Goal-directed imitation for robots: A bio-inspired approach to action understanding and skill learning[☆]

W. Erlhagen^{a,*}, A. Mukovskiy^a, E. Bicho^b, G. Panin^c, C. Kiss^c, A. Knoll^c, H. van Schie^d,
H. Bekkering^d

^a Department of Mathematics for Science and Technology, University of Minho, 4800-058 Guimarães, Portugal

^b Department of Industrial Electronics, University of Minho, 4800-058 Guimarães, Portugal

^c Informatics, Chair for Robotics and Embedded Systems, Technical University Munich, 85748 Garching, Germany

^d Nijmegen Institute for Cognition and Information, Radboud University Nijmegen, 6500 HE Nijmegen, The Netherlands

Available online 30 March 2006

Abstract

In this paper we present a robot control architecture for learning by imitation which takes inspiration from recent discoveries in action observation/execution experiments with humans and other primates. The architecture implements two basic processing principles: (1) imitation is primarily directed toward reproducing the outcome of an observed action sequence rather than reproducing the exact action means, and (2) the required capacity to understand the motor intention of another agent is based on motor simulation. The control architecture is validated in a robot system imitating in a goal-directed manner a grasping and placing sequence displayed by a human model. During imitation, skill transfer occurs by learning and representing appropriate goal-directed sequences of motor primitives. The robustness of the goal-directed organization of the controller is tested in the presence of incomplete visual information and changes in environmental constraints.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Imitation learning; Dynamic field; Mirror neurons; Action sequence

1. Introduction

There has been a growing interest in creating autonomous robots which are capable of developing motor and cognitive skills through real-time interactions with their environment. Recent research in movement learning suggests that trying to imitate an experienced teacher (e.g., a human) is a powerful means of speeding up the learning process (for reviews see [1,2]). This form of social learning in robots is not restricted to movements. It may be complemented by acquiring more abstract knowledge such as, for instance, structurally new motor behaviors composed of a set of parameterized motor primitives.

In this paper we summarize results of an interdisciplinary project which aimed at exploring new ways of imitation and

imitation learning in artefacts based on recent discoveries in cognitive psychology and neuroscience. The basic idea was to get new insights into the relevant functional mechanisms underlying imitation from behavioral and neuronal data. Central questions for robot imitation that we have addressed in our work concern “what to imitate” and how to solve the correspondence problem across dissimilar embodiments and task constraints [3]. Very often these differences simply do not allow for a matching on the level of movement trajectory or path. In the goal-directed theory of imitation proposed by Bekkering and colleagues [4,5] imitative behavior can be considered successful whenever the goal of an action in terms of the desired outcome of the movement is reproduced. The action means, on the other hand, may or may not coincide with the observed ones. The focus on the consequences of the movement requires that the imitator understands the demonstrator’s behavior as an intentional motor act directed at a specific goal (e.g., placing an object at a certain position). The “matching hypothesis” forwarded by Rizzolatti

[☆] This work was supported by the European grants *ArteSmit* (IST-2000-29686) and *JAST* (IST-2-003747-IP).

* Corresponding author. Fax: +351 253510401.

E-mail address: wolfram.erlhagen@mct.uminho.pt (W. Erlhagen).

Fig. 1. Bridge paradigm. The robot has to imitate a human grasping an object and placing it at one of the two targets behind the bridge obstacle.

and colleagues [6] based on their discovery of the mirror system states that an action is understood if its observation activates the motor representations controlling the execution of a similar goal-directed action (“motor simulation”). The proposed controller implements action understanding and goal-directed imitation as a continuous process which combines sensory evidence, contextual information, and a goal-directed mapping of action observation onto action execution. As a theoretical framework, we base our implementation work on dynamic fields [7] previously used to endow autonomous robots with certain cognitive capabilities (e.g., memory, decision making, [8–10]).

The complete control architecture including vision, cognition and path planning is validated in variations of a paradigm in which the robot system learns to imitate a grasping and placing sequence displayed by a human model. The learning is accompanied by structural changes in the controller representing knowledge transferred from the human to the robot during imitation.

