the bottom-up information enhanced by the parental actions.

An advantage of using the saliency model is that likely important locations can be detected without employing any top-down knowledge about the demonstrated tasks, the objects used in the tasks, or even the parents (e.g., a human model such as a face and skin color). Thus, it enables us to fairly analyze different tasks focusing on how important information is physically highlighted by parental action modifications. For a more detailed explanation about the model, refer to the authors' paper [12] and the Itti's originals [16], [17].

III. RESULTS AND DISCUSSION

A. Statistical Results of Highlighted Information

1) *Stacking-cups task*: Figs. 4 and 5 are the results for the stacking-cups task¹: Fig. 4 shows the attention proportion by the saliency model (a) before the task started (for 2 [sec]), (b) during it being demonstrated, and (c) after it was fulfilled (for 2 [sec]). The image locations attended to by the model were classified into four regions: the parent's face, his/her hands, the objects (i.e., the cups), and others (e.g., the parent's body, irrelevant objects, and background). The filled bars and the open ones represent the results for the Infant-Directed Interaction (IDI) and the Adult-Directed Interaction (ADI), between which the significant difference and the statistical trend are denoted by ‘***’ and ‘*’, respectively. Fig. 5 is the result concerning the saliency value for the cups. It represents how much the inherent features (i.e., the color, intensity, and orientation) contributed to the total saliency for the cups. For instance, a high contribution rate indicates that the cups were underlined by relative suppression of the parents' movement rather than by their additional movement to it (e.g., shaking the cups).

The results for the stacking-cups task are summarized as follows. In IDI compared to in ADI:

- (1.a) the parent's face attracted more attention during the task (non-parametric test; $Z = -2.556$, $p < 0.05$), but less attention after the task ($Z = -1.874$, $p = 0.061$),
- (1.b) the cups attracted more attention before the task ($Z = -2.045$, $p < 0.05$) and after it (parametric t-test; $t(14) = 1.846$, $p = 0.086$),
- (1.c) the task-irrelevant locations (i.e., others) attracted less attention before the task ($Z = -1.988$, $p < 0.05$), and
- (1.d) the inherent features contributed more to the saliency for the cups before ($Z = -2.040$, $p < 0.05$) and during the task ($Z = -3.045$, $p < 0.05$).

¹These results were duplicated from [12].

The reasons for the results are discussed in Section III-B, together with those for the sprinkling-salt task.

2) *Sprinkling-salt task*: Figs. 6 and 7 are the corresponding results for the sprinkling-salt task: Fig. 6 shows the proportion of the attended locations, and Fig. 7 the contribution rate of the inherent features to the saliency for the salt dispenser.

The statistical analysis on the sprinkling-salt task revealed the followings. In IDI compared to in ADI:

- (2.a) the salt dispenser attracted more attention before the task ($t(24) = 1.742$, $p = 0.094$) and after it ($Z = -3.027$, $p < 0.05$),
- (2.b) the task-irrelevant locations attracted less attention before the task ($Z = -2.029$, $p < 0.05$) and after it ($Z = -1.870$, $p = 0.061$), and
- (2.c) the inherent features contributed more to the saliency for the salt dispenser during the task ($t(24) = 2.038$, $p = 0.053$).

B. Similarities and Differences in Parental Action Modification

Comparing the results for the two tasks, we found some similarities as well as differences in the parental action modification.

1) *Commonality across the tasks*: Regardless of the type of the demonstrated task, the parents highlighted the objects used in the task before starting it and after fulfilling it when interacting with the infants. The results (1.b) and (2.a) showed the higher proportion of the attention to the objects in IDI than in ADI. If a task involves objects, it is definitely important to introduce them to infants. Knowing objects used in the task enables infants to appropriately follow the continuing demonstration and to understand the goal and/or the means of it. The parents therefore emphasized the objects regardless of the important aspects of the tasks.

This parental action modification caused a secondary effect: The parents relatively diminished the saliency for the task-irrelevant locations. The results (1.c) and (2.b) showed less attention to the irrelevant features especially before the task. Before the parents started demonstrating the task, the infants were not yet really engaged in the interaction. Because the infants' attention relies more on the bottom-up information, they might be visually exploring the environment. Therefore, the parents needed to educate the infants' attention by emphasizing the task-relevant locations (e.g., the objects), which resulted in the relatively less saliency for the irrelevant.

