

Towards Integrating Learning by Demonstration and Learning by Instruction in a Multimodal Robot

Stefan Wermter, Mark Elshaw, Cornelius
Weber, Christo Panchev, Harry Erwin
Centre for Hybrid Intelligent Systems
School of Computing and Technology
University of Sunderland
St Peter's Way, Sunderland
UK
www.his.sunderland.ac.uk

Abstract—Learning by demonstration and learning by instruction offers a potentially more powerful paradigm than programming robots directly for specific tasks. Learning in humans or primates substantially benefits from demonstration of actions or instruction by language in the appropriate context and there is initial neurocognitive cortical evidence for such processes. Cortical assemblies have been identified in the cortex that activate in response to the performance of motor tasks at a semantic level. This evidence supports that such mirror neuron assemblies are involved in actions, observing actions and communicating actions. Furthermore, neurocognitive evidence supports that cell assemblies are activated in different regions of the brain dependent on the action type being processed. Based on this neurocognitive evidence we have begun to design a neural robot in the MirrorBot project that is based on multimodal integration and topological organisation of actions using associative memory. As part of these studies in this paper we describe a self-organising model that clusters actions into different locations dependent on the body part they are associated with. In particular, we use actual sensor readings from the MIRA robot to represent semantic features of the action verbs. Furthermore, ongoing work focuses on integration of motor, vision and language representations for learning from demonstration and language instruction.

I. INTRODUCTION

Often robots are restricted in their general autonomous behaviour and only perform what has been preprogrammed. We begin to see initial research in learning by language instruction or action demonstration (e.g. Billard 2002 [2], Demiris and Hayes 2002 [5]). However, so far, demonstration and language instruction have only played a minor role in intelligent robotics. Some robots like the tour-guide robot Rhino [4] have been quite robust in terms of their localisation and navigation behaviour. However they do not use learning by demonstration or learning from language instructions. Although the conversation office robot jijo-2 [1] can be instructed to navigate to certain landmarks and the Minerva tour-guide [17] interacts by using simply preprogrammed speech, they are restricted in their ability to learn. Furthermore, the Kismet interactive robot [3] can recognise and represent emotions using a

sophisticated head but does not learn by imitation or instruction.

Learning through imitation has been a useful approach for primates and therefore is an active research into the area of learning in robots. For instance Demiris and Hayes 2002 [5] and Maistros and Hayes 2001 [10] have devised approaches based on the mirror neuron concept to achieve robot learning through imitation. Demiris and Hayes 2002 [5] use behaviour and forward models in their approach where a demonstrator robot was observed by the imitator robot performing actions and then is required to predict what is being performed. Maistros and Hayes 2001 [10] use the Scheme Theory to express the function of the mirror neurons to achieve learning by imitation of grasping actions. This was done by using as demonstrator-imitation scenario in a similar manner to that by Demiris and Hayes 2002 [5]. Billard 2002 [2] also considers the use of imitation to aid autonomous robot communication learning of a proto-language by using an unsupervised approach based on a dynamic recurrent associative memory architecture. Language is learnt through a student robot recreating the actions of the teacher robot, through the teacher robot describing what it observes and the student robot having a similar perspective to the teacher. Gaussier et al. 2001 [9] have considered the use of a neural network approach that is able to achieve learning and communication through imitation. In doing so they concentrate on low-level imitations that recreate simple movements that are found in infants.

In our approach (e.g. Wermter and Panchev 2002 [22], Wermter et al. 2001 [20], Wermter and Elshaw 2003 [21], Weber and Wermter 2003 [18]) we study learning in intelligent robots based on some evidence from the brain since it obviously supports learning from demonstration and learning from verbal instructions in humans. In this particular study here in the context of the MirrorBot project we focus on two constraints: first multimodal learning and integration of action, vision and language and second the topological arrangement of actions and

their visual and language counterparts.

First, multimodal learning, recently, a class of neurons has been found in the rostral part of the ventral premotor cortex (area F5) in monkeys that are active both when a monkey handles an object and when it observes an experimenter performing similar actions [14]. More recently, PET studies have implicated these 'mirror neurons' in the gesture recognition system of humans. This system involves Broca's area, a language area in humans, which is generally believed to be the human homologue of area F5 in monkeys. Therefore, we explore the role of mirror neurons and cell assemblies for multimodal integration of action, vision, and language in the MirrorBot project.