2. Experimental paradigm

To test the idea of a goal-directed organization of imitative behavior in a robot-human task we adopted a paradigm which has been originally developed for experiments with human subjects (van Schie and Bekkering, submitted for publication [22]). The paradigm contains an object that must be grasped and then placed at one of two laterally presented targets that differ in height. Importantly, the grasping and transporting behaviors are constrained by an obstacle in form of a bridge (see Fig. 1). Depending on the height of the bridge, the lower target may only be reached by grasping the object with a full grip and transporting it below the bridge. Placing the object at the higher target, on the other hand, may require combining a precision grip and a hand trajectory above the bridge.

The robot had to reproduce the observed or inferred action consequence (placed object). The work was conducted on a robot platform consisting of an industrial 6-degrees-of-freedom robot arm (KUKA, Germany) on which a four-fingered anthropomorphic robot hand (GRAALTECH, University of Genova, Italy) was mounted. A real-time vision system provided the information about the scene parameters and the human hand motion.

Fig. 2. Schematic diagram of the control architecture. The central part responsible for the selection of means and goals in the imitation task reflects recent neurophysiological findings in 4 interconnected brain areas. The functionality of the STS-PF-F5 pathway is to match action observation and action execution. The matching is controlled by the goal representations in area PFC.

3. The control architecture

Three interconnected modules (vision, cognition, path planning) define the robot control architecture (see Fig. 2).

3.1. Vision module

The *vision module* provides the environmental variables of the task setting (Cartesian position of bridge, object and goals) by means of a semi-automatic calibrated stereo camera system. All outputs are stored in the general configuration structure, globally available for the other modules of the controller. The demonstrator's hand and the object are identified and tracked in real time on the basis of a chroma-space blob segmentation in the YUV color space using a monocular camera view. The hand tracking algorithm is based on a mutual information optimization approach [11] which maximizes the consistency between an observed image and postures of a hypothetical hand model (26-degrees-of-freedom). The hand trajectory (above or below the bridge) and the placing goal (high or low) are classified on the basis of a distance measure relative to the respective object. The categorization of the grasping behavior (full or precision) is essentially based on the orientation of the palm relative to the object.

3.2. Cognitive module

In the *cognitive module* decisions about the action goal and the means to achieve that goal are made. Its layered architecture is biologically inspired, as it represents the basic functionality of neuronal populations in interconnected brain areas known to be involved in action observation/execution tasks (for details see [12]). The STS-PF-F5 pathway is believed to represent the neural basis for a matching between the visual description of an action in area STS and its motor representation in area F5. Direct physiological evidence for this hypothesis comes from the existence of “mirror neurons” in F5 that become active both when the animal makes a particular action and when it observes another individual making a similar action. The fundamental idea we adopt for the robotics work is that the matching takes place on the level of motor primitives that represent complete goal-directed motor behaviors such as, for instance, “grasping an object with a full grip” ([6], see also [1] for an excellent overview in the context of robotics research). Motor primitives do not encode the fine details of the movement and thus provide a sufficiently abstract level of description for imitation learning across dissimilar embodiments. Concretely for the bridge paradigm, we distinguish two types of grasping primitives (precision grip (PG) and full grip (FG)) and two types of transporting primitives for avoiding the obstacle (below (BT) or above (AT) the bridge).

The representations in the intermediate layer PF reflect recent neurophysiological findings in brain area PF that suggest a goal-directed organization of action means. Using a grasping–placing task, Fogassi and colleagues [13] described a population of grasping mirror neurons which showed a selective response in dependence of the final goal of the action (placing vs. eating) to which the grasping act belongs. For the bridge paradigm, we abstract this finding by assuming representations of specific combinations of primitives (e.g., PG-AT) which allow achieving a specific placing goal. Possible goals parameterized by their height relative to the bridge (spatial gap, Fig. 1), are assumed to be encoded in neuronal populations of the “prefrontal area” PFC. The reciprocal connections between PFC and PF are learned during the imitation experiments. Functionally, they allow one to override a direct matching between primitives in STS and F5 if necessary. Beside the direct stimulation by the vision system (placed object), the goal representations in PFC may be influenced by two additional information sources: (i) the task input represents memorized information about the number, identity (height) and probability of goals (for details of a computational implementation see [14]). It reflects the fact that in a known task setting the robot may engage in partial motor preparation even before the observation of the human model. (ii) The second input represents object cues (e.g. color) which may become associated with the goal during imitation.

