

Towards robots with teleological action and language understanding

Britta Wrede, Katharina Rohlfing, Jochen Steil, Sebastian Wrede, Pierre-Yves Oudeyer, Jun Tani

► To cite this version:

Britta Wrede, Katharina Rohlfing, Jochen Steil, Sebastian Wrede, Pierre-Yves Oudeyer, et al.. Towards robots with teleological action and language understanding. Humanoids 2012 Workshop on Developmental Robotics: Can developmental robotics yield human-like cognitive abilities?, Nov 2012, Osaka, Japan. hal-00788627

HAL Id: hal-00788627

<https://hal.inria.fr/hal-00788627>

Submitted on 14 Feb 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards robots with teleological action and language understanding

Britta Wrede and
Katharina Rohlfing
Citec

Bielefeld University
bwrede@techfak.uni-bielefeld.de
kjr@uni-bielefeld.de

Jochen Steil and
Sebastian Wrede
CoR-Lab

Bielefeld University
{swrede,jsteil@
cor-lab.uni-bielefeld.de}

Pierre-Yves Oudeyer
Flowers Team
INRIA, Bordeaux
pierre-yves.oudeyer@inria.fr

Jun Tani
Cognitive Neuro-Robotics Lab
KAIST
tani1216jp@gmail.com

Abstract—It is generally agreed upon that in order to achieve generalizable learning capabilities of robots they need to be able to acquire compositional structures - whether in language or in action. However, in human development the capability to perceive compositional structure only evolves at a later stage. Before the capability to understand action and language in a structured, compositional way arises, infants learn in a holistic way which enables them to interact in a socially adequate way with their social and physical environment even with very limited understanding of the world, e.g. trying to take part in games without knowing the exact rules. This capability endows them with an action production advantage which elicits corrective feedback from a tutor, thus reducing the search space of possible action interpretations tremendously. In accordance with findings from developmental psychology we argue that this holistic way is in fact a teleological representation encoding a goal-directed perception of actions facilitated through communicational frames. This observation leads to a range of consequences which need to be verified and analysed in further research. Here, we discuss two hypotheses how this can be made accessible for action learning in robots: (1) We explore the idea that the teleological approach allows some kind of highly reduced one shot learning enabling the learner to perform a meaningful, although only partially “correct” action which can then be further refined through compositional approaches. (2) We discuss the possibility to transfer the concept of “conversational frames” as recurring interaction patterns to the action domain, thus facilitating to understand the meaning of a new action. We conclude that these capabilities need to be combined with more analytical compositional learning methods in order to achieve human-like learning performance.

I. INTRODUCTION

One prominent goal in robotics is to endow robots with the capability of making sense of a situation. To achieve this, a popular approach is to build systems that perceive pieces of situation (objects, persons, their movements) first and then to combine them for a bigger picture on this situation. There is, however, increasing evidence in the developmental literature that children have the bigger picture first before they actually understand the individual pieces: Imagine a toddler who is eager to participate in a game that the older siblings are playing: S/he will not be able to understand every move in this game, let alone the proper rules but (a) s/he is definitely motivated to participate i.e. to act as the others do and (b) s/he knows that this participation is about holding

some cards, having a turn with dicing and the ultimate goal of receiving other cards. This toddler seems to have amazing skills according to which s/he analyzes the situation instantly and this enables her or him to produce a behavioral turn which will provide the child with data and solicit positive or negative input from the physical and/or social environment. This way, learning within an interaction can be stimulated [1]. We know from our previous research that disabled persons also draw from similar capabilities to engage in an interaction even if not every component is understandable [2].

Today's artificial cognitive systems lack the capabilities to actively participate in an interaction on a very low level of understanding having a very rough and preliminary idea of the goal. But we know that without such capabilities an interaction with a user cannot start or will break down soon when facing unforeseen situations.

