Aberystwyth University

An infant inspired model of reaching for a humanoid robot

Lee, Mark; Law, James; Shaw, Patricia; Sheldon, Michael

DOI:

[10.1109/DevLrn.2012.6400805](https://doi.org/10.1109/DevLrn.2012.6400805)

Publication date:

2012

Citation for published version (APA):

Lee, M., Law, J., Shaw, P., & Sheldon, M. (2012). *An infant inspired model of reaching for a humanoid robot*. 1-6. Paper presented at 2012 IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL), San Diego, CA, USA. <https://doi.org/10.1109/DevLrn.2012.6400805>

General rights

Copyright and moral rights for the publications made accessible in the Aberystwyth Research Portal (the Institutional Repository) are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Aberystwyth Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Aberystwyth Research Portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

tel: +44 1970 62 2400
email: is@aber.ac.uk

An infant inspired model of reaching for a humanoid robot

M. Lee, J. Law, P. Shaw, and M. Sheldon

Department of Computer Science

Aberystwyth University

Aberystwyth, Ceredigion, SY23 3DB, UK

Tel: +44 (0)1970 622424, Fax: +44 (0)1970 628536

Web: <http://www.aber.ac.uk/en/cs/>

in: Development and Learning and Epigenetic Robotics (ICDL), 2012 IEEE International Conference on. See also BIB_{TeX} entry below.

BIB_{TeX} :

```
@inproceedings{6400805,  
  author={Lee, M. and Law, J. and Shaw, P. and Sheldon, M.},  
  booktitle={Development and Learning and Epigenetic Robotics (ICDL),  
    2012 IEEE International Conference on},  
  title={An infant inspired model of reaching for a humanoid robot},  
  year={2012},  
  month={nov.},  
  volume={},  
  number={},  
  pages={1-6},  
  keywords={Grasping; Learning systems; Robot kinematics; Robot sensing systems; Torso;  
    Visualization; control engineering computing; humanoid robots;  
    learning (artificial intelligence); manipulators; psychology; sensors;  
    3D space grasping; biological process; complex learning space; coordination;  
    cumulative learning; gaze control; humanoid robot; iCub robot; infant inspired model;  
    manipulation behaviour; motor system; psychological process; reaching behaviour;  
    redundancy; robotics; saccade control; sensor control; spatial transformation;  
    tabula rasa; torso control; visual feedback; visually elicited reaching;},  
  doi={10.1109/DevLrn.2012.6400805},  
  ISSN={}, }
```

© 2012 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

An Infant Inspired Model of Reaching for a Humanoid Robot

Mark Lee, James Law *Member, IEEE*, Patricia Shaw, and Michael Sheldon

Abstract—Infants demonstrate remarkable talents in learning to control their sensor and motor systems. In particular the ability to reach to objects using visual feedback requires overcoming several issues related to coordination, spatial transformations, redundancy, and complex learning spaces, that are also challenges for robotics.

The development sequence from tabula rasa to early successful reaching includes learning of saccade control, gaze control, torso control, and visually elicited reaching and grasping in 3D space. This sequence is an essential progression in the acquisition of manipulation behaviour.

In this paper we outline the biological and psychological processes behind this sequence, and describe how they can be interpreted to enable cumulative learning of reaching behaviours in robots. Our implementation on an iCub robot produces reaching and manipulation behaviours from scratch in around 2.5 hours. We show snapshots of the learning spaces during this process, and comment on how timing of stage transition impacts on learning.

I. INTRODUCTION

REACHING in humans requires the coordination of several different muscle groups controlling the shoulder, elbow, and wrist. Each of these requires relations to be made between the range of proprioceptively sensed positions and the muscle movements needed to reach those positions. Furthermore, reaching to seen objects requires the space of possible reach positions to be mapped onto the visual space perceived by the eye, but this is not straightforward as multiple arm poses may be available to reach each seen position [1].

These issues pose problems for reach-learning in humanoid robots. Multiple kinematically dependent joints create large learning spaces; visual- and joint-spaces are not topographically related, requiring some kind of transformation; redundancy creates multiple joint poses for reaching to point targets, and these cause difficulties in generating smooth reaching trajectories without discontinuities. There have been many studies and experiments on robot reaching, using both neural models and AI based methods, e.g. [2], but very few perform hand/eye coordination learning on complex kinematics in real time without prior training.

