

Recent trends in online learning for cognitive robotics

Jochen Steil, Heiko Wersing

2006

Preprint:

This is an accepted article published in Proceedings of the European Symposium on Neural Networks. The final authenticated version is available online at: [https://doi.org/\[DOI not available\]](https://doi.org/[DOI not available])

Recent Trends in Online Learning for Cognitive Robotics

Jochen J. Steil^{1,2}, Heiko Wersing²

1-Neuroinformatics Group, University of Bielefeld, Germany
jsteil@techfak.uni-bielefeld.de

2-Honda Research Institute Europe,
Carl-Legien-Str. 30, 63073 Offenbach,
heiko.wersing@honda-ri.de

Abstract. We present a review of recent trends in cognitive robotics that deal with online learning approaches to the acquisition of knowledge, control strategies and behaviors of a cognitive robot or agent. Along this line we focus on the topics of object recognition in cognitive vision, trajectory learning and adaptive control of multi-DOF robots, task learning from demonstration, and general developmental approaches in robotics. We argue for the relevance of online learning as a key ability for future intelligent robotic systems to allow flexible and adaptive behavior within a changing and unpredictable environment.

1 Introduction

In hard- and software we currently observe technological breakthroughs towards cognitive agents, which will soon incorporate a mixture of miniaturized sensors, cameras, multi-DOF robots, and large data storage, together with sophisticated artificial cognitive functions. Such technologies might culminate in the widespread application of humanoid robots for entertainment and house-care, in health-care assistant systems, or advanced human-computer interfaces for multi-modal navigation in high-dimensional data spaces. Making such technologies easily accessible for every day use is essential for their acceptance by users and customers. At all levels for such systems *learning* will be an essential ingredient to meet the challenges in engineering, system development, and system integration. Neural network methods are of crucial importance in this arena.

Cognitive robots are meant to behave in the real world and to interact smoothly with their users and the environment. While off-line learning is well established to implement basis modules of such systems and many learning methods work well in toy domains, in concrete scenarios on-line adaptivity is necessary in many respects: in order to cope with the inevitable uncertainties of the real world, the limited predictability of the interaction structure, or to acquire new and to enhance preprogrammed behavior. Online learning is also a main methodological ingredient in the developmental approach to intelligent robotics, which aims at incremental progressing from simple to more and more complex behavior.

While learning is a multi-faceted phenomenon, – which is reflected in numerous different proposals how to relate and implement its various aspects–, in the

current context of online-learning the *time scale* at which learning can take place can be used to categorize approaches [44].

In many subfields of cognitive robotics at the slowest time scale learning methods are used in order to create initial system functions by off-line algorithms. Learning at this level does not involve any behavior of the robot and permits to create important initial subsystem functionalities which would have been much harder to obtain by explicit programming alone. While this level is very important initially in the construction of powerful robots, its contribution becomes frozen afterwards and is therefore not in the scope of the present paper.

At a faster scale those learning processes can be summarized where adaptive changes occur on-line, during (and based on) the actual behavior of the robot and refine its initial capabilities. The increased complexity introduced through the active behavior is then often compensated by requiring the adaptive changes to be (at least largely) local to each module, so that learning processes at this level can become “encapsulated” in a single functional module. Typical examples are subsystem calibrations like on-line color-recalibration or the refinement of control models and the respective algorithms can mostly be based on ideas of statistical learning. In this group of algorithms the majority of approaches have been developed in visual learning and will some will be reviewed in Section 2.

Finally, at the behavioral level it is the main challenge to make rapid “one-shot” learning feasible. This cannot rely exclusively on slow and repeated adaptations; instead, this level has to rapidly coordinate adapted subsystem functionalities in very structured, situation-specific and cross-modular ways. Clearly, coming up with working learning mechanisms at this “situated” level poses a significant research challenge. The notion of imitation learning has emerged as a very promising paradigm to cope with this challenge and some recent approaches will be covered in Section 3.