By teleological understanding we mean that a system can apply pragmatic or contextual knowledge about the task or goals. In contrast, by compositional understanding we mean an approach that is already able to parse an action into its components thus allowing for generalisation to new actions and situations. In the action modality, Gergely and Csibra (2003) [3] suggested that young children develop a teleological understanding of actions that is based on a situational analysis. Seeing a ball moving from one side to the other, one-year-olds are sensitive to whether this ball is moving in a rational way or not. A rational way is achieved when (1) the moving of the ball brings about a future goal state, and (2) the goal state is realized by the most rational action available to the ball within the constraints of the situation ([3], p. 289). Thus, a straight path that the ball takes from one side to a particular place is better acceptable than a curved path. Such a rationality assumption (ibid: 290) refers to the infants ability to understand actions without attributing intentional mental states to others. In the language modality, some scholars (e.g. [4], [5], [6], [7]) acknowledged that children need to establish some interaction protocols that provide a general pragmatic frame for them to act appropriately without knowing every linguistic detail. Social signals such as contingency and ostension ([8], [3], [7], [9]) guide children towards understanding of such frames and result in childrens ability to participate in an

interaction, even though they might not understand every detail of it. The term “frame”, thus, refers to a multi-modal interaction structure that is already established. Fogel et al. ([5] p. 3) characterize frames as “regularly recurring patterns of communication”. For example a “labeling-frame” consists of looking at and pointing to an object and labeling it possibly with an excited facial expression and an according intonation. Thus, once the structure is established in recurrent interactions, the only new element within it is the new label for a particular objects - simplifying the understanding process tremendously.

II. TELEOLOGICAL ACTION

We understand teleology in the sense that a child (or a robot) will pick up an element of the action (such as manner or goal) to be important in his/her observation which will then be applied for the action production. Although in the following we make no explicit difference between action understanding and production we are aware that a child (or robot) may be able to understand the “goal” of an observed action (either a “manner” goal or a “target position” goal), but may not be able to reproduce this goal at this point. Also, the concrete mechanisms underlying perception and production may be different but at a higher level we do not exclude that a common representation exists.

Developmental psychology provides evidence that young infants are able to understand actions without decomposing them. Rather, they seem to apply some sort of goal-oriented interpretation enabling them to carry out an action in order to achieve that goal - even if in the first attempts they do not achieve the goal.

In our attempt to characterize a teleological action, we think that two questions need to be separated: The first question concerns the development of teleological reasoning, i.e. what kind of representation of the goal is formed when experiencing an action in a situation. Another question concerns how the experience of several actions can contribute top-down to an analysis of a new situation. We would like to exemplify these questions on recent discussion about findings from developmental studies. Gergely and colleagues [10] compared childrens imitation abilities in two different conditions. More specifically, children saw the experimenter switching on light by pressing a button with his forehead. In one condition, the experimenter had his hands put visibly free on a table. In another condition, the experimenter seemed to be cold and wrapped himself and his hands in a blanket. Children imitated the action of switching the light on with the forehead in the hands-free condition predominantly. The authors suggested that it is due to a situational analysis that the children conduct: In the hands-occupied condition, they might conclude that while the experimenter chooses the most rational action of pressing the button with his forehead (because his hands are occupied), for the children, the most rational action would be to press the button with their hands. Only in the condition, in which the action of pressing the button with a forehead seems to be a rational one, seem the children to imitate it. Recently, however, this interpretation has been challenged by

a more low-level explanation. Conducting some controlling conditions, in the study by Beisert et al. [11], the authors found that children in the hands-occupied condition might have been distracted by the unfamiliarity of the experimenters appearance. Thus, it is not due to the fact that they judge an action as rational or not. Instead, children seem to imitate in cases, in which they have the cognitive resources to do so. The unfamiliarity of the experimenter distracted the childrens attention. Accordingly, only the goal (switching the light on) was emulated, regardless of the means (the forehead). We think that Beiserts findings actually perfectly fit with the idea of a teleological action: Children perceive an action as goal-directed, and depending on their cognitive resources they can perceive (and/or produce) further details of how to achieve the goal. If their resources are limited, they will choose to reach the goal with an action they have in their repertoire.¹ This repertoire is built upon the child's experience (the second question). Thus, teleological action understanding is about formulating a quick behavioral plan and to further learn from an interaction. However, we still know little about what action experience leads to this repertoire (the first question).