2) *Uniqueness depending on the task*: Although the results are not yet statistically verified, we found interesting phenomenon regarding differences in the parental action modification depending on the task.

Before starting the task, the parents highlighted the objects in different ways (refer to the leftmost bars in Figs. 5 and 7). In the stacking-cups task, the parents suppressed their body movement to underline the static state of the cups, whereas in the sprinkling-salt task, they produced additional movement to the salt dispenser. Figs. 8 and 9 give the sample scenes, where the objects attracted the attention due either to the inherent

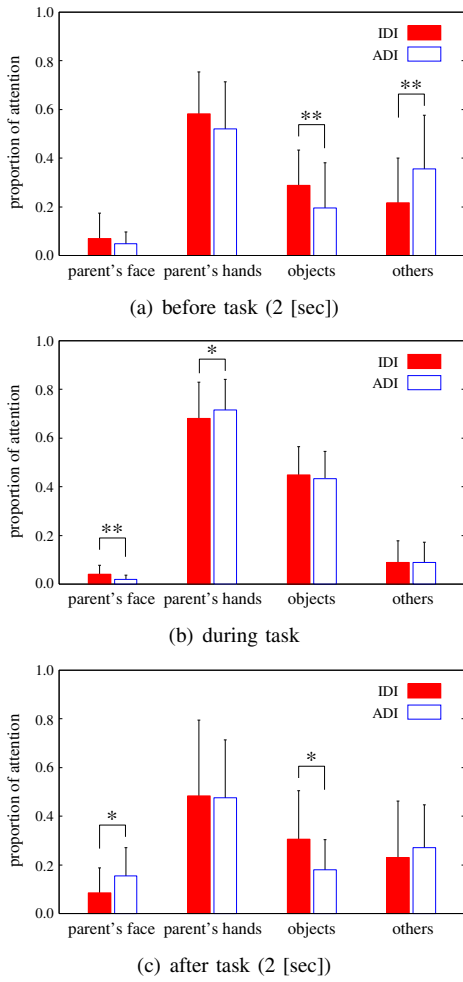


Fig. 4. Proportion of attended locations in *stacking-cups task*




Fig. 5. Contribution rate of inherent features to saliency for *cups*

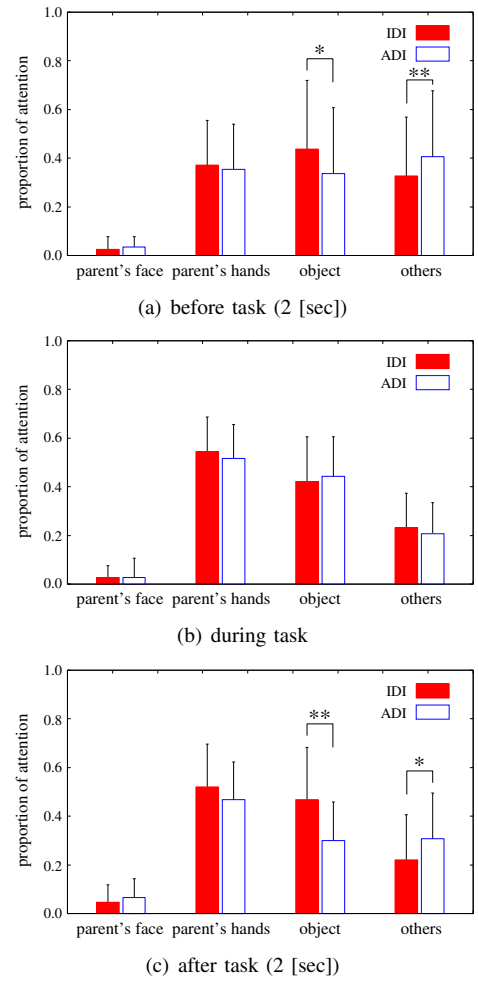


Fig. 6. Proportion of attended locations in *sprinkling-salt task*




Fig. 7. Contribution rate of inherent features to saliency for *salt dispenser*

features or the movement. The motion maps shown in Figs. 8 (c) and 9 (c) represent whether the movement contributed to the saliency. No salient point in Fig. 8 (c) (a black image) indicates no movement in the scene. In the stacking-cups task, the parents tended to take a long pause before starting the task, and did the same also after finishing it (refer to the rightmost in Fig. 5). They seemed to try to highlight both the initial and final states of the task by suppressing their movement. In the sprinkling-salt task, by contrast, their action of shaking the salt dispenser was effective in drawing the attention to it (see