3.2.1. The dynamics of decision making and learning

Each individual layer of the cognitive module is formalized by a dynamic field. The particular form we employed was originally introduced by Amari [7] as a mathematical model

for the dynamics of pattern formation in neuronal tissue. The main idea is that the interplay between excitatory and inhibitory interactions in local populations of neurons may sustain the population activity for extended periods of time. The build-up of self-sustained activity patterns in populations encoding action goal and means may thus be seen as the process of stabilizing and maintaining the task relevant information.

In the Amari model, the activity $u(x, t)$ of a neuron at field location x at time t is governed by the following integro-differential equation:

$$\tau \frac{\delta}{\delta t} u(x, t) = -u(x, t) + h + \sum_i S_i(x, t) + f_1(u(x, t)) \left[\int w(x - x') f_2(u(x', t)) dx' - w_{\text{inhib}} \int f_2(u(x', t)) dx' \right] \quad (1)$$

where $\tau > 0$ defines the time scale of the dynamics and $h < 0$ the resting level to which the field activity relaxes without external stimulation. The non-linear functions $f_i(u)$, $i = 1, 2$; are of sigmoid shape,

$$f_i(u) = \frac{1}{1 + \exp(-\beta_i(u - \theta_i))} \quad (2)$$

with threshold $\theta_1 > \theta_2$ and slope parameter $\beta_1 = \beta_2$. The strength of the excitatory connections w to field neighbors is expressed in terms of the distance between locations, that is, $w(x, x') = w(x - x')$. We used gaussian profiles with standard deviation σ_s and amplitude A_s . The feedback inhibition depends on the overall activity in the field and is controlled by the constant $w_{\text{inhib}} > 0$. Since the excitatory process is spatially restricted, the globally inhibitory process dominates at larger distances. Finally, the term $\sum_i S_i(x, t)$ represents the summed external input to the field which consists of excitation from connected layers and input from the vision module. The latter is modelled as gaussian functions of adequate intensity.

For the present work, we have adapted the model parameters to guarantee a bistable regime of the dynamics in which a transient input may act as a switch between a homogeneous rest state and a localized activity profile [7]. In Fig. 3, Panel A, we exemplify the evolution of such a profile in response to an input to a neuronal population representing the precision grip (PG). The build-up of excitation is accompanied by an increase in lateral inhibition which causes a suppression of activation in the population representing the grip alternative (FG). There is a threshold, $u_{\text{TH}} = 0$, for triggering a self-sustained pattern. Weak external inputs (e.g., task input) may only bring the activation close to that threshold. As shown in Panel B, this “preshaping” mechanism may nevertheless drastically alter the time course of the suprathreshold response triggered by a sufficiently strong input [14]. The observed speed-up of processing may, in turn, affect the decision processes in connected layers.

Crucial for our approach to learning by imitation is that the control architecture may autonomously evolve during practice by developing new task-relevant representations. We apply a

Fig. 3. Panel A: The temporal evolution of a self-stabilized activation pattern representing a decision for a PG-grip is shown. Note that at the time of stimulus onset, $t = 0$, the field appears to be already pre-activated by the constant input from the task layer in PFC. Panel B: The time course of the maximally excited field element is compared for the case with pre-activation (solid line) and without pre-activation (dashed line).