Second, examining the processing of action verbs that relate to the leg, face and arm Pulvermüller et al. 2000 [13] found that cell assemblies are associated through semantic information with the appropriate body part. Furthermore, it was noted by Rizzolatti et al. 2001 [16] that when subjects were required to observe actions made by the mouth, hand and foot that the foot was represented dorsally and mouth and hand ventrally in the brain. This neurocognitive topological evidence motivates our approach for self-organising associative memory in multiple regions of the brain.

In our approach neural learning of multimodal association of motor actions, vision and language will be a key element for learning robots. We understand learning by demonstration and imitation in this general sense of learning multimodal internal topological representations. In an initial experiment and architecture described in this paper we show how representation of demonstrating motor actions and language instructions can be integrated. In a second step we will outline an architecture for integration of motor actions, vision and language representations.

II. ASSOCIATING MOTOR ACTIONS WITH ACTION VERBS

We have begun to associate internal representations of demonstrated actions with a word description. This system learns to associate the semantic features that are found in the sensor readings that represent the action with a representation of the word. As can be seen at the bottom of Fig 1 the architecture firstly contains a self-organising network to associate the action sensor readings with the appropriate body part by clustering the actions in different regions. At the next processing level there is a self-organising network for each body part that uses the sensor reading vectors to associate the actual action verbs with different regions. To the right in the architecture, the words that are represented using their phonemes are clustered in a self-organising network. The upper-most self-organising network associates the action representations by using the locations on the body part self-organising networks and their appropriate word form representation from the

location on the word form self-organising network. Hence by associating the action representation with the word form the robot can describe the action with a word when it receives only the action representation and vice versa perform the action when it is given the word only.

In this system the input is used to produce the output by recreating the action from the sensor readings. The sensor readings provide information on the action such as the velocity of the separate wheels, the gripper activities and how the constituent sub-actions relate to the states of sensors such as break-beam and table sensors. If the robot receives the 'put' action sensor reading representation, it would be introduced into the trained body part network and activate the hand region of the output layer. The hand self-organising network would then position the sensor readings input in the 'put' region of the output layer. As the robot is describing the word form there is no necessary input from the word self-organising network into the association self-organising network. However, as the network has previously learnt to associate this action with the appropriate word form the 'put' region of the network is activated. The robot will then state using its language synthesis that the action semantic features provided are those for 'put'. On the association self-organising network, the winner-take-all mechanism removes ambiguities in the representation, allowing for only one action.

This describes the pathway from the internal action representation via the association area to the language description. It would be used to make the robot speak from observing actions. In a similar but opposite pathway the word input representation can lead via the association area - to the sensory robot action. This would be used to make the robot execute a verbal command.

This approach offers some brain-inspired regional modularity by having multiple self-organising networks each performing a subtask of the overall task. These networks are linked in a distributed overall memory organisation. Furthermore, this architecture includes components that are analogous to brain regions at a higher level. For instance, the SOMs that take the action representations and cluster these are related to the sensory motor cortical areas of the brain. The approach also takes into account the neurocognitive evidence of Pulvermüller (2003) [11] in that cell assemblies in different regions are associated with specific action verbs as a functional unit, with the association being based on the action verbs relationship with the appropriate body part.

This architecture links in some concepts of the mirror neuron theory. The relationship of mirror neurons to language was pointed out by Rizzolatti and Arbib 1998 [14] who found that neurons located in the F5 area of a primate's brain were activated by both the performance of the action and its observation. The recognition of motor actions comes from the presence of a goal and so the

motor system does not solely control movement [8]. The role of these mirror neurons is to associate action representations with vision or language representations. The mirror neuron system was a critical discovery as it shows the role played by the motor cortex in action depiction [16]. By using the sensor readings as input the mirror neuron concept is considered since the understanding of the action can come from either performing the action or a stored representation is linked to observing the action.