We present an approach to reach learning in humanoid robotics that draws heavily from the psychological literature and is inspired by the development and behaviour of very early infants. We identify several key factors that we consider important principles to be included in our models:

- *Motor babbling*. This is spontaneous, internally motivated, action that generates sensorimotor data during infancy. We show that it is not random activity but is functional in relating previous action to current and new

sensory-motor patterns. This has close links to the role of play behaviour.

- *Proprioception*. Proprioception develops in the pre-natal stages and, along with motor babbling, is likely to enable learning of muscle control. Proprioception develops before vision and visual guidance in reaching, and is a key factor in learning reaching motions.
- *Proximal to distal development*. Infant development follows a cephalocaudal pattern, with eye and head control appearing before arm and torso. Furthermore, upper arm control appears before forearm control and grasp learning, and this sequence has important ramifications.
- *Coarse to fine development*. Infant abilities appear at first coarse, and are refined over time. This relates to the sensory resolution and motor control abilities, as well as to the development of skills. These embedded constraints are central to developmental growth.

We view developmental sequences as the key to skill learning, and various other works show close relevance. Grunewald recognised the cephalocaudal progress of infant growth [3] and used this in skill development in robotics [4], the ITALK project has produced a robot development map [5] similar to [6], and Asada and colleagues are researching into a range of robotic models with strong emphasis on human cognitive growth [7], including the earliest stage possible; fetal development [8]. Others report on developmental approaches to reaching, including staged release [9], and experiments with proximo-distal maturation show that developmental constraints produce better learning [10].

In the following sections we will describe the development of reaching in infancy, our approach to implementing a similar sequence of development on our robotic platform, and give a series of snapshots showing the data structures built through the learning process as the robot develops reaching behaviour.

II. A DEVELOPMENTAL MODEL

In the first few months from birth infants orientate to sounds and attractive visual stimuli. They make ballistic attempts to reach towards stimulating targets but usually fail to make contact. This “pre-reaching” behaviour leads on to successful contact with objects at around 15 weeks [11], [12]. During this stage it seems that infants do not view the hand during reaching and vision is only used for target location [13]. This means that proprioception is important for arm guidance and it seems that proprioceptive development in the womb provides a more mature, although possibly incomplete, spatial framework by the time visual space is first experienced [14].

Limb movements are jerky for much this early period. The cerebellum appears to be responsible for the production of

smooth action but is very under-developed at birth. This is believed to be the cause of the marked under-damped oscillations of the arm, which gradually reduce as the cerebellum matures (over the relatively long period of 2 years).

Before 4 months there is no independent control over the fingers and grasps are formed only after contact as haptic experiences. Hand control for grasping develops later than reaching. This is an example of the cephalocaudal direction of development that is so prominent in infants [6]. It is also seen in early reaching, which involves trunk and shoulder movement, but with fingers locked. This principle of distal freezing of motor systems is an important feature and is a significant way of solving the problems associated with multiple and redundant degrees of freedom.

Only after 8 or 9 months does object size really affect approach and grasping. From this point the visually sensed object size modulates the hand aperture. Also at this age, the shift from proximal to distal control of reaching is started. It seems this is not solely due to maturational change but the trajectory of development depends heavily upon experience and patterns of behaviour [13], [15].

Another contribution to the mastery of arm control is the use of stereotypical motor patterns that have the effect of reducing the number of degrees of freedom during the early stages. By close coupling groups of muscles it is possible to reduce the number of control variables while producing a set of effective space covering actions [16]. It has also been observed that humans have a tendency to avoid extremes in arm configurations, probably because such positions considerably reduce the options for the next move. Similarly, it has been shown that people adapt their initial pose and grasp for the final arm configuration in an action task [17]. For example, subjects will choose a grasping configuration on a handle such that their hand ends up in a non-awkward position when releasing or using the object. These considerations should influence our model so that any constraining or cost function applied to reduce the DoF problem should be applied to the final configuration, not the starting configuration.

From these findings we can summarise some key points: A view of the hand is much less important when reaching than when grasping objects or other manipulations. Grasp learning follows successful reaching and involves learning object properties (affordances), finger control, tactile and other experiences. For earlier infants, who don't have much grasp control (i.e. use of fingers) proprioception may provide enough information for reaching actions.