Fig. 9) but not showing the state of it. Since the initial and final states of the dispenser (i.e., where and how it was placed) were not really important but the means to achieve the task was more crucial, the parents were supposed to just educate the infants' attention by producing the additional movement to it. Although it has not been statistically verified yet, these results suggest that the parental actions differ in terms of the way of highlighting the objects depending on the importance of the goal of a demonstrated task.

Regarding the attention to the parent's face, we found an




Fig. 8. Cups highlighted by suppression of parent’s body movement




Fig. 9. Salt dispenser highlighted by additional movement to it

interesting change in the stacking-cups task (1.a), but not in the sprinkling-salt task. In the stacking-cups task, the parent’s face attracted more attention during the task, but less after it. During the task demonstration, the parents often addressed the infants to indicate the meaningful segments of the actions and to alert the infants to the coming event [13]. For example, they talked to and/or smiled at the infants just after finishing moving each cup. This had the effect of communicating the sub-goal of the task. In contrast, in the sprinkling-salt task, it was not really significant. Because the task was shorter than the stacking-cups task and there was no notable sub-goal in it, the parents did not need to inform about it. Another reason may be the importance of the means of the task. The movement of handling the salt dispenser is crucial in achieving the task and thus should not be interrupted by any other factors. In contrast to the flexibility of the cup-handling movement, interruptions to the dispenser’s movement significantly affect the result. Therefore the parents rarely addressed the infants while demonstrating the sprinkling-salt task so that the infants could focus on the dispenser’s movement. Although a further analysis is required, we consider that the different importance of the means action may cause the different frequency of the parental social signals.

IV. ISSUE OF ‘WHAT TO IMITATE’ IN ROBOT IMITATION

Knowing the important aspects of a demonstrated action is a significant issue to be addressed in robot imitation. This is stated as ‘what to imitate,’ and has been intensively investigated in developmental robotics [7]–[9]. Here we discuss the issue with regard to (1) the robot’s attention and (2) the decision on imitating the goal versus the means. Our contributions to these arguments are also presented.

A. What Should a Robot Attend to?

If a robot is not provided with any a priori knowledge about a demonstrated task (e.g., the goal or the means), objects used

in the task (e.g., the color, the shape, or the size), or even the demonstrator (e.g., a human or another robot), how can it detect the relevant features to imitate from its sensory signals?

Breazeal and Scassellati [8], [9] pointed out the importance of the robot’s attention. An appropriate attention system allows a robot to decide which aspects of the sensory information are relevant or circumstantial, and to selectively direct its computational resources to the relevant ones. Scassellati [19] further proposed that besides the inherent saliency of signals, a robot can exploit:

- social cues given by a demonstrator,
- constraints in its embodiment,
- cross-modality of the perceptual information, and
- developmental progress both in the internal and the external complexity.

Social cues such as gaze direction and a pointing gesture presented by a demonstrator may directly indicate the important targets in the environment. The ability to comprehend such cues enables a robot to make a reasonable choice of the sensory signals [20], [21]. If a robot knows its physical bodily constraints, it would also reduce the perceptual space. Because possible actions are limited, a robot can make an assumption about the action being demonstrated. The redundancy of perceptual signals from several modalities is also helpful. Since multi-modal signals are sometimes correlated (e.g., a speech signal and a moving mouth), a strong cue detected in one modality can be used to extract the corresponding one in another modality.