A. Representation of Goals

How are goals represented and how are concepts of goals formed? In accordance with Mandler [12] we believe that the “goal” of an action is a concept inflicted upon the action by the adult teacher who has already learned that it is helpful to understand the world with respect to “goals”. Consider for example a person moving an object from point A to point B. What is the “goal” of this action? The goal could be for the object to be in the final position B. However, the goal could also be to simply remove the object from its current position A. Yet another goal could be that the object needs to be moved in order to change its state e.g. shaking to make snowflakes move in the water as in a snowdome. There is an infinite number of potential goals which an action may have. It is therefore almost impossible to derive its goal (or meaning) from one single observation. So how do infants learn the meaning or goals of actions?

Mandler [12] argues that although young infants are already able to produce predictions about ongoing actions and to behave in a goal-oriented manner, the abstract concept of a goal only forms later. Instead, she proposes basic conceptual elements with which children might perceive a motion as meaningful. More specifically, in the example outlined above, infants will most likely perceive a START and an END of this motion as meaningful units. It is still of question whether and

¹However, note that the results of this experiment can not explain if the infants did indeed not perceive the manner or if they were simply not able to incorporate it in the computation of their behavior generation. One could also argue that infants perform actions according to an optimality criterion, i.e. to produce the action with the least effort possible which depends on the manner. However, if the cognitive load is too high the infant may not be able to take a different manner (i.e. using the head) into account but rather chooses the manner that is generally accompanied with the least effort (i.e. using the hands).

how these units can be learned or are innate concepts that the children are equipped with.

Following this line of argumentation, we believe that in a very early representation of action the goal will not be explicitly represented. Rather, physical effects that occur in the scene during the action can be noticed and represented. Thus, in order to derive the invariant "goal" of an action the learner needs to look out for effects that always occur with this action. We elaborate later on how such potentially goal-bearing, meaningful effects may be perceived.²

B. Teleological processing

We assume that infants separate between goal and manner and - if under cognitive load - are only capable of reproducing one aspect (e.g. only the goal, not the manner). While at first sight this may appear as a disadvantage or compensation strategy it might well be that this resource-boundedness actually helps the infant to learn more efficiently, as it produces an action that is partially "correct" and also partially "wrong". This will provoke the tutor (or the physical environment) to give the learner a corrective feedback e.g. by pointing out the missing parts (or by reinforcing the achieved part). This is what we understand as the "Action production advantage".

III. ACTION FRAMES

When compared to similar findings in language, for action learning, the scholars have just begun to acknowledge the value of recurrent structures and how their recognition might facilitate action understanding. While in language learning, "frames" refer to an established interaction structure that is based on regularly recurring patterns of communication, in action learning the value of recognizing the recurrent structure is implicit to the assumption of a rational action. In other words, one can assume that infants will consider a movement as rational if they recognize it as a familiar movement within a particular situational configuration. We think that this aspect needs to be more recognized: The memory of events is the basis against which we perceive and predict actions. However, to date, we know little about how young infants remember structural properties of the motion stream [13] - yet this capability of making sense of actions seems to be basic [12]. Loucks et al. [13] (p. 233) suggest that "as people take action to carry out their object-directed intentions, certain physical and temporal properties predictably coalesce within the motion stream [] The time course of this sequences gives the concatenation of elements a ballistic quality, with characteristics acceleration and deceleration parameters". However, we need more evidence in developmental studies of how children's memory allow to form such elements.

²Another issue related to this is the question how the learner knows that this is the same action as before, especially if some aspects have changed. We argue that it is the communicative frame - be it verbally by specifying a label e.g. "Look, I *shake* it." or be it non-verbally by providing a repetition frame e.g. through prosody - that indicates whether the new action is the same or a different one. In the latter case it is likely that we find some kind of contrastive marker, indicating what it is that is different from the previous action (e.g. "Look, now I *don't shake* it - see, there are no snowflakes if you move it carefully").

These insights strongly suggest that we need to endow artificial systems with the capability to memorise and recognise recurrent structures. How the recognition of recurrent structures is integrated with the learning system we do not know. But we speculate that there is an intimate relationship especially in the (frequent) cases where a structure is only partially known and the unknown part needs to be learned and associated with meaning.