A very important issue in all online learning approaches is the way the results of learning are stored. Many motor control architectures use learning to change the parameterization of basis behaviors and therefore only implicitly store the learning result. Other approaches employ primitive graphical mappings, hash tables, or – more sophisticated– associative neural mappings to store co-occurrence of sensory inputs and motor outputs for later reuse. The interplay between memory, association and reward is investigated in a model of conditioned learning with latent inhibition by Gomond and Salotti([15] in this issue). Models that use a more complex representation during learning have to find a compromise for the classical stability-plasticity dilemma, to control the tradeoff between learning flexibility and generalization. A common approach that is clearly motivated from biological learning models is the separation into short-term memory (STM) and long-term memory (LTM). This has been applied to online learning for visual object representations [24] and models of word acquisition [38]. Duro et al. ([7] in this issue) propose a two-level memory architecture with STM and LTM using evolutionary methods to organize the transfer between the levels and apply this to tasks of robot navigation.

2 Online Learning in Cognitive Vision

One of the main problems of online learning in cognitive vision is the natural high dimensionality of visual sensorial input. This poses a challenge to most current learning architectures, since generalization is difficult to achieve in high-dimensional spaces. Another problem is that most established trainable classifier architectures like e.g. multi layer perceptrons (MLP) or support vector machines (SVM) do not allow online training with the same performance as for offline batch training. Consequently, only few work has been done in the recent time on online learning of complex visual stimuli like real objects in natural environments, if we compare this to the large body of work on model-free object recognition architectures that can be trained offline. The capability of online learning makes a fundamental difference to offline learning, since it enables an interactive process between teacher and learner. The immediate feedback on the current learning state can induce an iterative and active learning process that reduces the amount of training data that has to be presented and allows an iterative error correction based on user feedback.

To make online learning feasible, the complexity of the sensorial input has been reduced to simple blob-like stimuli [34, 23, 11], for which only positions are tracked. Based on the positions, interactive and online learning of behavior patterns in response to these blob stimuli can be performed [23]. In the following we will focus on recent work that has tried to extend the framework for online learning to the actual learning and discrimination of complex objects with more visual structure.

Garcia et al. [13] have applied the coupling of an attention system using features like color, motion, and disparity with a fast learning of visual structure for simple colored geometrical shapes like balls, pyramids, and cubes. Shape is represented as low-resolution feature maps computed based on convolutions with Gaussian partial derivatives. Based on the shape and feature map representation the system can learn and direct attention to particular objects.

Another approach to tackle the problem of high dimensionality of visual object representations are histogram-based methods. Steels and Belpaeme [42] have studied the dynamics of learning shared object concepts based on color histograms in an interaction scenario with a dog robot. Roy and Pentland [38] have investigated a computational model of unsupervised word acquisition that learns directly from raw multi-modal sensory input. The visual object representation is based on multidimensional receptive field histograms [41] for shape representation and color histograms. The learning proceeds online by using a short-term memory for identifying re-occurring pairs of acoustic and visual sensory data, that are then passed to a long-term representation of extracted audiovisual objects.

Arsenio [2] has proposed a developmental learning approach for humanoid robots based on an interactive object segmentation model that can use both external movements of objects by a human and internally generated movements of objects by a robot manipulator. Using a combination of tracking and segmenta-



Fig. 1: Presentation scenario for the online learning model of Kirstein et al. The right side shows candidate regions that are passed to the ASDF segmentation stage of Götting et al. and then used to build up an incremental object representation. The system can discriminate freely rotated complex shapes, even if they have similar color structure.

tion algorithms the system is capable of online learning of objects by storing them using a geometric hashing [35] representation. Based on a similarity threshold, objects are separated into different classes. An earlier version using only color histograms [2] was later extended by shape features [3], that use the comparison of all pairs of oriented edges in the segmented objects. From his experiments it seems, however, that the discriminatory power is limited to a small number of objects and still strongly depends on the color histogram representation. What is more important is his integration of the online object learning into a model for tracking objects and learning task sequences and to recognize objects employed on such tasks from human-robot interaction.

An interesting approach to supervised online learning for object recognition was proposed by Bekel et al. [6]. Their VPL classifier consists of three major stages. First the input is separated into raw appearance clusters using a vector quantization method, and in the second stage a local principal component analysis is performed for dimension reduction. The final stage is a supervised classifier using a local linear map architecture [36]. The image acquisition of new object views is triggered by pointing gestures on a table, and is followed by a short training phase, which takes some minutes. Currently the main drawback is the lack of an incremental learning mechanism to avoid the complete retraining of the architecture. The recognition model has been integrated into a cognitive vision system that integrates a wide variety of visual functions like localization, object tracking/recognition and action recognition [52].