III. EXPERIMENTAL DESIGN

We believe the developmental timeline for infants is an important tool for modeling and recognize the various constraints that pertain at different levels of development [6]. In the following sections we briefly describe how we use constraints to model the stages described in the previous section, and the resulting behaviour on the iCub. Due to space constraints the implementation details for each stage are left to the referenced papers, although the results described here are unique both in themselves and in describing the complete reach-learning process.

Following our earlier experiments, we allow eye/head coordination and eye saccading to develop independently of the construction of a proprioceptive mapping of limb space. The eye saccade learning is as described in detail in [18] and involves head movement compensation. For the growth of the proprioceptive reach space we arranged that the arm would have restricted movement on the joints for elbow and upper arm rotation, and a "rest" position was defined with the arm retracted and the hand near the head. A reach action consisted of a movement from the rest location to a specified spatial target field on or above the table surface in front of the robot. A range of target locations were generated for the volume of space around the table by motor babbling in the proprioception learning stage, (this can be done in simulation and then the locations can be transferred if motor babbling is considered unsafe on unconstrained physical hardware).

When sufficient experience has been obtained to build the gaze and reach maps the independence constraint between vision and proprioception can be lifted. This facilitates the interaction of hand and eye in behaviour known as hand regard activity. This behaviour helps by coordinating visual gaze space with the proprioceptive space of the arm/hand. Up to this point progress has been very similar to our previously described experiments [19], [20], [21].

At this stage of development the robot is able to reach to a gaze point and look at a hand position. But we notice that the gaze space is a much larger space than the reach space. This is mainly because the maximum reach is determined by the arm length which is much less than the visual range. Another important point is that the reach and gaze geometry are closely coupled in the sense that they are both grounded or referential to a point on the body centre line somewhere near the neck. This means that, regardless of the configuration of the rest of the body below the shoulders, if a stimulus is seen to be within the reachable range of the gaze/reach mappings then it can be reached. Conversely, if a stimulus is unreachable (i.e. seen but has no mapping into reach space) it can become possible to reach it by moving the head/shoulders/arm into a position where it becomes reachable. This effect can also stimulate the recruitment of locomotion to achieve distant desirable goals. However, as locomotion is not yet available, we notice that torso movement (which develops early, [11]) can be used to extend the reach space.

For the iCub, torso movements are available as tilt (forward) and rotation (about the body centre line). A torso/visual mapping can be constructed by noting the effect of torso movements on the gaze point. This process is exactly the same as the head/visual mapping which provides gaze compensation for head movements, and is described in [18]. Now, with a torso map developed, it is possible to reach to a target in a two step process: use the torso map to bring the target into a reachable location in gaze space; then use the gaze/reach mapping to generate a reach action.

The gaze space is an approximately spherical system with variables H, V and D , for left-right, up-down, and distance relative to the centre of the head. An arm configuration can be defined in terms of an n valued vector, K , for the n joint angles. Then the reach space is populated with a set of

configurations, K_i , each mapped into a gaze point, $[h_i v_i d_i]$. If motor babbling has produced a sparse but even coverage of the reach volume then we can find a K_j for an unmapped gaze point $[h_j v_j d_j]$ by interpolation between two near neighbours. Assume that $[h_1 v_1 d_1]$ and $[h_2 v_2 d_2]$ are local to $[h_j v_j d_j]$ and each are mapped, to K_1 and K_2 respectively. Then distance metrics can be computed between the vectors K_1 and K_2 and between K_1 and K_j and the resulting interpolation ratio is then applied to the elements of K_1 and K_2 to obtain a new configuration K_j on the basis of linear piecewise interpolation. If the new reach location proves to be inaccurate then its configuration can be stored, together with the mapping to $[h_i v_i d_i]$, to increase the population density of the reach space. Eventually there will be sufficient K points in the reach space that linear interpolation is effective everywhere but the space is still relatively sparse.

As described in section II, very early reaching behaviour arises before any hand control has been established and so we set the hand to be normally open with the fingers flat. If the front of the hand makes good contact with an object then an automatic finger close is executed. This provides a kind of grasp reflex which is maintained, even while the iCub performs other actions, and is only released by removal of the object, either by accident or external interaction. Unlike object contact, the release is not a significant sensed event.

