

Introduction: The Fourth International Workshop on Epigenetic Robotics

Luc Berthouze

Neuroscience Research Institute
Tsukuba AIST Central 2, Umezono 1-1-1
Tsukuba 305-8568 Japan
Luc.Berthouze@aist.go.jp

Giorgio Metta

LIRA-Lab, DIST, University of Genova
Viale Causa, 13
16145 Genova, Italy
pasa@liralab.it

1. Introduction

As in the previous editions, this workshop is trying to be a forum for multi-disciplinary research ranging from developmental psychology to neural sciences (in its widest sense) and robotics including computational studies. This is a two-fold aim of, on the one hand, understanding the brain through engineering embodied systems and, on the other hand, building artificial epigenetic systems. *Epigenetic* contains in its meaning the idea that we are interested in studying development through interaction with the environment. This idea entails the embodiment of the system, the situatedness in the environment, and of course a prolonged period of postnatal development when this interaction can actually take place. This is still a relatively new endeavor although the seeds of the *developmental robotics* community were already in the air since the nineties (Berthouze and Kuniyoshi, 1998; Metta et al., 1999; Brooks et al., 1999; Breazeal, 2000; Kozima and Zlatev, 2000). A few had the intuition – see Lungarella et al. (2003) for a comprehensive review – that, intelligence could not be possibly engineered simply by copying systems that are “ready made” but rather that the development of the system fills a major role. This integration of disciplines raises the important issue of learning on the multiple scales of developmental time, that is, how to build systems that eventually can learn in any environment rather than program them for a specific environment. On the other hand, the hope is that robotics might become a new tool for brain science similarly to what simulation and modeling have become for the study of the motor system. Our community is still pretty much evolving and “under construction” and for this reason, we tried to encourage submissions from the psychology community. Additionally, we invited four neuroscientists and no roboticists for the keynote lectures. We received a record number of submissions (more than 50), and given the overall size and duration of the workshop together with our desire to maintain a single-track format, we had to be more selective than

ever in the review process (a 20% acceptance rate on full papers). This is, if not an index of quality, at least an index of the interest that gravitates around this still new discipline.

2. Invited speakers

In a very insightful contribution, **Nadel** focuses on the key role of imitation in the development of a sense of agency – *a sense of being the owner of one’s own action*. In particular, she discusses how shared motor representations (i.e., between action generation, action simulation, action recognition, and action imitation) relate to neonatal imitation. Remarking that early imitation is already selective of human action (an idea reminiscent of **Gergely**’s keynote in last year’s workshop), she puts forward two important perceptual-motor couplings: perceptual-motor couplings as primitives of imitation, and perceptual-motor couplings as results of being imitated. These would lead to two classes of perceptions – perceptions resulting from one’s own actions, and perceptions that one cannot modify – which Russell (1996) hypothesized as being at the origin of the sense of agency. *Reciprocal imitation* – when infant and mother reciprocate imitation – is found to start early after birth, and may contribute to fill the seemingly huge gap between recognizing actions, and coding messages with communicative intent. In fact, experiments with autistic children show that repeated imitative sessions improve imitation, recognition of being imitated, and non verbal communication. Because robots possibly meet what autistic children can accept as social environment, robotic systems such as *Robota* could be used to stimulate their perception-action couplings, and may eventually lead them toward acceptance of human presence and further social use of imitation.

The next three invited speakers all deal with motor development. **von Hofsten**, very rightfully, comments that “motor development has all too often been considered as a set of milestones with little significance for the psychology of the child” (Hof-

sten, 2004). Instead, recent evidence on both developmental and adult studies (see **Fadiga**'s presentation) show how much the motor system shapes high cognitive abilities such as language. A new "picture" of the brain is being delineated in which cognition and motor control form a continuum. This idea may have underlined the dynamical system approach to the study of development by Thelen and Smith (1998). **von Oftsen** clearly addresses questions very relevant to Epigenetic Robotics by insisting on the concept of embodiment, on the dynamics of the interaction with the environment, and on the prospective control of movement. It emerges that development determines increasingly more powerful prediction abilities, and cognition could be possibly seen as acquiring the ability to "reason" and imagine the future course of events.