Kirstein et al. [24] have developed an approach for the supervised online learning of object representations based on a biologically motivated architecture of visual processing. They use the output of a recently developed topographical feature hierarchy [51] to provide an appearance-based representation of three-dimensional objects using a dynamical vector quantization approach. Unlike most other rapid learning approaches to object recognition this architecture is not based on a dimension reduction principle, but uses a series of hierarchical

feature detection and spatial integration stages, similar to the ventral visual pathway of humans. The output is a topographical feature map representation that achieves a robust response to modest transformations of the input like rotation, scaling, and translation. The efficiency of the representation is achieved by sparse coding that ensures that object views are represented using only sparse activation in the high-dimensional feature space [51]. The output of the feature detection stage is passed to an online learning memory model that consists of a rapidly learning short-term memory (STM) and a long-term memory (LTM) for later consolidation of object representations. The STM is realized as an incremental vector quantization model that adds object view representatives based on a threshold on the Euclidean similarity to already stored templates. Using this STM, the architecture is capable of online learning of 50 complex objects within three hours with an average classification error of 5%. In these experiments objects were freely presented by hand in arbitrary rotations, but the background was constrained to be black in combination with a black glove to facilitate a simple entropy-based segmentation of the objects. Based on the current object view content of the STM, the LTM can be trained continuously, albeit on a slower time scale, using an incremental adaptation of a learning vector quantization (LVQ) training algorithm. As was shown in [24], online learning is also feasible for a realistic setting of a limited STM, that can only hold recent information of the last 10 objects. By using an appropriate incremental node addition procedure and temporally changing learning rates for the object representatives in the LTM, the system can handle the stability-plasticity dilemma quite efficiently and does not suffer from catastrophic forgetting problems like standard MLP or SVM architectures.

Götting et al. ([16], in this issue) have developed the adaptive scene-dependent filter (ASDF) segmentation approach to extend the online learning object recognition architecture by Kirstein et al. [24] to a less constrained scenario that would also be the natural interaction with a humanoid robot like ASIMO [18]. The setting is given as a human teacher presenting objects in front of a stereo camera system, without any further constraints on the environment or the clothing of the teacher (see Figure 1). Using stereo processing, the focus is centered on the shown object with the concept of peripersonal space [14] in the proximity of the camera head. Based on the raw distance map around the candidate object region, a relevance map is computed that covers the object only coarsely. For each pixel location in the candidate region, a local feature vector is computed based on RGB color channels, distance, and local orientation. Using a dynamic vector quantization model [17], first an unsupervised segmentation is computed using the local feature vectors and then object segments are chosen by the overlap relation of the found segments to the relevance map. The segmentation is sufficiently fast to allow a processing frame rate of 6 frames per second for the complete preprocessing and object learning system on a system of two dual processor desktop computers. Using the ASDF approach, the learning and recognition performance shows almost no degradation compared to the more constrained setting shown in [24], and learning of 50 objects is possible with a

classification error less than 5%. The complete online learning recognition system has recently been integrated into a cognitive vision system architecture using also speech recognition for a particularly intuitive interactive training dialogue during object presentation.

3 Learning by demonstration and imitation

In practice, all current cognitive robot architectures implemented in real robots rely on hand-wired designs of the control flow to assure a proper sequencing of behaviors. Often such flows can be formalized in form of state machines, which trigger functional modules with respect to a given overall state of the robot. Learning on the architectural level in this context takes an evolutionary pathway: it consists of memorizing action patterns, which are defined to be reasonable by the designer, rather than changing the connectivity between the elementary functional units. However, with the number of modules and states of the system the number of potential state-transitions increases exponentially and the functional interdependencies quickly become intractable. Even when using states and state-machine concepts, it is difficult to ensure stability on the transition level, which means that the system always is in a defined state and recovers to a reasonable behavior. Reconfiguration of the control flow then is a tedious endeavor always at risk of complete functional failure of the system.