correlation based learning rule for the synaptic connections, $a(x, y)$, between any two neurons x and y in any two model layers that is compatible with the field format [15]:

$$\tau_s \frac{\delta}{\delta t} a(x, y, t) = -a(x, y, t) + \eta f_2(\bar{u}_1(x)) f_2(\bar{u}_2(y)) \quad (3)$$

where $\eta > 0$, $\tau \ll \tau_s$ and $\bar{u}_1, \bar{u}_2$ denote the equilibrium solutions of the relaxation phase in layer 1 and layer 2, respectively. Note that a transient phase of the dynamics could have been chosen as well without changing the results of the present study. Important for establishing a goal-directed organization of the control architecture is that an internally generated reinforcement signal representing a successful path planning toward the desired goal posture (see below) defines the time window for the learning. For simplicity, we have chosen a function which takes on the value 1 during the learning period and 0 otherwise. As a result, the metric for the learning appears to be defined by the similarity in the end state of the action [3,16].

3.3. Path planning

For generating overt behavior, the abstract motor primitives represented in layer F5 have to be translated into the right kinematics. We employ a global planning method in posture space which is inspired by Mel's biologically plausible network model [17]. In the network, each of the locally interconnected nodes represents a stored posture, that is, a n -dimensional array of joint angles $\Theta_i = (\theta_{i1}, \dots, \theta_{in})^T$, $i = 1, \dots, N$, where N represents the number of network nodes covering the workspace W . It is assumed that a model for the forward kinematics of the arm/hand system exists, that is, the nodes are associated with the arm/hand location in Cartesian space. Each node Θ_i is connected with its k nearest neighbors Θ_j , $j = 1, \dots, k$ defined by the euclidean metric $d(i, j) = \|\Theta_i - \Theta_j\|_2$. The connection weights $w_{ij} = w(\Theta_i, \Theta_j)$ are assumed to decrease exponentially with euclidean distance, $w(i, j) = \exp(-d(i, j))$. Starting with an external activation of a set of goal postures P_G , activation spreads in each time step to inactive nodes by summing the excitation, v_j from the active neighbors multiplied by the respective connection weight, $v_i(t) = \sum_{j=1}^k v_j(t-1) * w(i, j)$ and $v_{P_G}(0) = 1$ (there is no interaction between already activated units). When

the wavefront reaches the node corresponding to the initial posture, the activation dynamics is stopped. The sequence of postures defining a suitable path from the initial state to the goal state is then found by back propagation in time along the maximally excited nodes. Inverse kinematics is used to define three distinct sets of goal postures P_{G^i} , $i = 1, 2, 3$, associated with the Cartesian location of the object to be grasped, X^1 , and location of the two placing targets, X^2 and X^3 , respectively. Moreover, additional information from the vision and the cognitive module of the control architecture is integrated before starting the wavefront operations [18]. Posture nodes which are impossible due to the obstacles are inhibited. The forbidden set P_O of all inhibited nodes is found by explicitly testing for spatial overlap in Cartesian space between the to-be-assumed posture and the bridge obstacle using forward kinematics. Moreover, the ensemble of nodes which can become part of the wavefront is further constrained by the motor primitives in F5. For instance, we use again forward kinematics to check whether a particular node Θ_j represents "all links of the robot arm in a high position" as required by a trajectory above the bridge. This pre-selection of a set of compatible postures, P_{F5} , restrict the global planning process to the subset $P_{F5} \cap P_O = \emptyset$.

A real-time path planning for artefacts with higher degrees of freedom is possible due to a technique known in the literature as refinement procedure (e.g., [19,20]). Starting with a distribution of posture nodes covering only sparsely the whole work space, new nodes are successively introduced whenever needed. This refinement procedure and the successive re-planning may be necessary, for instance, around an obstacle or in the vicinity of nodes already defining a possible (albeit not smooth) path.

4. Experimental results

A set of imitation experiments within the bridge paradigm has been performed which differ in the amount of visual information available to the robot and in task constraints. The aims were (1) to exemplify what kind of knowledge may be transferred from the human to the robot by autonomously developing new representations, and (2) to illustrate the advantages of a goal-directed organization of the control architecture in terms of robustness, compared to more traditional via-point matching models.