Movements are organized in "actions", and perceptual development can then be understood from "the action capabilities of the child and what objects and events afford in the context of those actions". In his talk, **Fadiga** asks questions related to the neural basis of motor control. In particular, experiments using different techniques (TMS and fMRI to name a few) are presented, supporting a "motor theory of perception". One example is language: already years ago, Liberman and Mattingly (1985) proposed a theory in which the elemental components of language are motor acts. It is quite significant that, after the discovery of mirror neurons in the monkey, we now have experiments that link motor responses of speech to speech perception (Fadiga et al., 2002). Besides, mirror neurons have been linked to the development of imitation and to learning by imitation in many computational works [see past editions of this workshop]. Mirror neurons may have been one of the most influential topics in robotics and certainly also in developmental robotics.

Finally, **Konczak** addresses key issues in the emergence of voluntary motor behavior. He takes us through the developmental timeline of reaching in infants (from a week to 2-3 years) to highlight some key events in the co-development of the nervous and motor systems: early pre-programmed ballistic reaching motions, parallel development of the approach and transport phases of the reach, exploitation of external forces resulting in more economical movements. The author then reports on a study in which subjects had to perform goal-directed forearm movements against an assisting viscous force. Aimed to demonstrate that the nervous system uses a neural representation of the inverse arm dynamics to control upper extremity motion, this study also showed that even in children that successfully engage in goal-directed activities (between 6 and 10 years of age), the neural representations of their arm dynamics still lacked precision and stability. Finally, Konczak dis-

cusses the issue of *time-criticality*, namely, that the organism undergoes critical periods of development during which the nervous system *expects* certain sensory inputs. The deprivation of such stimuli will negatively affect later sensorimotor (and even cognitive) development.

3. Regular contributions

3.1 Long papers

3.1.1 Development of perceptual features

This year, a few papers are devoted to the development of perceptual features (mainly visual but also multi-modal). We believe that this is important for the Epigenetic Robotics community because the development of cognition is deeply intertwined with the development of perception. Lately, a more "ecological" approach to the study of perception led to some progress, in particular, through the integration of multiple cues (visual-auditory) as well as motoric ones (Metta and Fitzpatrick, 2003). It is interesting to note that such re-interpretation of certain problems in robotics has been accompanied also by new discoveries in neuroscience such as the "mirror neurons" (Gallese et al., 1996), new theories on the relationship between motor control and language (Rizzolatti and Arbib, 1998), and object permanence (Graziano et al., 1997a).

In his paper, **Arsenio** develops a framework for learning about objects inspired by Mahler's theory of child development (Mahler, 1979), in a dynamical and complex way, through and with the support of a human caregiver. The role of the caregiver is seen as providing useful constraints and structuring information for learning. As far as we are aware, the author shows for the first time how the caregiver/teacher can actively show objects to the robot, which then adds them into its object "store" for future recognition. Most interesting is the possibility of showing a sort of cross-modal transfer by learning certain object features not by directly providing the robot with segmented data but rather by using natural cues for introducing concepts to the robot (e.g., geometric shapes).

Reflecting recent findings in neuroscience showing how deeply intertwined and multimodal brain areas are, **Fitzpatrick and Arsenio** use the matching of cross-modal cues to segment objects. It is important to stress that the (good) segmentation of sensory data is paramount to the acquisition of an appropriate "training set" for learning about the world. As in the previous paper, the authors show how different cues can be reliably integrated in learning about objects presented to the robot by humans as well as for the robot to learn about its own body. In the latter case, the robot employs proprioception into the same

cross-modal matching procedure. Neurophysiology has shown that similar links between bodily and sensory information can be found in the brain (Graziano et al., 1997b), and psychophysics has shown that our sense of self is very robust to many kind of transformations (e.g. spatial distortions) but latencies. The feeling of our body parts fails if we are tricked in situations where there are large delays between motor commands and visual feedback. A similar argument is made in experiments by **Nadel and Henning and Striano**.

The paper by **Driancourt** deals with the development of low-level (at least for now) features and their grouping following biologically-plausible principles. The author sets out the goal of building a complete system and rightfully addresses the problem of constructing a sound perceptual system. The experiments presented here do not deal with the actual control of movement but show how the first layers of processing, grouping, and feature extraction can be carried out through an efficient implementation similar in spirit to Grossberg's ART models. The paper presents links to the biological processes of cortical column formation, and simple and complex cell development in the visual cortex. As in other approaches, the author shows that the statistics of data, such as joint entropy, can drive suitable learning processes.