How can actions or the situation provide a recurrent structure that help the learner to better understand a new action? This may have different dimensions. One recurring cue may be the structure of a situation: consider for example a situation where only one agent is present (be it in the form of an agent toy such as a teddy bear or in the form of a real agent e.g. the tutor) but no manipulable object in his/her reach. This is a frame indicating that an intransitive action has to happen, where the movement of the agent him/herself is of relevance. However, if there is an object present, a transitive action can be expected, where the important thing will happen to the object. A different dimension of action frames might be related to the structure of actions itself: consider for example the process of manipulating an object. This will always involve a sequence of reaching, grasping, manipulating. From this one might condense the restriction, imposed by physics, that in order to manipulate an object it needs first to be grasped. A more sophisticated notion of such an action structure has been explored and formalised by Pastra & Aloimonos [14]. Although we can formulate hypotheses on how actions and situations may frame actions and thus render them more understandable, these hypotheses need to be carefully tested in experiments with infants.

IV. FAST LEARNING

How can we achieve the capability in an artificial system to learn some meaningful part of an action based on only one first observation? We argue that it needs to be capable of perceiving some sort of physically observable effect of the action which it will try to imitate (or emulate) based on its existing experience.

A. Goal identification

One bottom-up way for the learner to infer or detect the (potential) goal of an action is to identify recurring features in repeated action demonstrations. From a modeling perspective it is reasonable to assume that bottom-up biases exist that help the learner to identify goal-relevant features such as:

- **Goal**-position of manipulated object (the standard "goal-directed" action where the spatial goal matters e.g. when putting a cup on a saucer, or opening of a bottle where the goal is that the cap is away from the bottle)
- **Source**-position of manipulated object (e.g. move the object away)
- **Orientation** of manipulated object (e.g. turning an object upside-down)
- **Movement** itself (e.g. shaking leading to movement of snowflakes)

- **Visibility** of manipulated object (e.g. a cup that “vanishes” when nested in another cup)
- **Form** of manipulated object (e.g. open vs closed; scrunched; folded vs unfolded;)
- **Sound** (on vs off; intensity; frequency)
- **Social reaction** of a person (e.g. smiling, waving, nodding, speaking,)

A further consideration then concerns the process how such features may be observed. Again, we propose very basic strategies, such as:

- Detection of physical **effects** (for goal identification)
- Detection of **deviations** from “normal”, that is highly frequent, patterns (e.g. a normal trajectory might be a straight one between two points in space, if the trajectory is deviating from this by a cyclic shaking pattern this will be noticed and interpreted as important) (for manner identification)

In order to produce the identified goal (which might be wrong) the system will choose the most frequent means it knows in order to achieve this goal (which might be wrong as well) (note that in contrast, Csibra & Gergeley [3] argue that infants follow a rationality principle: it is most rational to move from A to B in a straight line instead of taking a deviation; it is most rational to push a button with the hands instead of the head). Instead, here we argue for a frequency effect. For example:

Goal	Means
Position of object	Movement with hand in a direct line to B
Orientation of object	Grasp object with one hand and turn it with the other hand
Visibility of object	Movement with hand towards larger object in environment

B. Carry out action

In order to achieve the action production advantage we need a system that is capable of carrying out an action regardless of its current learning state. This invites to concentrate on movements and actions that are relevant to engage in interaction and facilitate communication with a tutor on the one hand and which can be applied to a wide range of situations on the other hand. This early repertoire of action shall rely on minimal requirements for internal modeling and rather have the form of direct sensory-motor patterns that are easy to learn and to explore. An example are pointing gestures with the arm, which reference directions and can initiate and maintain much interaction, but need not be sophisticated in exactly referencing a particular object e.g. with an extended finger. Such elementary actions can easily be learned without representing exact positions in space and without an explicit kinematic model of the body in a holistic fashion [15]. This means also that there is a mechanism allowing the system to take part in an interaction e.g. by taking into account minimal turn-taking rules. The challenge is to devise respective simple but behavioral meaningful and flexible sensori-motor representations of such elementary actions that can further be specified and differentiated as development progresses.