As a result of the earlier hand regard behaviour the system is able to spatially correlate visual stimuli with hand positions and vice versa. Thus, when an object is presented for the first time it is likely to be detected in periphery vision and a saccade will bring the object to fixation. This fixation location in gaze space will stimulate a corresponding target for a reaching action and a reach will be initiated. At this early stage it would be expected that some reaches would miss the object and others would contact it. Some of those that make contact will also grasp the object through the grasp reflex. In accord with infant stereotypical motor patterns [22] the reach actions are completed by a return of the arm to a “home” or quiescent location in proximity to the body. (Such home positions are equivalent to the mouth, as mouthing is almost a default behaviour for any object acquired by the hand.)

After a period of early reaching, experience will have been gained on “disturbing a stimulus” (by moving it or knocking it completely out of the environment) and “holding” (with kinesthetic and possibly tactile signals). The next constraint to be lifted is the reflexive grasp and we do this by allowing the fingers to close to a given aperture and by activating a “hand empty” sensor. The hand now has potential for more control; smaller movements of the fingers can be related to visual movement or properties of objects and better grasps can be produced by matching the aperture to objects. Better approach and poise are also now within new control possibilities. Also the release of a grasped object now becomes an experienced event and so this allows objects to be dropped deliberately and thus the sophisticated skill of moving an object from one place to another is now available to learn.

In the system as described, the gaze and reach spaces record the locations of stimuli (objects) and their various properties. This is in effect a short term memory which remembers objects

during saccades and reaches but a decay function ensures that after a long period without attention such recent sensory events are erased. Consequently some form of memory is required to record actions and experiences that have proved useful and can be recalled in relevant situations. We have implemented a schema learning mechanism which provides memory and motivation functions [23]. A schema encodes the context in which an action may be performed together with the result of that action. These schemas can then be chained together to carry out sequences of actions (for example, reaching toward an object, grasping it, then moving it to a new location and finally releasing it). Schemas are selected for execution based upon an intrinsic motivation algorithm which considers the novelty of currently experienced stimuli combined with their similarity to previous experiences, resulting in actions being selected which are likely to elicit new information about the world. Example schemas are shown in the next section.

IV. EXPERIMENTAL RESULTS

Following the cephalocaudal development of the infant, the robot begins by learning the eye movements required to saccade to a visual target. Learning is conducted through our developmental framework using constraints to restrict learning of sensorimotor mappings [24]. Fig.1 shows the learnt mappings between sensor and motor spaces for making eye saccades, built up by a process of motor babbling. When a stimulus is received on the retina, the mapping between the point of stimulation and the associated motor movement is followed, triggering a saccade that fixates on the stimulus.

Next, a constraint is released enabling the learning of neck control. This could be a physical constraint, such as the lack of sufficient torque in the neck, or an emergent constraint, such as the prerequisite for accurate eye saccades as a basis for learning head movements [25]. Fig. 2 shows the learnt mapping between neck muscles and the impact of these on the visual space.

The gaze space is represented by combining the motor maps from the eye and neck system. Each field in the gaze map corresponds to a relative pan and tilt movement required to fixate on that field, and contains the eye and neck movements required to do so. When performing a gaze shift to a target, the proportion of movement allocated to each system is governed by the relationship given in [26]. Although the eye and head joints are not co-located, our experiments indicate that treating them as such gives sufficient accuracy when performing gaze shifts. Depth in the gaze space is treated separately, and is calculated by the vergence angle between the two eyes.

The ego-centric gaze space shares a reference point, the torso, with the reach space. This supports the mapping of reaches to gaze direction, but also provides a space in which to represent the robot’s environment. We use this space as a visual memory as well as for learning hand-eye coordination [21].

Reaching movements are mapped onto the gaze space using a combination of motor babbling and hand regard. Following the literature on early infant reaching, constraints are imposed on the type of reaches possible. In the early stages, the

Fig. 1. Maps for eye saccade control (best viewed in colour). Coloured circles indicate learnt fields over which stimuli or actions (in the case of the visual and motor maps) are considered identical. Matched colours indicate links, or mappings, between maps. A stimulus in a visual field will trigger the associated motor movement, causing the eye to saccade to the stimuli.