Finally, **Olsson, Nehaniv and Polani** deal, in an information theoretic manner, with the question of which information should be extracted given a set of sensors. This question has been often discussed in animal science, especially in regard to the visual system. Animals reared in environments with orientationally restricted contours seem to develop selective visual fields, in a sort of adaptation to the ecological niche. The authors use an information metric to compute information distances between sensors. Clusters of sensors are formed when the informational distance between them is small. The authors experiment with an AIBO robot which is placed in a contour oriented environment – consisting mainly of vertical contours – and subsequently in an unconstrained environment. Their results show a process of *unfolding* of the sensors: while sensors are initially clustered – the informational distance between sensors in a same horizontal position is zero – they eventually get separated in the metric projection when they start to distinguish between different information (i.e., unconstrained environment). Given that the environment of most mammals is found to include mostly horizontal and vertical contours, this mechanism could explain the data obtained in animal research. From an engineering point of view, this approach could be used to optimize the information gain of a given set of sensors for a given environment.

3.1.2 Social interaction

Kaplan and Hafner propose a survey on the development of joint attention. They reviewed both the developmental psychology literature and the robotic implementations showing the current state of affair with respect to the implementation of a hypothetical "intentional stance" a la Dennet. In particular, they show a series of challenges (attention detection, manipulation, social coordination, and intentional stance) to the construction of robotic artifacts that show sophisticated joint attention abilities. They set also precise criteria on what to consider truly joint attention and list situations that might be mistaken for joint attention, in particular, when considering robotic implementations. The authors identify in the discussion also which questions and potential directions are still open for research. Interestingly this line of research lead directly to addressing the problem of imitation and human-robot interaction already considered as crucial by many in our community (Demiris and Johnson, 2003; Nehaniv and Dautenhahn, 2002).

Dominey and Boucher delve into the investigation of language acquisition (and the relationship between sentence structure and meaning) from video and audio recordings. They follow a constructive approach and try to see to what extent the acquisition of language in an embodied system requires "highly pre-specified" organs as in (Chomsky, 1995). Also interesting, is the use of events in the form of object collisions, movements, and their spatial relationships in general.

Finally, **Marom and Hayes** tackle the difficult problem of establishing the hypothetical t_0 of our models: i.e. they ask how much innate knowledge should be inserted versus the amount of learning and development by interaction with the environment. They analyze the performance of learning systems as the quantity of initial design effort (innate) and social interaction is varied in a 2D design space. Social interaction is further subdivided depending to the amount of manipulation the teacher "communicate" in interacting with the robot. Working within this design space, the authors show how experiments with various amount of innate vs. acquired effort can lead to similar results: e.g. the more effort devoted by the teacher compensates for the lack of initial abstraction. This paper is an incipit into the direction of formalizing and studying the performances of certain learning systems where complex social interaction is considered.

3.1.3 Motor learning

Yamamoto and Fujinami nicely complement **Konczak's** contribution on motor learning and discuss the concept of *differentiation* – the decompo-

sition of movement into several sections to be executed in different phases – as a follow-up to *synchronisation*, a concept they discussed in their contribution to last year’s workshop. To illustrate their point, they studied two full-body motions (kneading, and samba shaking) and provide a detailed analysis of the changes in phase relationships as skill increases. While synchronisation can be explained in terms of entrainment between central pattern generators (CPG), differentiation would require temporal differentiation of oscillators. This could be achieved through the modulation of the CPGs through sensory input. The critical issue is then to find controlling points such that stability isn’t lost. The zeros of angular momentum or joint torque are suggested as suitable candidates.

3.1.4 Cognitive modeling

Balkenius and Björne’s contribution focuses on cognitive modeling, and the issue of attention in normal and autistic children. Taking the stance that a model of autistic children should have its basis in a model of normal cognitive development, they construct a simple cognitive model and show that autistic performance can result from the simple parameter modification of a normal functioning model. To deal with a task first described by Akshoomoff and Courchesne (1992) in which both normal and autistic children were tested on a task involving mixed visual and auditory stimuli with forced attention shifts, Balkenius and Björne constructed a simple model with three components: a contextQ system that learns association between stimuli and response based on reinforcement, a context module that controls in what context each stimulus-response association should be used, and an automation system that learns to produce stimulus-triggered contextual shifts. Their model successfully replicates human data, with differences between normal and autistic children accounted for by the variation of a parameter describing the influence of the automation system on the context.