C. Develop compositional structure

A further desideratum accruing from these considerations is the capability of the learning architecture to switch between teleological and compositional learning. While the system needs to be able to build an initial structure of an action it then needs to apply the emerged structure towards new incoming action demonstrations and parse them accordingly. This does not necessarily need to happen in an explicit way. It is possible that depending on the current representation different processing takes place implicitly.

V. COMPARISON OF COMPOSITIONAL AND TELEOLOGICAL LEARNING

Why is teleological learning not only beneficial but even necessary? Why shouldn't a learning system start with compositional learning? By teleological learning we mean a process that exploits the production advantage through continuous engagement in interaction frames and which progresses from teleological behavior to more differentiated (compositional) perception, situated action and situation understanding. We argue that it is the use of pragmatic and action frames that help the learner to bootstrap an initial action representation and to tie observed new cues to a meaning-frame, i.e. some kind of category that helps the learner to better understand. Such categories may be syntactic categories such as objects, verbs or adjectives or they may be semantic categories such as agent or patient.

We further argue that the teleological approach may achieve one-shot learning resulting in a partially correct action representation by focussing its resources on one single aspect of the action. This way the complexity of decomposing an observed action in all its parts and potential meanings is highly reduced.

The predominant approach in movement learning towards motor skills and complex actions currently is compositional. It considers movement primitives in low dimensional task spaces, i.e. elementary movement patterns that can be either goal directed [16], [17], [18] or oscillatory [19]. The respective learning methods have developed sophisticated schemes to represent and generalize the manner by encoding the movement in dynamical systems that follow learned transients towards goals and can even model the relevance of the manner according to demonstrations through probability based methods [20], [18]. These primitives are assumed to form the basic building blocks of a later to be composed complex action, however, this very process of composition is hardly considered and taken for granted. In practice the composition is far from easy to achieve [21] and typically employs very heuristic schemes [22] or needs in itself sophisticated modeling approaches [23]. The transformation from task spaces into the motor domain is solved model-based through standard inverse kinematics. Movement primitive based methods need to be complemented by approaches to segment complex actions into basic primitives, which is also far from easy e.g. using NMF [24] or sequential primitives ([25], or taking into account the already learned primitives [26]. On the other hand, it has been shown that neural networks (e.g. MTRNNs [27]) are capable

of evolving a structure that allows to learn action primitives and their sequencing in the same neural structure.

This compositional approach can neither explain how in the course of development the coordination of motor function comes about, which needs to orchestrate a highly complex body in extremely high dimensional motor spaces. Nor does it offer any route to interaction and communication and thereby mostly ignores the behavioral function of the action, as was recently observed also in [28]. More recently, first learning approaches to directly tackle the coordination problem in high dimensions towards complex sensory-motor behavior have been very successful [29], [30], [31], [32] motivated by the observation that infants also perform reaching movements from the very beginning [33]. While impressive, these approaches share the feature that they need a set of “goals” to drive learning, i.e. in motor coordination tasks a number of spatial positions to reach. These approaches provide some means to follow a more behavior and goal oriented, i.e. a more teleological approach to action learning.

Fig. 1 illustrates this restriction: in the example the target position is specified (‘B’) whereas the other potential goals (source, manner, form) of the action are not specified in the sense that they can take any value (*). Although this is generally not represented in an explicit way, it is often encoded implicitly in the learning algorithm which, for example, does not allow to switch between different types of goals.




Fig. 1. Implicit structure underlying compositional learning approaches: each entry needs to be specified, generally only one entry is specified explicitly while the others are specified implicitly through the learning algorithm.

In contrast, in the teleological approach the learner would simply select (based on which strategy ever) one “goal” and try to carry it out with the easiest or most frequent means it has at its disposal (cf. Fig. 2. If this turns out to be unsuccessful – which can be tested through interaction with the tutor – the hypothesis can be refined or corrected according to the feedback. Also, new constraints can be added if necessary.

In short, what is needed is an approach that facilitates the acquisition of structure, i.e. the discovery that in an action not only the end position matters but also the way in which it is demonstrated or executed and that constraints may exist that modulate the action such as obstacles.

Based on the structure that teleological learning has thus achieved, compositional learning strategies can be applied.