To make learning in such situations feasible, the approach of imitation learning is very appealing [5, 1, 10, 27, 33], in particular for humanoid robots with many degrees of freedom and potentially very complex behavior [40]. The basic idea is to find a “template” for a successful behavior by observation of a (human) instructor. This requires to endow the robot system with sufficient perceptive capabilities to visually acquire the action to imitate; to transfer the observed action into an internal representation, which accounts as well for the system’s parameters and copes with the different accessibility of sensor data and the possibly different “instrumentation” with actuators and finally to be able to physically execute a suitable action by an actuator [44]. One early approach to formalize this linking between perception and motor action has been provided in [40] and further pursued e.g. in [39].

Learning along these lines aims at memorizing successful action paths realizable within the given control flows by monitoring the instructor, because systematic exploration in the real world is practically impossible. It requires a representation to store the imitated behavior for reuse, while expecting that lower level flexibility and robustness will allow to apply it later in similar situations. This representation might be implicit in parameter settings, however, in order to learn multiple behaviors, classification of situations and a reference to an appropriate behavior needs some explicit memorizing mechanism. Though the concept of imitation appears in principle very appealing, concrete realizations require complex behavioral systems and the availability of many partial sub-skills. Therefore many authors have concentrated on partial imitation in trajectory learning, reaching movements, or other particular behaviors like grasping.

3.1 Trajectory learning and adaptive control of multi-DOF robots

While for classical standard fabrication robots standard tools from control theory are available to drive the robot along pre-specified paths, the control of the often highly redundant complex kinematics for humanoid robots is by far more difficult. In particular the problem to coordinate different DOFs to generate goal directed and smooth movements have motivated many authors to use imitation of human movements as basic templates. All these approaches rely on recording a number of reference trajectories, defining or training suitable movement primitives offline, and on achieving online capabilities by changing the parameterization of the movement generating dynamic system according to the observable current context. For implementing the basis behaviors, Ijspeert et al. use nonlinear differential equations parameterized by target coordinates for goal reaching movements [20], or employ parameterizable oscillator networks for generating cyclic movements [19]. Billard and Mataric use specifically tailored recurrent networks inspired by biological mirror systems to imitate arm movements [9]. Though the movement primitives in the aforementioned references are defined beforehand, the respective robots can react with a reasonable and flexible suitable behavior to online observed input. Recently, D'Souza et al. [48] have also demonstrated for the 41 DOF ATR humanoid robot that it is possible to learn a complete inverse dynamics model truly online while behaving and to use this in adaptive control for tracking a reference behavior. The basis for this approach is an online learning scheme, which acquires incrementally low dimensional local projection models which are combined into a locally weighted output sum.

A further learning model of robot behavior has been developed by Tani [45] based on a partially recurrent neural network with parameterized bias (PBRNN), which is trained conventionally with backpropagation to associate perceptive inputs to desired motor outputs by evaluating the prediction error for perception. The PBRNN has for instance been used to imitate cyclic human arm movements [22], to connect speech input with motor behavior [22], or to generate joint attention in a visually driven imitation game [21].

3.2 Task learning from demonstration

Most of the earlier imitation approaches were devoted to learning by demonstration [12], [25], [47], aiming to reproduce reaching behavior and grasping with simple grippers. In these models, the imitation concentrates to provide a desired sequencing of basic sub-skills to achieve an observed target behavior on the task level. In [44, 28, 43], a more sophisticated system including modules for visual attention, speech recognition, integration of visual and linguistic input, scene memory and anthropomorphic grasping for instructing a robot to grasp every day objects has been presented to provide an architecture for imitation grasping. Only recently more systematic approaches to extract the important information about the task to be solved by imitation have been formulated using an optimization framework [8] or symbolic terms [53].

4 The developmental approach

A different though overlapping source of ideas leading to online learning has recently emerged under the paradigm of developmental learning [49],[4],[50]. It combines earlier approaches to learn partial skills like gaze tracking, hand-eye coordination, visual tracking of objects, or sensori-motor mapping for simple manipulation with motivation from developmental psychology to provide approaches for incremental scaling of robotic architectures from simple to complex behavior. A pronounced approach in this direction has been sketched in [26] as developmental pathway which progresses from a level of self-awareness defined by capabilities as eye vergence control and arm-head coordination over world awareness given by visually initiated reaching and control of grasp to imitation very much in the sense described above. Different authors have concentrated on steps along this pathway like joint attention and gaze direction [31], [46], basic sensori-motor control [30] or linking vision and motor control for reaching and grasping [32], [29]. An interesting example of online learning in developmental approaches to vision is the adaptation of the mapping from visual to motor space for learning the gaze control of a saccading active camera head [37]. The ability of online learning allows a rapid autonomous recalibration of the visual system to different optical parameters, and can also cope with a complete mirroring of the the up-down direction induced by a prism.