Chen and Weng deal with *object permanence* – the understanding that objects continue to exist, even when they cannot directly be perceived. Using their robot SAIL, they revisit the *drawbridge* experiment of Baillargeon et al. (1985) in which infants’ looking times are observed to be longer in so-called “impossible” events. As discussed by **Schlesinger** in the last two editions of this workshop, these experiments raise the questions of whether these longer looking times result from an understanding of the physics of the system, or simply from pure perceptual processes. Similar to Schlesinger, Chen and Weng believe in a perceptual perspective and constructed a biologically-inspired computational model with a

task-independent developmental program to implement a general computational theory for novelty detection and novelty-based value system. They created 12 subjects by developing SAIL under different environments and performed a group-wise statistical analysis on the looking times. Their results, consistent with those of Baillargeon suggest that visual preference for novelty may explain the phenomenon, and thus be considered the origin of the conceptual knowledge of object permanence, rather than the outcome.

Finally, **Prince, Hollich, Helder, Mislivec, Reddy, Salunke and Memon**’s contribution deals with the issue of *contingency* or *synchrony* which were important points in **Gergely**’s keynote speech last year as well as **Nadel**’s contribution this year. The authors’ aim is to propose a formal perceptual-level model of a mechanism which has been linked to a vast array of critical cognitive developments. To measure synchrony in audio-visual information, they used an algorithm by Hershey and Movellan (2000) – where synchrony is defined as Gaussian mutual information – and extended it to estimate the degree of synchrony. With the aim of assessing whether this general-purpose synchrony detection mechanism could account for infant synchrony detection, they applied their algorithm to four tasks of increasing complexity: (a) punctuate visual movements with synchronous audio presentations of a word, (b) continuous visual movements of a face with associated speech stream, (c) continuous visual movements with irrelevant speech stream and (d) substitution of the face continuous visual movements by visual movements of an oscilloscope. Comparing the performance of their algorithm with results from infant studies, they found that while the model generally performed well, there were a few notable differences with human data, for example, in locating an audio source when faced with two motion sources. Nonetheless, the authors suggest that this study open interesting perspectives in the study of *agency*, perhaps through reciprocal imitation as suggested by **Nadel**.

3.2 Short papers

Fitzpatrick explores some processes such as the merging of concepts and the construction of causal relationships. Starting from certain building blocks, his system, implemented on the humanoid robot Cog (recently retired), builds higher level perceptual components to represent simple tasks (and their causal structure). This concept of clustering and aggregation is recurrent with the paper by **Driancourt** and allows by following certain “developmental rules” to achieve increasingly complex behaviors or perceptual features. Interestingly, this paper uses some modules whose primitive for learning are based on action it-

self, i.e., by acting, the robot learns about the objects it encounters.

D’Este addresses the question of ”sharing meaning with machines”, which is of practical importance in building machines we can communicate with. Interestingly, the author reports of the grounding of language through touch (in deaf-blind children) which is an area pretty much neglected in robotics, or at least, where the robots’ sensory systems are more lacking (resolution, quality of the sensors, etc.). She analyzes various levels of sharing that allow the transfer of meaning between dissimilar embodiments: sharing concept, attention, mind, etc. Some of these ”shared” abilities are seen as necessarily to be in place to support communication.

Oudeyer and Kaplan ask the question of how to design a controller that allows self-development (”active development” in their own words) of a robotic agent beyond what usually possible by manual construction of the developmental sequence (as we have seen in the previous papers). They call the algorithm Intelligent Adaptive Curiosity since it would bear the functions of driving exploration and ”intelligently” bringing the robot in appropriate parts of the state space where there is room for learning. The algorithm has some resemblance to the problem of exploration in reinforcement learning and of other models of value systems. They show a simulation experiment to validate the approach which is eventually based on the maximization of the learning progress (defined after an opportune measure).

Yoshikawa, Hosoda and Asada address the issue of *binding* – the correspondence of sensations between different modalities. They suggest that learning a multimodal representation of the body would be a good first step for robots because morphological constraints in self-body observation could make binding tractable. The issue, however, is to make sure that the system can match its foci of visual and tactile attention. The authors propose a model to learn the *cross-anchoring* between double-touching – the touching of one’s body with one’s body parts – and self-occlusion – the covering of a body part by another. Morphological constraints are used to drive the learning of this mapping through Hebbian learning. The effectiveness of the method is verified with computer simulations.