Fig. 2. Maps for head contributions to gaze shift. Motor movements are mapped to corresponding shifts in the visual space.

elbow joint is fixed, and “swiping” movements are made using the joints in the shoulder. Reaches are initiated from a “pre-reaching” pose with the hand near the head. This enables the robot to reach to objects on a line similar to the gaze direction, and limits collisions with other objects. The gaze-reach mapping is between two 3-dimensional spaces corresponding to shoulder proprioception and the gaze space. Fig. 3 shows a 2-dimensional projection of this mapping.

With constraints limiting elbow movement, the range of reach distances is very limited. The infant overcomes this by using movements of the torso to bring objects into range. Fig. 4 shows a mapping of torso rotation to a shift in gaze position. By rotating the torso the shoulder can be moved closer, or further, from objects to alter the distance for reaching. As the reach postures are mapped to vision through the gaze space, movement of the torso has no impact on eye-hand coordination.

At this stage, the robot is capable of gazing to objects, orientating itself to bring the objects into reaching distance, and making reaching motions toward them from the “pre-reaching” position. Using the schema learning mechanism it now starts to build composite actions from these beginnings.

When the robot sees an object it checks for schemas excited by that stimulus and finds that the most excited schema is one in which it remembers seeing its own hand in the location the object now occupies (Fig. 5a). Upon executing this the robot finds that when an object is present in the location it reaches its hand towards it receives an unexpected touch sensation.

Pre-conditions	Action	Post-conditions
	Reach to 35,-66	Hand at 35,-66

(a) Initially excited schema

Pre-conditions	Action	Post-conditions
Obj. 1 at 35,-66	Reach to 35,-66	Obj. 1 at 35,-66 Hand at 35,-66 Touching obj. 1

(b) Extended schema with new information

Pre-conditions	Action	Post-conditions
Obj. \$a at \$x,\$y	Reach to \$x,\$y	Obj. \$a at \$x,\$y Hand at \$x,\$y Touching obj. \$a

(c) Generalised schema

Fig. 5. The creation of a schema representing the act of touching objects

A new schema is then formed to represent this knowledge (Fig. 5b), which can then be generalised in to a form which represents reaching out and touching objects in any position (Fig. 5c).

The new touching schema is executed a number of times due to the novelty of the experiences involved. However after a short while the excitation drops below that of the next most excited schema, which in this case is the grasping schema. The grasping schema is excited by the memory of the robot touching its own hand when performing a grasp with no objects present, which it is reminded of by the touch sensation it receives from the object it has reached towards (Fig. 6a).

Fig. 3. 2 dimensional projection of reach maps in gaze and motor space

Fig. 4. Maps for torso rotation impact on the gaze direction

Pre-conditions	Action	Post-conditions
	Grasp	Touching hand

(a) Initially excited schema

Pre-conditions	Action	Post-conditions
Obj. 1 at 35,-66 Touching obj. 1	Grasp	Obj. 1 at 35,-66 Hand at 35,-66 Holding obj. 1

(b) Extended schema with new information

Pre-conditions	Action	Post-conditions
Obj. \$a at \$x,\$y Touching obj. \$a	Grasp	Obj. \$a at \$x,\$y Hand at \$x,\$y Holding obj. \$a

(c) Generalised schema

Fig. 6. The schema memory learns to go from touching to grasping objects

Executing this whilst touching an object results in the robot successfully grasping the object and receiving the sensation of holding an object. A new schema is then created to represent this new information (Fig. 6b). As with the new touching schema this grasping representation can also be generalised as shown in Fig. 6c, which represents the act of grasping an object in any location.

In the last stage of our current implementation the touching and grasping schemas can be chained together to form a plan of action which allows the robot to reach towards and then grasp an object at any location (Fig. 7).

This completes the process of attaining visually elicited reaching. Learning is driven by novelty in the early stages, giving way to goal directed behaviour only when suitable goals have been found through 'play'. The sequence shows

Pre-conditions	Action	Post-conditions
Obj. \$a at \$x,\$y	Reach to \$x,\$y	Obj. \$a at \$x,\$y Hand at \$x,\$y Touching obj. \$a

Pre-conditions	Action	Post-conditions
Obj. \$a at \$x,\$y Touching obj. \$a	Grasp	Obj. \$a at \$x,\$y Hand at \$x,\$y Holding obj. \$a

Fig. 7. Chaining of touching and grasping schemas

cumulative learning of skills from sensorimotor mapping to action planning. A key indication of the power of this approach is that the whole sequence described here can be run on the iCub robot in just 2.5 hours.