4.1. Copying the means

In the first set of experiment, a complete visual description of the teacher's actions in terms of the grasping and transporting behavior exists and the vision system identifies the placing goal. Although the robot has the knowledge how to grasp, transport and place objects in its motor repertoire, it does not know how to combine under the constraints of the bridge paradigm the specific motor primitives to achieve the same end state. One strategy could be trying to copy the primitives displayed by the human demonstrator. The visual description of the observed motions in layer STS resonates via the matching mechanism in the mirror circuit with the congruent motor representations of

Fig. 4. Copying the means. Panel A: The teacher shows a complete grasping–placing sequence, here precision grip (PG), followed by a transport above the bridge (AT). Panel B: Cognitive module. The peaks of activation in layer F5 represent the means (motor primitives) selected by the robot to reproduce the same goal. Panel C: The robot reproduces the observed placing at the lower target using the same means.

the robot. If the covert path planning toward the desired posture necessary to achieve the action goal turns out to be successful, the observed action sequence becomes associated with the goal representation in layer PFC by the learning procedure described above. Fig. 4 illustrates the result of this learning by imitation in an example in which the robot copies the demonstrated precision grip and the trajectory above the bridge to place the object at the higher goal. In the various layers of the neural field model, the task specific information is encoded by activity profiles representing a steady state of the dynamics.

4.2. Discerning motor intention

The second set of experiments has been designed to reflect a major challenge for all robotics systems cooperating in

Fig. 5. Inference task. Panel A: Only the grasping behavior is observable. Panel B: The stable state in layer PFC of the field model represents the inferred (lower) goal. Panel C: To reproduce the inferred action effect the robot combines a full grip (FG) followed by a trajectory below (BT) the bridge as represented in the motor layer F5 in Panel B.

cluttered environments with other agents. Due to occluding surfaces, for instance, only partial visual information about the action displayed by the partner may be available and the observing robot has to infer the action goal. The proposed control architecture implements the idea that a goal-directed motor simulation together with the integration of prior task knowledge underlies the capacity of discerning motor intention. Consistent with this view, it has been recently reported that mirror neurons may fire under appropriate experimental conditions (e.g., with additional contextual cues) even if the goal of the motor act is hidden from view [6]. In the concrete example shown in Fig. 5, only the demonstrator's grasping of the object with a full grip was observable. However since the robot is familiar with the task, links between goal representations and associated goal-directed sequences have been established in previous trials. In addition, the constant task input results in a pre-activation below threshold, u_{TH} , of all task-relevant representations. As a result of the robot's "expectation" about possible goals and means, the evolving activation in STS encoding the observed FG-grip is sufficient to trigger first the sequence FG-BT and subsequently the representation of the associated lower goal (Panel B in Fig. 5). As depicted in Panel C, the robot shows its action understanding by combining a full grip and a trajectory below the bridge to reproduce the inferred action effect.

Fig. 6. Goal directed imitation. Panel A: Conflict in the grasping behavior, i.e. the teacher uses a full grip for placing the object at the higher goal. Panel B: As shown in layer F5 of the field model, the robot decides to use a precision grip to reproduce the observed action effect.

4.3. Goal directed imitation

The third set of experiments illustrates that the learned link from the mirror circuit to the goal representation is crucial. The bar of the bridge is removed for the human but not for the robot (Panel A in Fig. 6). Because of this change in the environmental constraints, the demonstrator now uses a full grip for placing the object at the higher target. For the robot, a direct matching on the level of motor primitives would result in a collision with the bridge. As shown in the snapshot of the field model in Panel B of Fig. 6, the decisions in layer F5 represent the motor primitives PG and AT previously associated with the higher goal (compare Fig. 4). This choice is the result of the primacy of the goals over the means implemented in the control architecture. The goal representation is triggered by direct input from the vision system. Through the learned links to layer PF, it biases the decision processes in the mirror circuit, thus overriding the direct mappings from the visual motion description in STS. Technically, we exploit here differences in time course with the goal representation being processed faster in a known task setting compared to the representations in STS (for a detailed discussion of the biological context see [12]).