Weber, Elshaw, Zochios and Wermeter’s contribution deals with imitation. With respect to **Demiris and Johnson’s** contribution of last year (Demiris and Johnson, 2003) they add a language component to achieve *learning by imitation*. With an aim of modeling the multimodal nature of the mirror neuron system MNS – shown by others to bridge action observation, generation and language (see Rizzolatti and Arbib, 1998, for example)– they propose a hierarchical approach combining

Kohonen self-organizing maps and Helmholtz machines to model specific aspects of learning using the MNS with regards to demonstration learning.

Gomez, Lungarella, Eggenberger Hotz, Matsushita and Pfeifer propose a concurrent developmental release of three different modalities of degrees of freedom – visual DOFs (resolution of the sensor), motor DOFs (number of joints) and neural DOFs (size of the neural area) – as a way to preserve the so-called principle of *ecological balance* during development. Their case-study on the acquisition of foveation by a robotic hand-arm-eye system shows that a *starting small* strategy – an initial reduction of all DOF and a subsequent release – leads to faster learning than when all DOFs are immediately incorporated.

Baillie’s contribution deals with issues related to extending the famous *Talking Heads* experiment by Steels and Kaplan (2002) to a more general scenario with two autonomous robots in a complex and unconstrained visual environment. The author discusses specific issues related to the triggering and sharing of attention, perceptual segmentation, categorization, speech recognition, etc. A Universal Robotic Body Interface (URBI) is under development and may, one day, be useful to our community.

Finally, **MacDorman, Chalodhorn, Ishiguro and Asada** discuss the tight coupling between recognition and response through a robotic implementation in which the same networks – non-linear PCA neural networks – are used for both processes, and develop by the same mechanisms. At this stage, the method is limited to periodic motion, which may be a significant limit for a future generalization of the framework.

3.3 Posters

De Vylder, Jansen, and Belpaeme present a preliminary work along the direction of the acquisition of concepts through an imitation game. This game allows, by repetitive plays, to isolate a shared repertoire of behaviors that can be used to establish (emerge) a form of proto-conceptualization. The authors point out also to the relevance of joint attention (analyzed in detail by **Kaplan and Hafner** in narrowing down the ”perceptual information flow”).

Stoytchev puts into a developmental perspective the issue of tool use: Results in animal studies suggest that a system’s ability to use tools may be an adaptation mechanism to overcome limitations imposed on its anatomy. He briefly sketches his approach to learning tool affordances through exploratory behaviors.

Henning and Striano’s contribution deal the sensitivity to interpersonal timing in very young infants, an issue also discussed by **Nadel**. The authors describe studies with 3-month old and 6-month old

infants in which they investigate the effects of a one-second delay on mother-child interaction. The authors report that infants were reliably more attentive to their mothers' image when she was interacting live compared to when she was delayed. They conclude that timing provides a cue to establish when to consider others' behavior as feedback onto ones' own, an important feature for robot that learn to communicate with other robots or humans.

Rattenberger, Poelz, Prem, Hart, Webb and Ross deal with the important issue of incrementally building capabilities of increasing complexity. In last year's introduction to the proceedings, we remarked on the apparent difficult for epigenetic robotics project to move past emergent behaviors, to *successively emergent behaviors*. The authors attempt the deal with the issue by drawing inspiration from Jerne (1973)'s immune network theory. Their experimental results show that this architecture can support the emergence of complex perceptual structures and control components through interaction with the environment.

Finally, in three closely-related posters, **Lacerda, Klintfors, Gustavsson, Lagerkvist, Marklund and Sundberg** outline an *ecological theory of language acquisition* (ETLA in brief) that views the early phases of the language acquisition process as an emergent consequence of the interaction between the infant and its linguistic environment. The apparent difficulty of acquiring the seemingly infinitely-complex world of meanings is somehow facilitated by the simplicity of stimulus available to young infants (the so-called *poverty of the stimulus* argument). Linguistic development can proceed from the continuous adaptation of the adult to the perceived linguistic development of the infant. In itself, this principle is very similar to **Pfeifer's** principle of ecological balance which **Gomez et al** also discuss in their contribution. The authors are working on a mathematical implementation of the ETLA, based on multi-sensory unsupervised hierarchical pattern recognition and general memory processes. The results of the model will be tested against experimental data from language comprehension tests presented in one of the posters.

4. Conclusions

Two main observations can be made from this year's contributions:

- Issues discussed in one workshop are being revisited/re-analyzed (e.g., object permanence, joint attention, contingency, motor learning, information theoretic approach to sensory processing) perhaps indicating that this community has come to agree on some critical issues.
- Some studies are truly capturing the essence of

what epigenetic robotics is about, with models being explicitly compared with human data in more than just superficial ways (**Balkenius, Chen, Prince**, etc.). We certainly hope this trend will confirm in the future year and we will certainly try to encourage it as much as possible.