Fig. 2. Structure underlying teleological learning approaches: only one entry needs to be specified in order to enable the system to carry out the action. This yields in a corrective feedback of the tutor which can immediately be added to the teleological representation.

Note, that this is a circular relationship: at every point in time it must be possible to extend the “filled” compositional structure by a teleological approach.

VI. CONCLUSION

We have argued that infants follow a teleological action learning approach which is characterised by the ability to assign some sort of meaning to an observed action (which may be wrong) facilitated through conversational and action frames and the ability to imitate part of the action (possibly leading to the non-achievement of the goal).

This perspective on action and language learning has a range of consequences when considered as a basis for robotic learning approaches which need to be explored in future in order to better understand their functions. In this paper we have explored two potential consequences: (1) Communicational and action frames allow the learner to make use of recurrent structures, be they syntactic, semantic or pragmatic, in order to assign meaning to a newly encountered entity. Especially the concept of action frames is largely unexplored, although first attempts exist [14]. Further research needs to tackle the question how such action frames look like in more detail and how they can be exploited for learning. (2) Through the capability to extract only one part (e.g. only manner or only goal) from a single action demonstration and to reproduce it by means of extrapolation from previous experience the learner builds an action concept in an incremental way, by eliciting feedback of the tutor and thus gaining specifically tailored information to increase his/her action concept.

Both capabilities have not yet been shown by existing compositional approaches but might be achieved through different applications. We argue that in order to achieve a system that is able to learn new actions from the very first steps towards a level of expertise based on existing capabilities a combination of teleological and compositional approaches is needed.

REFERENCES

- [1] B. Wrede, K. J. Rohlfing, M. Hanheide, and G. Sagerer, *Towards Learning by Interacting*, ser. LNAI. Springer, 2009, pp. 139 – 150.