5. Discussion

We have presented a control architecture for imitation and learning by imitation which is strongly inspired by recent insights about the processing principles underlying these capabilities in humans and other primates. In general, our approach emphasizes the role of factors in imitation which are usually considered cognitive such as goal inference or decision making. The experiments with the robot system illustrate that an organization of imitative behavior toward reproducing the goal of an observed action complements purely motor approaches which focus on a matching on the trajectory or path level [2,16]. The primacy of the goal over the action

means allows coping with differences in embodiment and task constraints known as the correspondence problem in robot imitation [3]. The emphasis on “end state granularity” [3] as a measure to classify successful imitative behavior does not exclude, however, that the transfer of knowledge on the level of action means may also be essential (compare Fig. 4). Most importantly, learning to understand an observed behavior as a goal-directed action enables the robot to reuse the stored information in new contexts and to acquire more abstract knowledge associated with that action. For instance, association between specific object properties (e.g., color) and where to place an object of a particular category may be learned within the proposed control architecture (see Fig. 2).

The idea that the movement production system is essentially involved in action understanding has been proposed in the context of robotics research before (for a review see [2]). For instance, Demiris and Hayes [21] used internal forward models known from motor control theory to predict the sensory consequences of observed actions in an imitation task. However, the questions how to cope with differences in embodiment, task constraints or even motor skills have not been systematically addressed. In principle, the control architecture proposed here allows for learning to understand the purpose of a hand movement not strictly in the repertoire of the imitator. The only condition is that the observed action effect may be achieved using proper action means [12].

In the present imitation paradigm, the goal-directed sequence is still relatively simple as it combines only four motor primitives. In a joint effort with our experimental colleagues we are currently testing the learning of more complex sequences composed of a richer set of existing movement primitives (e.g., a sequence composed of two grasping and placing behaviors). Conceptually, the implementation work does not require any changes in the overall control architecture. However, some modifications in the field dynamics (Eq. (1)) and the learning dynamics (Eq. (3)) have to be introduced. First, to prevent a competition between motor primitives belonging to the same category (e.g., grasping) but to different parts of the sequence, the representations should be transient in nature. A straightforward solution is to add a “forgetting dynamics” which results in a destabilization of the self-sustained activation profiles [9]. Second, to allow for an explicit representation of the temporal order of primitives in layer PF a predictive Hebbian learning rule [15] should be used.

Acknowledgements

We would like to thank Leonardo Fogassi and Giacomo Rizzolatti for numerous discussions about this work.

References

- [1] S. Schaal, Is imitation learning the route to humanoid robots? *Trends in Cognitive Science* 3 (6) (1999) 233–242.
- [2] S. Schaal, A. Ijspeert, A. Billard, Computational approaches to motor learning by imitation, *Philosophical Transactions of Royal Society of London B* 358 (2003) 537–547.
- [3] C.L. Nehaniv, K. Dautenhahn, The correspondence problem, in: K. Dautenhahn, C.L. Nehaniv (Eds.), *Imitation in Animals and Artifacts*, MIT Press, 2002, pp. 41–61.