A critical issue is the scheduling of the release of constraints. In connected work we have investigated how the timing of constraint release impacts on learning of gaze control [18]. Those results showed a trade off between timing of constraint release and the rate of learning. If there are no sequencing constraints, then sub-systems are allowed to learn in parallel and learning is found to be slow, due to added physical and computational complexity. Correspondingly, connectivity between maps is sparse. If constraints remain in place for a prolonged period, learning of the unconstrained system is initially fast and connectivity is high, but at the expense of improvement in the constrained system. However, learning saturates as the space becomes increasingly explored. By releasing the constraint on a sub-system at an intermediate

time, learning of mappings in both systems is increased. Preliminary results suggest that the optimal time to release a constraint to maximise learning depends on the interaction of the codependent learning rates of the systems involved. This is a matter for further investigation.

V. CONCLUSIONS

We have described the nature of development of reaching in human infancy, and how stages in the development provide useful insights for learning to reach in humanoid robotics. We have taken these ideas and implemented them on a robotic platform using our framework for developmental learning. Results show snapshots of the sensorimotor mappings and schemas learnt along the developmental trajectory in a cephalocaudal manner from making eye saccades to reaching and grasping. The work shows the value of several principles we draw from the developmental literature.

Motor babbling is a key element in learning. The limited abilities of the infant mean that goal-driven learning is absent or restricted, and intrinsic activity, in the form of motor babbling, plays a significant role in early development. But babbling, which has close links to play behaviour, is more than random behaviour, generating vital sensorimotor data and rehearsing prior action and experience.

Proprioceptive space is an important and under-rated perceptual substrate in early learning. Before vision has developed sufficiently, proprioception provides the main feedback on limb positioning. This allows limb movements to be learnt, to some extent, prenatally. Once vision has matured, motions learnt proprioceptively can be refined with visual feedback.

Certain abilities, such as gaze control, must be refined before others, such as reaching. This is manifest in the cephalocaudal sequence of development. Furthermore, constraining distal joints until control over proximal ones has been learnt, structures the learning task. In this case, by restricting motion at the elbow joint the robot is able to learn shoulder control with a straightforward mapping.

The resolution of sensor and motor abilities in the infant are initially coarse, and gradually become finer with neural development and learning. Viewing this phenomena in terms of constraints allows us to reflect the developmental trajectory in robotic systems. As the robot masters coarse abilities, relevant constraints can be lifted to allow further refinement.

Using these lessons from human development we have built a robotic system that learns to reach in a way that overcomes many of the hurdles to humanoid reaching. Our experiments can be seen to expose some of the “logic” that appears to be behind the infant’s development in early sensory-motor learning. We believe this continued approach will offer further valuable models and solutions for robotics.