5 Perspectives for cognitive learning architectures

Online learning in cognitive robot architectures is a complex yet increasingly important topic. It gains even more attention with the advent of humanoid multi DOF robots which interact with humans in a multi-modal fashion. Looking at biological systems in comparison to contemporary robot control architectures, the former consist of interconnected loops stabilized by numerous mechanisms of error-tolerance implemented by adaptivity, self-repair, and default fallback behaviors while the latter lack exactly this robustness. In biology, it is often difficult to distinguish functionally divided modules properly and any attempt to modularize function and structure devoted to it is quite an idealization. This suggests that a major challenge to be met for approaching higher complexity in cognitive robots is to successfully adopt methods of dealing with systems that we cannot analyze in their full detail and in closed form. Learning architectures can provide a key to scale robot systems to a reasonable complexity allowing for smooth interaction. Though on the system level systematic investigation of learning architectures is very difficult and comparative evaluation in interaction is a completely open issue, there is a lot of progress in the subfields of imitation learning on the trajectory level, joint attention, learning of reaching and grasping, and visual acquisition of objects and scene models. On the architectural level, a great source of inspiration are the biological mirror systems in the motor cortex, which inspire memory systems and associative networks in various ways. The most difficult step towards truly cognitive architectures currently

seems to be the integration of the various powerful methods available in flexible architectures into a unified system. This step should enable situated learning on the system level and in this way go beyond partial skill acquisition of simple imitation behaviors.

References

- [1] P. Andry, P. Gaussier, S. Moga, J. P. Banquet, and J. Nadel. Learning and communication via imitation: An autonomous robot perspective. *IEEE SMC*, 31(5):431–442, 2001.
- [2] A. Arsenio. *Cognitive-developmental learning for a humanoid robot: a caregiver's gift*. PhD thesis, MIT, 2004. AI Technical Report 2004-006.
- [3] A. Arsenio. Developmental learning on a humanoid robot. In *Proc. Int. Joint Conf. Neur. Netw. 2004, Budapest*, pages 3167–3172, 2004.
- [4] M. Asada, K. F. MacDorman, H. Ishiguro, and Y. Juniyoshi. Cognitive developmental robotics as a new paradigm for the design of humanoid robots. *Robotics and Autonomous Systems*, 37:185–193, 2001.
- [5] P. Bakker and Y. Kuniyoshi. Robot see, robot do : An overview of robot imitation. In *Proceedings of the AISB workshop on Learning in Robots and Animals*, pages 3–11, Brighton, UK: The Society for the Study of Artificial Intelligence and the Simulation of Behaviour, 1996.
- [6] H. Bekel, I. Bax, G. Heidemann, and H. Ritter. Adaptive computer vision: Online learning for object recognition. In *German Pattern Recognition Symposium*, pages 447–454, 2004.
- [7] F. Bellas, J.A. Becerra, and R.J. Duro. Results with a two level memory management system in the multilevel darwinist brain. In *Proc. ESANN*, 2006.
- [8] A. Billard, Y. Epars, S. Calinon, G. Cheng, and S. Schaal. Discovering optimal imitation strategies. *Robotics and Autonomous Systems*, 47(2-3), 2004.
- [9] A. Billard and M. Mataric. Learning human arm movements by imitation: Evaluation of a biologically-inspired connectionist architecture. *Robotics and Autonomous Systems*, 941:1–16, 2001.
- [10] A. Billard and M. J. Mataric. A biologically inspired robotic model for learning by imitation. In *Proc. 4. Int. Conf. on Autonomous agents*, pages 373–380, Barcelona, 2000.
- [11] Cynthia Breazeal and Brian Scassellati. A context-dependent attention system for a social robot. In *Proc. 16th Int. Joint Conf. on Artif. Intell. (IJCAI-99-Vol2)*, pages 1146–1153, 1999.
- [12] R. Dillmann, M. Kaiser, and A. Ude. Acquisition of elementary robot skills from human demonstration. In *In International Symposium on Intelligent Robotic Systems*, pages 185–192, Pisa, Italy, 1995.
- [13] Luiz-Marcos Garcia, Antonio A. F. Oliveira, Roderic A. Grupen, David S. Wheeler, and Andrew H. Fagg. Tracing patterns and attention: Humanoid robot cognition. *IEEE Intelligent Systems*, 15(4):70–77, 2000.
- [14] C. Goerick, H. Wersing, I. Mikhailova, and M. Dunn. Peripersonal space and object recognition for humanoids. In *Proceedings of the IEEE/RSJ International Conference on Humanoid Robots (Humanoids 2005), Tsukuba, Japan*, 2005.
- [15] N. Gomond and J.-M. Salotti. Extended model of conditioned learning within latent inhibition. In *Proc. ESANN*, 2006.
- [16] M. Götting, J.J. Steil, H. Wersing, E. Körner, and H. Ritter. Adaptive scene-dependent filters in online learning environments. In *Proceedings Eur. Symp. Neur. Netw. ESANN, Bruges*, 2006.
- [17] G. Heidemann and H. Ritter. Efficient vector quantization using the WTA-rule with activity equalization. *Neural Processing Letters*, 13(1):17–30, 2001.