- [4] H. Bekkering, A. Wohlschläger, M. Gattis, Imitation of gestures in children is goal-directed, *The Quarterly Journal of Experimental Psychology* 53A (2000) 153–164.
- [5] A. Wohlschläger, M. Gattis, H. Bekkering, Action generation and action perception in imitation: An instantiation of the ideomotor principle, *Philosophical Transactions of Royal Society of London B* 358 (2003) 501–515.
- [6] G. Rizzolatti, L. Fogassi, V. Gallese, Neurophysiological mechanisms underlying the understanding and imitation of action, *Nature Reviews* 2 (2001) 661–670.
- [7] S. Amari, Dynamics of pattern formation in lateral-inhibitory type neural fields, *Biological Cybernetics* 27 (1977) 77–87.
- [8] C. Engels, G. Schöner, Dynamic fields endow behavior-based robots with representations, *Robotics and Autonomous Systems* 14 (1995) 55–77.
- [9] E. Bicho, P. Mallet, G. Schöner, Target representation on an autonomous vehicle with low-level sensors, *The International Journal of Robotics Research* 19 (5) (2000) 424–447.
- [10] M. Quoy, S. Moga, P. Gaussier, Dynamical neural networks for planning and low-level robot control, *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 33 (4) (2003) 523–532.
- [11] A. Viola, W.M. Wells III, Alignment by Maximization of Mutual Information, in: 5th IEEE International Conference on Computer Vision, Cambridge, MA, 1995, pp. 16–23.
- [12] W. Erlhagen, A. Mukovskiy, E. Bicho, A dynamic model for action understanding and goal-directed imitation, *Brain Research* 1083 (1) 175–189.
- [13] L. Fogassi, P.F. Ferrari, B. Gesierich, S. Rozzi, G. Rizzolatti, Parietal lobe: From action organization to intention understanding, *Science* 308 (2005) 662–667.
- [14] W. Erlhagen, G. Schöner, Dynamic field theory of movement preparation, *Psychological Review* 109 (2002) 545–572.
- [15] P. Dayan, L.F. Abbott, *Theoretical Neuroscience*, The MIT Press, 2001.
- [16] A. Billard, Y. Epars, S. Calinon, S. Schaal, G. Cheng, Discovering optimal imitation strategies, *Robotics and Autonomous Systems* 47 (2004) 69–77.
- [17] B.W. Mel, *Connectionist Robot Planning: A Neurally-Inspired Approach to Visually-Guided Reaching*, Academic Press, 1990.
- [18] Y. Koga, K. Kondo, J. Kuffner, J.-C. Latombe, Planning motions with intentions, in: *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, ACM Press, 1994, pp. 395–408.
- [19] L. Kavradi, J.-C. Latombe, Randomized preprocessing of configuration space for path planning: Articulated robots, in: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Munich, Germany, 1994, pp. 764–1771.
- [20] N. Roy, S. Thrun, Motion Planning through Policy Search, in: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Lausanne, Switzerland, 2002.
- [21] J. Demiris, G. Hayes, Imitation as a dual-route process featuring predictive and learning components: A biologically plausible computational model, in: K. Dautenhahn, C.L. Nehaniv (Eds.), *Imitation in Animals and Artifacts*, MIT Press, 2002, pp. 327–361.
- [22] H. van Schie, H. Bekkering, Neural mechanisms for goal-directed actions reflected by slow wave brain potentials (submitted for publication).

Wolfram Erlhagen is an Associate Professor at the Department of Mathematics for Science and Technology and director of the research centre *Officina Mathematica* at the University of Minho in Portugal. He studied Mathematics and Physics at the Universities of Bonn and Bochum in Germany, obtaining a doctoral degree in Mathematics in 1996 with work on pattern formation in active media. Before coming to Portugal he was leader of the Theoretical Biology group at the Institute for Neuroinformatics, Ruhr-University Bochum (1996–2000) and invited researcher at the CNRS in Marseille, France (1997). Dr Erlhagen's research interests cover the fields of spatially extended dynamical systems, neurocomputational and cognitive modelling, and architectures for autonomous robots. In close cooperation with experimental groups he applies his

theoretical investigations to problems of motor planning, visual perception and reasoning with the ultimate goal to bridge Cognitive Sciences to Robotics.

Albert Mukovskiy received his M.Sc. degree in Biophysics in 1992 from the Moscow State University, Russia. In 1999 he joined as a research fellow the Institute for Neuroinformatics at the Ruhr-University of Bochum, Germany. Since 2002 he is research fellow at the Department of Mathematics for Science and Technology at the University of Minho, Portugal. His research interests cover computational neuroscience, motor control, and more specifically the interaction between perception and action in imitation learning.