III. SELF-ORGANISATION ON THE ROBOT

In order to have greater objectivity and to incorporate self-organising maps into a robot control system, sensor readings were taken from the MIRROR-neuron Robot Agent (MIRA) (see Fig 8). The MIRA robot is based on a Peoplebot and was set up to perform various actions that are associated in humans with the leg, head or hand. The leg verb actions were 'turn left', 'turn right', 'forward' and 'backward'; head action verbs were 'head up', 'head down', 'head right' and 'head left'; and finally the hand verbs were 'pick', 'put', 'lift', 'drop' and 'touch'. One action can be made of several basic actions. For instance, the hand verb action 'pick' included the following sub-actions (i) slowly move forward to the table; (ii) tilt camera downward to see table; (iii) lift gripper to table height; (iv) open gripper; (v) close gripper on object; (vi) stop forward motion; and (vii) lift gripper. The MIRA robot performing the 'pick' action is shown in Fig 8, top. This sequence of sub-actions corresponds in principle (although not in detail) to motor schemata since a complex action is represented as a sequence of basic actions. Sensor readings were taken for such sequences of basic actions.

In order to provide sufficient and varied training and test data the actions were performed 20 times (15 training and 5 test) under diverse conditions. For instance, the speed the robot was travelling at and the angle that the camera was tilted or panned to were varied. The sensor readings were taken 10 times a second while MIRA performed these actions including the state of the gripper, the velocity of the wheels and the angle that the robot's camera was at. The full list of the sensor readings is given in Table I.

To reduce the size of the input for the self-organising networks to a manageable level, 10 sets of the readings were taken over time to represent the action. This was achieved by taking the first, last and eight equi-distant sets of readings and combining them to create a single input for a sample. This procedure concatenates the whole time series to one data point and bypasses problems of short-term memory. We normalised the sensor readings for such variables as velocity of left wheel, velocity of right wheel, x coordinate of robot, y coordinate of robot, and the pan and tilt of the camera.

For the self-organising network to cluster actions based on the appropriate body part the input layer had 120 units,




Fig. 1. The self-organising associative architecture.

one for each of the preprocessed sensor readings. The output layers had various sizes (from 8 by 8 units to 13 by 13 units) and the networks were trained between 50 to 500 epochs at intervals of 50 epochs. There were 260 samples in total, 195 for training and 65 for testing. The location of each of the training and test samples on the self-organising maps were identified based on the units that had the highest activation.

Fig 2 shows a self-organising network that was 12 by 12 units. Once this network architecture was trained for 50 epochs there was clear clustering into the three body parts (see Fig 2). The hand action words such as 'pick', 'touch', 'lift' were at the bottom of the training and test output layers in the hand body part region, with the head actions slightly below and to the right of the leg region. Although one unit within the head region contained both head and leg action samples with the highest activation, the percentage for head samples was much higher on both test and training data. For the training and test data the percentage of head action samples with the highest activation for that unit was over 60% for training samples and 70% for test samples. Due to the major difference between the head and leg action percentages for this unit, only the head percentage is shown on Fig 2.

For the training data 100% of the samples which correspond to the head and hand actions fell in the appropriate region and 88% of the leg data. For test data the percentage

TABLE I
SENSOR READINGS TAKEN BY ROBOT DURING ACTIONS.

Sensor Reading	Value
Left Wheel Velocity	Real number
Right Wheel Velocity	Real number
x coordinate of robot	Real number
y coordinate of robot	Real number
Break-beam state of gripper	No beam broken, Inner broken, Outer broken, Both broken
Gripper state	Fully open, closed, between open and closed
Gripper at highest or lowest position	Yes No
Gripper moving	Yes No
Table sensors activated	Yes No
Gripper opening	Yes No
Pan of camera	Integer
Tilt of camera	Integer




Fig. 2. The percentages for the test samples for the body parts that have highest activation for each unit on a 12 by 12 units network after a training time of 50 epochs. (Black - Hand, White - Head, Grey - Leg)

was even better with 100% for hand and head and 90% for leg. It is interesting to note that within the hand verb region there was a good division into the actual action classes. In Fig 3 'pick' was located in the lower right of the map, 'put' in the lower left, 'drop' in the unit above 'pick', 'touch' at the top of the hand region and most of the 'lift' samples were located in a unit just below 'touch'. For the other two classes there was some splitting into the individual actions but not on the scale of the hand class (see Fig 4 and Fig 5).

For such an architecture on both training and test data the clusters were in very similar positions on the output layer, which points to the ability of the network to generalise on data it has not seen before. When considering the percentage of test data that fell in the regions identified by

the training data the percentages were very high. For the hand actions 100%, head actions 95% and leg actions 88% of the test data fell into the appropriate training region. Therefore, if the self-organising network was used in the control of a robot it can perform successfully in an on-line manner clustering semantic features of the action to the appropriate region of the output layer.