- [18] M. Hirose, Y. Haikawa, T. Takenaka, and K. Hirai. Development of humanoid robot ASIMO. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Maui, Hawaii, 2001.
- [19] A. J. Ijspeert, J. Nakanishi, and S. Schaal. Learning rhythmic movements by demonstration using nonlinear oscillators. In *IEEE International Conference on Intelligent Robots and Systems*, pages 958–963, 2001.
- [20] A. J. Ijspeert, J. Nakanishi, and S. Schaal. Trajectory formation for imitation with nonlinear dynamical systems. In *Proc. IROS*, pages 752–757, 2001.
- [21] M. Ito and J. Tani. Joint attention between a humanoid robot and users in imitation game. In *Proc. 3rd Int. Conf. on Development and Learning*, 2004.
- [22] M. Ito and J. Tani. On-line imitative interaction with a humanoid robot using a dynamic neural network model of a mirror system. *Adaptive Behavior*, 12(2):93–115, 2004.
- [23] Tony Jebara and Alex Pentland. Action reaction learning: Automatic visual analysis and synthesis of interactive behaviour. In *Int. Conf. Computer Vision Systems*, 1999.
- [24] S. KIRSTEIN, H. WERSING, and E. KÖRNER. Rapid online learning of objects in a biologically motivated recognition architecture. In *27th Pattern Recognition Symposium DAGM*, pages 301–308. Springer, 2005.
- [25] Y. Kuniyoshi, M. Inaba, and H. Inoue. Learning by watching: extracting reusable task knowledge from visual observation of human performance. *IEEE Trans. on Robotics and Automation*, 10(6):799–822, 1994.
- [26] Manuel Lopes, Jos'e Bernardino, and Jos'e Santos-Victor. A developmental roadmap for task learning by imitation in humanoid robots. In *AISB - Third International Symposium on Imitation in Animals and Artifacts*, Hatfield, Uk., 2005.
- [27] Maja J. Mataric, Odest C. Jenkins, Ajo Fod, and Victor Zordan. Control and imitation in humanoids. In *AAAI Fall Symposium on Simulating Human Agents*, North Falmouth, MA, Nov. 3-5 2000.
- [28] P. McGuire, J. Fritsch, H. Ritter, J. J. Steil, F. Röthling, G. A. Fink, S. Wachsmuth, and G. Sagerer. Multi-modal human-machine communication for instructing robot grasping tasks. In *Proc. IROS*, pages 1082–1089. IEEE, 2002.
- [29] G. Metta and P. Fitzpatrick. Better vision through manipulation. *Adaptive Behavior*, 11:109, 2003.
- [30] Giorgio Metta, Giulio Sandini, Lorenzo Natale, and Francesco Panerai. Sensorimotor interaction in a developing robot. In *First International Workshop on Epigenetic Robotics*, 2001.
- [31] Yukie Nagai. *Understanding the Development of Joint Attention from a Viewpoint of Cognitive Developmental Robotics*. PhD thesis, Osaka University, 2004.
- [32] Lorenzo Natale. *Linking Action to Perception in a Humanoid Robot: a Developmental Approach to Grasping*. PhD thesis, University of Genova, 2004.
- [33] Monica M. Nicolescu and Maja Mataric. Experience-based representation construction: Learning from human and robot teachers. In *PROC IROS*, pages 740–745, 2001.
- [34] Monica N. Nicolescu and Maja J. Mataric. Learning and interacting in human-robot domains. *IEEE Trans. Systems, Man, and Cybernetics*, 31(5):419–430, 2001.
- [35] I. Rigoutsos and H. Wolfson. Geometric hashing: An overview. *CSAE: Computational Science & Engineering, IEEE Computer Society*, 4:10–21, 1997.
- [36] H. Ritter, T. Martinetz, and K. Schulten. *Neural Computation and Self-Organizing Maps: An Introduction*. Addison-Wesley, Reading, MA, 1992.
- [37] Tobias Rodemann, Frank Joubin, and Edgar Körner. Saccade adaptation on a 2 dof camera head. In *Third Workshop on SelfOrganization of Adaptive Behavior (SOAVE 2004) Ilmenau*, pages 94–103, 2004.