In closing, we thank the National Institute of Communication Technology (NICT) of Japan for their generous support of this workshop. Genova University has been a patient local sponsor through the process of arranging this workshop. We also thank the program committee members for their efforts in reviewing submissions.

References

- Akshoomoff, N. and Courchesne, E. (1992). A new role for the cerebellum in cognitive operations. *Behavioral neuroscience*, 106(5):731–738.
- Baillargeon, R., Spelke, E., and Wassermann, S. (1985). Object permanence in five-month-old infants. *Cognition*, 20(3):191–208.
- Berthouze, L. and Kuniyoshi, Y. (1998). Emergence and categorization of coordinated visual behavior through embodied interaction. *Machine Learning*, 31(1/2/3):187–200.
- Breazeal, C. L. (2000). Sociable machines: Expressive social exchange between humans and robots. Unpublished PhD, MIT, Cambridge, MA.
- Brooks, R. A., Breazeal, C. L., Marjanovic, M., and Scassellati, B. (1999). The cog project: Building a humanoid robot. In *Lecture Notes in Computer Science*, volume 1562, pages 52–87. Elsevier.
- Chomsky, N. (1995). *The minimalist program*. Cambridge, MA: MIT Press.
- Demiris, Y. and Johnson, M. (2003). Distributed, predictive perception of actions: A biologically inspired robotics architecture for imitation and learning. *Connection Science*, 15(4):231–243.
- Fadiga, L., Craighero, L., Buccino, L., and Rizzolatti, G. (2002). Speech listening specifically modulates the excitability of tongue muscles: a tms study. *European Journal of Neuroscience*, 15(2):399–402.
- Gallese, V., Fadiga, L., Fogassi, L., and Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain*, 119:593–609.
- Graziano, M., Hu, X., and Gross, C. (1997a). Coding the location of objects in the dark. *Science*, 277:239–241.

- Graziano, M., Hu, X., and Gross, C. (1997b). Visuo-spatial properties of ventral premotor cortex. *Journal of Neurophysiology*, 77:2268–2292.
- Hershey, J. and Movellan, J. (2000). *Advances in Neural Information Processing Systems*, volume 12, chapter Audio-vision: Using audio-visual synchrony to locate sounds, pages 813–819. Cambridge, MA: MIT Press.
- Hofsten, C. V. (2004). An action perspective on motor development. *Trends in cognitive sciences*, 8(6):266–272.
- Jerne, N. (1973). The immune system. *Scientific American*, 229(1):52–60.
- Kozima, H. and Zlatev, J. (2000). An epigenetic approach to human-robot communication. In *Proceedings of the IEEE International Workshop on Robot and Human Communication*, Osaka, Japan.
- Lieberman, A. and Mattingly, I. (1985). The motor theory of speech perception revised. *Cognition*, 21(1):1–36.
- Lungarella, M., Metta, G., Pfeifer, R., and Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15(4):151–190.
- Mahler, M. (1979). *The selected papers of Margaret S. Mahler: Separation-Individuation*. New York: Jason Aronson Publishers, reprinted 1994 ed. vol. 2 edition.
- Metta, G. and Fitzpatrick, P. (2003). Early integration of vision and manipulation. *Adaptive behavior*, 11(2):109–128.
- Metta, G., Sandini, G., and Konczak, J. (1999). A developmental approach to visually-guided reaching in artificial systems. *Neural networks*, 12(10):1413–1427.
- Nehaniv, C. and Dautenhahn, K. (2002). *Imitation in animals and artifacts*, chapter The correspondence problem, pages 41–61. Cambridge, MA: MIT Press.
- Rizzolatti, G. and Arbib, M. (1998). Language within our grasp. *Trends in neurosciences*, 21(5):188–194.
- Russell, J. (1996). *Agency: its role in mental development*. Hove, East Sussex, UK: Erlbaum (UK) Taylor & Francis. See section: action monitoring (pp. 76 and sq.).
- Steels, L. and Kaplan, F. (2002). *Linguistic evolution through language acquisition: formal and computational models*, chapter Bootstrapping grounded word semantics, pages 53–73. Cambridge: Cambridge University Press.
- Thelen, E. and Smith, L. B. (1998). *A dynamic system approach to development of cognition and action*. Cambridge, MA: MIT Press.