REFERENCES

- [1] O. Sporns and G. Edelman, “Solving Bernstein’s Problem: A Proposal for the Development of Coordinated Movement by Selection,” *Child Development*, vol. 64, no. 4, pp. 960–981, 1993.
- [2] F. Nori, L. Natale, G. Sandini, and G. Metta, “Autonomous learning of 3D reaching in a humanoid robot,” in *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*. IEEE, 2007, pp. 1142–1147.
- [3] R. Grupen, “A developmental organization for robot behavior,” Dept Of Computer Science, Massachusetts University, Amherst., DTIC Document ADA439115, 2005.
- [4] J. Coelho, J. Piater, and R. Grupen, “Developing haptic and visual perceptual categories for reaching and grasping with a humanoid robot,” *Robotics and Autonomous Systems*, vol. 37, no. 2, pp. 195–218, 2001.
- [5] A. Cangelosi, G. Metta, G. Sagerer, S. Nolfi, C. Nehaniv, K. Fischer, J. Tani, T. Belpaeme, G. Sandini, F. Nori *et al.*, “Integration of action and language knowledge: A roadmap for developmental robotics,” *Autonomous Mental Development, IEEE Transactions on*, vol. 2, no. 3, pp. 167–195, 2010.
- [6] J. Law, M. Lee, M. Hülse, and A. Tomassetti, “The infant development timeline and its application to robot shaping,” *Adaptive Behaviour*, vol. 19, pp. 335–358, 2011.
- [7] M. Asada, K. Hosoda, Y. Kuniyoshi, H. Ishiguro, T. Inui, Y. Yoshikawa, M. Ogino, and C. Yoshida, “Cognitive developmental robotics: a survey,” *Autonomous Mental Development, IEEE Transactions on*, vol. 1, no. 1, pp. 12–34, 2009.
- [8] H. Mori and Y. Kuniyoshi, “A human fetus development simulation: Self-organization of behaviors through tactile sensation,” in *Development and Learning (ICDL), 2010 IEEE 9th International Conference on*. IEEE, 2010, pp. 82–87.
- [9] P. Savastano and S. Nolfi, “Incremental learning in a 14 dof simulated icub robot: Modeling infant reach/grasp development,” *Biomimetic and Biohybrid Systems*, pp. 250–261, 2012.
- [10] S. Ramirez-Contró, A. Cangelosi, and D. Marocco, “Developing motor skills for reaching by progressively unlocking degrees of freedom on the icub humanoid robot,” in *Proceedings of the Post-Graduate Conference on Robotics and Development of Cognition*, 2012.
- [11] M. R. Fiorentino, *A Basis for Sensorimotor Development, Normal and Abnormal: The Influence of Primitive, Postural Reflexes on the Development and Distribution of Tone*. Springfield, Ill: Charles C Thomas, 1981.
- [12] C. von Hofsten, “Prospective control: A basic aspect of action development,” *Human Development*, vol. 36, pp. 253–270, Sep. 1993.
- [13] N. E. Berthier and R. L. Carrico, “Visual information and object size in infant reaching,” *Infant Behavior and Development*, vol. 33, no. 4, pp. 555 – 566, 2010.
- [14] B. Hopkins, *Prenatal development of postnatal functions*. Praeger Pub Text, 2005.
- [15] C. von Hofsten and L. Rönqvist, “Preparation for grasping an object: A developmental study,” *Journal of Experimental Psychology: Human Perception and Performance*, vol. 14, pp. 610–621, 1988.
- [16] M. Latash, *Synergy*. Oxford University Press, USA, 2008.
- [17] D. Rosenbaum, E. Halloran, and R. Cohen, “Grasping movement plans,” *Psychonomic bulletin & review*, vol. 13, no. 5, pp. 918–922, 2006.
- [18] J. Law, P. Shaw, and M. Lee, “A biologically constrained architecture for developmental learning of eye-head gaze control on a humanoid robot,” *Autonomous Robots*, Under review.
- [19] M. H. Lee, Q. Meng, and F. Chao, “Developmental learning for autonomous robots,” *Robotics and Autonomous Systems*, vol. 55, no. 9, pp. 750–759, 2007.
- [20] F. Chao, M. Lee, and J. Lee, “A developmental algorithm for ocular-motor coordination,” *Robotics and Autonomous Systems*, vol. 58, pp. 239–248, 2010.
- [21] M. Hülse, S. McBride, J. Law, and M. Lee, “Integration of active vision and reaching from a developmental robotics perspective,” *IEEE Transactions on Autonomous Mental Development*, vol. 4, no. 2, pp. 355–367, 2010.
- [22] N. E. Berthier, R. K. Clifton, D. D. McCall, and D. J. Robin, “Proximodistal structure of early reaching in human infants,” *Experimental Brain Research*, vol. 127, pp. 259–269, 1999.
- [23] M. T. Sheldon and M. Lee, “PSchema: A developmental schema learning framework for embodied agents,” in *Proceedings of the IEEE Joint International Conference on Development and Learning, and Epigenetic Robotics*. Frankfurt, Germany: IEEE, 2011.
- [24] M. Lee, Q. Meng, and F. Chao, “Staged competence learning in developmental robotics,” *Adaptive Behaviour*, vol. 15, no. 3, pp. 241–255, 2007.
- [25] P. Shaw, J. Law, and M. Lee, “A comparison of learning strategies for biologically constrained development of gaze control on an icub robot,” *Autonomous Robots*, Under review.
- [26] X. Wang and J. Jin, “A quantitative analysis for decomposing visual signal of the gaze displacement,” in *Proceedings of the Pan-Sydney area workshop on Visual information processing - Volume 11 (VIP ’01)*, D. D. Feng, J. Jin, P. Eades, and H. Yan, Eds., vol. 11. Darlinghurst, Australia: Australian Computer Society, Inc., 2001, pp. 153–159.