- [38] Deb Roy and Alex Pentland. Learning words from sights and sounds: a computational model. *Cognitive Science*, 26(1):113–146, 2002.
- [39] S. Schaal, A. Ijspeert, and A. Billard. Computational approaches to motor learning by imitation. *Transaction of the Royal Society of London: Series B, Biological Sciences*, 358:537–547, 2003.
- [40] Stefan Schaal. Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, 3(6):233–242, 1999.
- [41] Bernt Schiele and James L. Crowley. Recognition without correspondence using multidimensional receptive field histograms. *International Journal of Computer Vision*, 36(1):31–50, 2000.
- [42] L. Steels and F. Kaplan. AIBO’s first words. The social learning of language and meaning. *Evolution of Communication*, 4(1):3–32, 2001.
- [43] Jochen J. Steil, Frank Röthling, Robert Haschke, and Helge Ritter. Learning issues in a multi-modal robot-instruction scenario. In *Proc. IROS*, volume Workshop on ”Robot Programming Through Demonstration”, Oct 2003.
- [44] Jochen J. Steil, Frank Röthling, Robert Haschke, and Helge Ritter. Situated robot learning for multi-modal instruction and imitation of grasping. *Robotics and Autonomous Systems*, Special Issue on ”Robot Learning by Demonstration”(47):129–141, 2004.
- [45] J. Tani. Learning to generate articulated behavior through the bottom-up and the top-down interaction processes. *Neural Networks*, 16(1):11–23, 2003.
- [46] J. Triesch, C. Teuscher, G. Deak, and E. Carlson. Gaze following: why (not) learn it? *Developmental Science*, 9(2):125–147, 2006.
- [47] J. Triesch, J. Wiegardt, E. Maël, and C. v. d. Malsburg. Towards imitation learning of grasping movements by an autonomous robot. In *Gesture-Based Communication in Human-Computer Interaction*, volume 1739 of *LNAI*, pages 73–84. Springer, 1999.
- [48] Sethu Vijayakumar, Aaron D’Souza, and Stefan Schaal. Incremental online learning in high dimensions. *Neural Computation*, 17(12), 2005.
- [49] J. Weng. The developmental approach to intelligent robots. In *Integrating Robotic Research: Taking The Next Leap*, AAAI Spring Symposium Series, 1998.
- [50] J. Weng. Developmental robotics: Theory and experiments. *International Journal of Humanoid Robotics*, pages 199–236, 2004.
- [51] H. Wersing and E. Körner. Learning optimized features for hierarchical models of invariant recognition. *Neural Computation*, 15(7):1559–1588, 2003.
- [52] Sebastian Wrede, Marc Hanheide, Sven Wachsmuth, and Gerhard Sagerer. Integration and coordination in a cognitive vision system. In *Int. Conf. on Computer Vision Systems*, 2006.
- [53] R. Zöllner and R. Dillmann. Using multiple probabilistic hypothesis for programming one and two hand manipulation by demonstration. In *IROS Workshop on programming by demonstration*, 2003.