- [2] B. Wrede, S. Kopp, K. J. Rohlfing, M. Lohse, and C. Muhl, "Appropriate feedback in asymmetric interactions," *Journal of Pragmatics*, vol. 42, pp. 2369–2384, 2010.
- [3] G. Gergely and G. Csibra, "Teleological reasoning in infancy: the naive theory of rational action," *Trends in Cognitive Science*, vol. 7, no. 7, 2003.
- [4] J. S. Bruner, *Childs talk: Learning to use language*. New York: Norton, 1983.
- [5] A. Fogel, A. Garvey, H. C. Hsu, and D. West-Stroming, *Change Processes in Relationships: A Relational-Historical Research Approach*. Cambridge: Cambridge University Press, 2006.
- [6] M. Tomasello, *Constructing language*. Cambridge: A usage-based theory of language acquisition, 2003.
- [7] A. Senju and G. Csibra, "Gaze following in human infants depends on communicative signals," *Current Biology*, vol. 18, pp. 668–671, 2008.
- [8] G. Gergely and J. Watson, "Early social-emotional development: contingency perception and the social biofeedback model," *Early social cognition: understanding others in the first months of life*, pp. 101–137, 1999.
- [9] A. Cangelosi, G. Metta, G. Sagerer, S. Nolfi, C. Nehaniv, K. T. Fischer, B. J., S. T., F. G., W. L., R. B., T. K., D. E., S. K., J., and A. Zeschel, "Integration of action and language knowledge: A roadmap for developmental robotics," *IEEE Transactions on Autonomous Mental Development*, vol. 2, pp. 167–195, 2010.
- [10] G. Gergely, H. Bekkering, and I. Kiraly, "Rational imitation in preverbal infants," *Nature*, vol. 415, p. 755, 2002.
- [11] M. Beisert, N. Zmyj, R. Liepelt, F. Jung, W. Prinz, and M. M. Daum, "Rethinking rational imitation in 14-month-old infants: A perceptual distraction approach," *PLoS ONE*, vol. 7, 2012.
- [12] J. M. Mandler, "On the spatial foundations of the conceptual system and its enrichment," *In: Cognitive Science*, vol. 36, pp. 421–451, 2012.
- [13] J. T. Loucks and D. Baldwin, "When is a grasp a grasp? Characterizing some basic components of human action processing," in *Action meets word. How children learn verbs*. Oxford: Oxford University Press, K. Hirsh-Pasek and R. M. Golinkoff, Eds. 228–261, 2006.
- [14] K. Pastra and Y. Aloimonos, "The minimalist grammar of action," *Philos Trans R Soc Lond B Biol Sci*, vol. 367, no. 1585, pp. 103–17, 2012.
- [15] A. Lemme, A. Freire, G. Barreto, and J. Steil, "Kinesthetic teaching of visuomotor coordination for pointing by the humanoid robot icub," *Neurocomputing*, to appear.
- [16] S. Calinon, F. Guenter, and A. G. Billard, "On Learning, Representing, and Generalizing a Task in a Humanoid robot," *IEEE transactions on systems, man, and cybernetics. Part B, Cybernetics: a publication of the IEEE Systems, Man, and Cybernetics Society*, vol. 37, no. 2, pp. 286–98, Apr. 2007.
- [17] A. Ijspeert, J. Nakanishi, and S. Schaal, "Learning Attractor Landscapes for Learning Motor Primitives," *Advances in Neural Information Processing Systems*, pp. 1547–1554, 2003.
- [18] M. Muehlig, M. Gienger, S. Hellbach, J. J. Steil, and C. Goerick, "Task-level imitation learning using variance-based movement optimization." Kobe, Japan: IEEE, 2009, pp. 1177 – 1184.
- [19] S. Degallier and A. Ijspeert, "Modeling discrete and rhythmic movements through motor primitives: a review," *Biological cybernetics*, vol. 103, no. 4, pp. 319–338, 2010.
- [20] S. Calinon, *Robot Programming by Demonstration: A Probabilistic Approach*. EPFL Press, 2009.
- [21] B. Nemeč and A. Ude, "Action sequencing using dynamic movement primitives," *Robotica*, vol. 30, no. 5, p. 837, 2012.
- [22] R. Reinhart, A. Lemme, and J. Steil, "Representation and generalization of bi-manual skills from kinesthetic teaching," in *IEEE Humanoids*. IEEE, 2012.
- [23] D. Kulić, C. Ott, D. Lee, J. Ishikawa, and Y. Nakamura, "Incremental learning of full body motion primitives and their sequencing through human motion observation," *The International Journal of Robotics Research*, vol. 31, no. 3, pp. 330–345, 2012.
- [24] O. Mangin and P.-Y. Oudeyer, "Learning to Recognize Parallel Combinations of Human Motion Primitives with Linguistic Descriptions using Non-Negative Matrix Factorization," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2012)*, 2012.
- [25] M. Pardowitz, R. Haschke, J. J. Steil, and H. Ritter, "Gestalt-based action segmentation for robot task learning," in *IEEE-RAS 7th International Conference on Humanoid Robots (HUMANOIDS)*, 2008, pp. 347 – 352.
- [26] A. L. A. Soltoggio and J. Steil, "Primitive-based iterative decomposition of complex trajectories by maximum error in kinesthetic learning." IEEE, to appear.
- [27] Y. Yamashita and J. Tani, "Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment," *PLoS computational biology*, vol. 4, no. 11, p. e1000220, 2008.
- [28] F. Guerin, N. Kruger, and D. Kraft, "A survey of the ontogeny of tool use: from sensorimotor experience to planning," *IEEE Trans. Autonomous Mental Development*, to appear.
- [29] A. Baranes and P. Oudeyer, "Active learning of inverse models with intrinsically motivated goal exploration in robots," *Robotics and Autonomous Systems*, 2012.
- [30] —, "Intrinsically motivated goal exploration for active motor learning in robots: A case study," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*. IEEE, 2010, pp. 1766–1773.
- [31] M. Rolf, J. J. Steil, and M. Gienger, "Goal babbling permits direct learning of inverse kinematics," *IEEE Trans. Autonomous Mental Development*, vol. 2, no. 3, pp. 216 – 229, 09/2010 2010.
- [32] —, "Online goal babbling for rapid bootstrapping of inverse models in high dimensions," in *IEEE Int. Conf. Development and Learning and on Epigenetic Robotics*, IEEE. Frankfurt: IEEE, 08/2011 In Press.
- [33] V. Hofsten and C., "Action in development," *Developmental Science*, vol. 10, pp. 54–60, 2007.