Our approach presented in this paper is concerned with the fourth item in the above, i.e., the developmental constraints. The gradual increase both in the internal and the external complexity is suggested to facilitate robot learning [22]. In imitation learning, the early immaturity of the perceptual ability allows a robot to extract more prominent and thus meaningful information while ignoring noises. The saliency-based attention model realizes it only applying the primitive features, and thus enables a robot to decide where to attend in a fully bottom-up way. An important point is here that the attention model is coupled with the demonstrators’ proper teaching, which is the developmental constraint in the environment. In our experiment, the saliency model, which is rather weak in itself, could extract the relevant features of the demonstrated actions because the parental demonstration directed to infants had the effect of physically highlighting the important aspects of them. We have also been investigating human-robot interaction focusing on how people want to teach a robot and whether a robot equipped with the saliency model can induce parent-like teaching of human partners [23]. Our findings suggest that the saliency model has the potential to overcome the ‘where to attend’ problem by encouraging people to properly teach actions to a robot.

B. What Should a Robot Reproduce, ‘Goal’ or ‘Means’?

Even if a robot could detect the relevant sensory information from a task demonstration, how can it know what to reproduce in the imitation, i.e., the goal or the means? Billard and

colleagues (e.g., [24], [25]) have been working on this topic intensively. They proposed a probabilistic model by which a robot can decide whether to imitate the trajectory of a demonstrator's hand, its relative trajectory to an object, or the goal of the task. Their further effort enabled their robot to incorporate social cues presented by a demonstrator in efficiently finding the features to imitate [26]. However, their model requires the candidate features to be defined beforehand, that is, it can decide 'what to reproduce' but does not address the 'where to attend' problem.

Our approach has the potential to overcome both of the issues. The saliency model coupled with the parent-like proper teaching enables a robot to appropriately shift its attention so as to extract the relevant features of the demonstrated actions. Moreover, the parent-like teaching can tell the important aspects of the actions. For example, the difference in the way of highlighting the objects (i.e., whether by suppressing or adding movement) and that in the frequency of social signals might indicate whether the goal or the means should be imitated. In order to more strongly support our hypothesis, we will pursue our analysis of parental actions.

V. CONCLUSION

We hypothesized that parents differently modify their actions when interacting with infants depending on how important the goal or the means of the actions is. Parental action modification has a great potential to convey the important aspects of a demonstrated task. Our experimental results revealed that when focusing on the goal of a task, parents tend to suppress their body movement to highlight the initial and final states of the task, while focusing on the means they produce additional movement. That is, the importance of the goal significantly affects the parental actions.

Although we also speculated that in a means-crucial task, parents underline the movement to achieve the task, it was not statistically verified in the current analysis. We, however, consider that the different aspects of the highlighted information, e.g., when the objects and the parent's hands attract the saliency-model's attention and how long and intensively they do, will tell the effect of the parental modification to emphasize the means of the task. We therefore intend to more closely investigate our data as well as to extend our analysis to other tasks.