Estela Bicho is an Assistant Professor at the University of Minho, Portugal, where she teaches courses in Control and Robotics and heads the research group on "Intelligent Autonomous Robotics & Dynamical Systems". She received the diploma (1992) and M.Sc. (1994) degrees in Electronics and Telecommunications from the University of Aveiro, Portugal. She was then a teaching assistant at the Department of Industrial Electronics at the University of Minho, and member of the "Equipe de Dynamique" at the "Centre de Recherche en Neurosciences Cognitives", CNRS, Marseille, France, where she pursued her doctoral studies on autonomous robotics and dynamical systems. She obtained her Ph.D. degree in Robotics, Automation and Control in 1999 from the University of Minho. The Ph.D. work received the Portuguese IBM honour price (1999). Her research focuses on cognitive control architectures for single and multi-robot systems.

Giorgio Panin received the M.Sc. degree in electronics engineering and the Ph.D. degree in robotics from the University of Genoa, Italy, in 1998 and 2002, respectively. He joined the European IST Project Artesimit in 2002 as postdoctoral fellow at the Informatics Faculty of the Technical University of Munich, Germany. Since 2004 he is Research Assistant at the TU-Munich, Chair for Robotics and Embedded Systems; his current research interests include anthropomorphic robot hand control, design and implementation of advanced real-time computer vision systems and human-machine interfaces, and imitation learning systems.

Csaba Ferenc Kiss graduated as an automatization and software engineer in 1995 at the University of Oradea (Romania). He joined the European IST Project Artesimit in 2001 as a research fellow at the Informatics Faculty of the Technical University of Munich, Germany. From 2000 to 2001 he worked as Research Assistant for the Bavarian Research Centre for Knowledge-based Systems in Munich, then as a Ph.D. candidate at the Chair for Robotics and Embedded Systems at TU-Munich. He is currently finishing his dissertation in the field of computer vision and human-machine interfaces.

Alois C. Knoll received the diploma (M.Sc.) degree in Electrical Engineering from the University of Stuttgart, Germany, in 1985 and his Ph.D. (summa cum laude) in computer science from the Technical University of Berlin, Germany, in 1988. He served on the faculty of the computer science department of TU Berlin until 1993, when he qualified for teaching computer science at a university (habilitation). He then joined the Technical Faculty of the University of Bielefeld, where he was a full professor and the director of the research group Technical Informatics until 2001. In March 2001 he was appointed to the board of directors of the Fraunhofer-Institute for Autonomous Intelligent Systems. Since autumn 2001 he has been a professor of computer science at the Computer Science

Department of the Technical University of Munich. His research interests include sensor-based robotics, multi-agent systems, data fusion, adaptive systems and multimedia information retrieval. Dr. Knoll is a member of the German Society for Computer Science (Gesellschaft für Informatik (GI)) and the IEEE.

Hein van Schie is a Postdoc at the Nijmegen Institute for Cognition and Information. He received his M.Sc. in Psychology (1997), and a Ph.D. in Cognitive Neuroscience (2003) at the University of Groningen. Dr. Van Schie's research interests focuses on understanding the functional and neural mechanisms that support cognitive function in language, action, and perception, using behavioral and neuroimaging techniques.

Harold Bekkering is Professor in Cognitive Psychology at the Dutch Radboud University Nijmegen and coordinator of the research Master Cognitive Neuroscience. He is also vice-director of the research school NICI (Nijmegen Institute for Cognition and Information). He received his M.Sc. (1990) in Psychology in Nijmegen and his Ph.D. (1995) in Maastricht from the Department of Human Movement Science. During his Ph.D. he was a fellow scientist at Washington University (St. Louis, USA). In 1995 he became a senior scientist at the Max-Planck-Institute for Psychological Research in Munich (Germany). He then moved to the University of Groningen to become associate Professor in Experimental Psychology (2001) before he accepted the chair in Cognitive Psychology in Nijmegen (2002). Dr Bekkering's research interests cover the fields of Developmental and Cognitive Psychology, Cognitive Neuroscience and the implementation of Cognition in Robotics. His work concentrates on the interaction between perception, semantics and action.