REFERENCES

- [1] C. L. Nehaniv and K. Dautenhahn, Eds., *Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions*. Cambridge University Press, 2007.
- [2] J. Call and M. Carpenter, "Three sources of information in social learning," in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds. MIT Press, 2002, pp. 211–228.
- [3] M. Carpenter and J. Call, "The question of 'what to imitate': Inferring goals and intentions from demonstrations," in *Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions*, C. L. Nehaniv and K. Dautenhahn, Eds. Cambridge University Press, 2007, pp. 135–151.
- [4] J. Call, M. Carpenter, and M. Tomasello, "Copying results and copying actions in the process of social learning: chimpanzees (pan troglodytes) and human children (homo sapiens)," *Animal Cognition*, vol. 8, pp. 151–163, 2005.
- [5] M. Carpenter, J. Call, and M. Tomasello, "Twelve- and 18-month-olds copy actions in terms of goals," *Developmental Science*, vol. 8, no. 1, pp. F13–F20, 2005.
- [6] L. M. Hopper, S. P. Lambeth, S. J. Schapiro, and A. Whiten, "Observational learning in chimpanzees and children studied through 'ghost' conditions," *Proceedings of the Royal Society: Biological Sciences*, vol. 275, no. 1636, pp. 835–840, 2008.
- [7] C. L. Nehaniv and K. Dautenhahn, "Like me? measures of correspondence and imitation," *Cybernetics and Systems: An International Journal*, vol. 32, pp. 11–51, 2001.
- [8] C. Breazeal and B. Scassellati, "Challenges in building robots that imitate people," in *Imitation in Animals and Artifacts*, K. Dautenhahn and C. L. Nehaniv, Eds. MIT Press, 2002, pp. 363–389.
- [9] —, "Robots that imitate humans," *Trends in Cognitive Sciences*, vol. 6, no. 11, pp. 481–487, 2002.
- [10] R. J. Brand, D. A. Baldwin, and L. A. Ashburn, "Evidence for 'motionese': modifications in mothers' infant-directed action," *Developmental Science*, vol. 5, no. 1, pp. 72–83, 2002.
- [11] K. J. Rohlfing, J. Fritsch, B. Wrede, and T. Jungmann, "How can multimodal cues from child-directed interaction reduce learning complexity in robots?" *Advanced Robotics*, vol. 20, no. 10, pp. 1183–1199, 2006.
- [12] Y. Nagai and K. J. Rohlfing, "Can motionese tell infants and robots 'what to imitate'?" in *Proceedings of the 4th International Symposium on Imitation in Animals and Artifacts*, 2007, pp. 299–306.
- [13] —, "Parental signal indicating significant state change in action demonstration," in *Proceedings of the 7th International Conference on Epigenetic Robotics*, 2007, pp. 205–206.
- [14] R. Pulverman, K. Hirsh-Pasek, S. Pruden, and R. M. Golinkoff, "Precursors to verb learning: Infant attention to manner and path," in *Action meets word: How children learn verbs*, K. Hirsh-Pasek and R. M. Golinkoff, Eds. Oxford Press, 2006, pp. 134–160.
- [15] S. Pruden, K. Hirsh-Pasek, and R. Golinkoff, "Current events: How infants parse the world and events for language," in *Understanding events: How humans see, represent, and act on events*, T. Shipley and J. Zacks, Eds. Oxford University Press, 2008.
- [16] L. Itti, C. Koch, and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 11, pp. 1254–1259, 1998.
- [17] L. Itti, N. Dhavale, and F. Pighin, "Realistic avatar eye and head animation using a neurobiological model of visual attention," in *Proceedings of the SPIE 48th Annual International Symposium on Optical Science and Technology*, 2003, pp. 64–78.
- [18] M. Schlesinger, D. Amso, and S. P. Johnson, "Simulating infants' gaze patterns during the development of perceptual completion," in *Proceedings of the 7th International Conference on Epigenetic Robotics*, 2007, pp. 157–164.
- [19] B. Scassellati, "Knowing what to imitate and knowing when you succeed," in *Proceedings of the AISB'99 Symposium on Imitation in Animals and Artifacts*, 1999, pp. 105–113.
- [20] C. Breazeal and B. Scassellati, "A context-dependent attention system for a social robot," in *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, 1999, pp. 1146–1151.
- [21] I. D. Goga and A. Billard, "Attention mechanisms for the imitation of goal-directed action in developmental robots," in *Proceedings of the 6th International Workshop on Epigenetic Robotics*, 2006.
- [22] Y. Nagai, M. Asada, and K. Hosoda, "Learning for joint attention helped by functional development," *Advanced Robotics*, vol. 20, no. 10, pp. 1165–1181, 2006.
- [23] Y. Nagai, C. Muhl, and K. J. Rohlfing, "Toward designing a robot that learns actions from parental demonstrations," in *Proceedings of the 2008 IEEE International Conference on Robotics and Automation*, 2008, pp. 3545–3550.
- [24] S. Calinon, F. Guenter, and A. Billard, "Goal-directed imitation in a humanoid robot," in *Proceedings of the International Conference on Robotics and Automation*, 2005.
- [25] A. G. Billard, S. Calinon, and F. Guenter, "Discriminative and adaptive imitation in uni-manual and bi-manual tasks," *Robotics and Autonomous Systems*, vol. 54, no. 5, pp. 370–384, 2006.
- [26] S. Calinon and A. Billard, "A framework integrating statistical and social cues to teach a humanoid robot new skills," in *Proceedings of the ICRA2008 Workshop on Social Interaction with Intelligent Indoor Robots*, 2008.