Turning to the hand, head and leg self-organising networks, when considering the clustering of the specific body part actions for all three types of action, the size of network that performed best was 8 by 8. For the hand network the training time that produced the best clustering was 50 epochs, for the head network it was 150 epochs and for the leg self-organising network it was 100 epochs. As can be seen from Fig 3 to Fig 5 there was clear clustering into different regions for the hand, head and leg actions.




Fig. 3. The percentage for the test samples for the specific hand actions that have highest activation for each unit on a 8 by 8 units network.

IV. ASSOCIATING VISION AND MOTOR REPRESENTATIONS

Our next step is to describe an associator neural network to localise a recognised object within the visual field. This is an essential basic skill for robotic learning by demonstration which we solve by a purely neuronal approach. The model, depicted in Fig 6, is thus a centrepiece of a larger model which can on the one hand perform actions like grasping and on the other hand is connected to neurally implemented language areas.

The idea extends the use of lateral associator connections within a single cortical area to their use between different areas [18]. The first cortical area is the visual area V1 which codes for an internal representation, "what", of images. The weights connecting it to the image are trained




Fig. 4. The percentage for the test samples for the specific head actions that have highest activation for each unit on a 8 by 8 units network.




Fig. 5. The percentage for the test samples for the specific leg actions that have highest activation for each unit on a 8 by 8 units network.

by a sparse coding Helmholtz machine. Earlier, intra-area lateral connections have been implemented within V1 to endow the simple cells with biologically realistic orientation tuning curves as well as to generate complex cell properties. We extend the lateral connections to also span a second cortical area, the "where" area which is laterally connected to the simulated V1. The lateral weights are trained to associate the V1 representation of the image to the location of an object of interest which is given on the "where" area. The lateral weights are thus object specific associative weights which can complete a representation of an image with the location of the object of interest.




Fig. 6. Model architecture. The hidden representation "what" of the image including the target object is associated to the location "where" of the target which is relevant for motor action.




Fig. 7. Example representations on the image, "what" and "where" areas. The *image* is originally in color, where in the upper row, the orange fruit target is artificially generated. The networks of the upper and lower row were also trained and activated with different parameters.

Fig 7, shows the network activities after initialisation with sample stimuli of an orange and relaxation to a steady state. In both cases the "where" area neuron's activations were initialised to zero initially (not shown). The relaxation procedure which spans the "what" and the "where" area then completes the pattern by displaying the location of the object of interest as a Gaussian activity hill.

Once that an object of interest appears in the visual field, it is first necessary to localise its position within the visual field. Then, usually the centre of sight is moved toward it, and a grasping movement prototype will be activated which is related to the specific affordance [15].

We have made initial experiments connecting the "where" area to motor neuron's output which control the robot camera's pan-tilt motors. The task is to move the camera so that the orange fruit is located in the centre of the "where" area (Figs. 6,7). This is achieved by a simple algorithm. Weights from every unit of the "where" area to the camera's pan and tilt units were trained based on the

error of a movement: if after a tilt movement the camera would face, e.g., too much upward, then the unit which elicited that movement had its weight to the tilt motor unit changed, so that at the next trial it would face a little less upward. Fig 8, bottom, shows the camera pointing toward an orange which is moved across its "visual field". This implements the MIRA robot's reaction to the command "Bot show orange".

Additionally, using reinforcement learning, we have very recently implemented the robot "docking" at a table so that it can grasp an object which lies at the border of the table with its short grippers [19]. The input to the reinforcement-trained network is the perceived target location (from the "where" area) and the rotation angle of the robot w.r.t. the table. Outputs are the four motor units and a critic unit which carries a value function on the input space. A positive reinforcement signal is given if two conditions are met: (i) the target is perceived at the middle of the lower edge of the visual field (where also the gripper is perceived by the camera which is at a fixed position) and (ii) the rotation angle is zero (which is defined such that the robot is approaching the table perpendicularly). The weights to the value function unit and those to the motor units develop concurrently such that an optimal strategy toward reaching the target will be performed. Fig 8, top, shows the robot perpendicularly at the table, at the goal position. The data delivered during these actions will be used for the training and verification of mirror neurons.

V. ASSOCIATING VISION, LANGUAGE AND MOTOR REPRESENTATIONS FOR LEARNING BY DEMONSTRATION

As the next steps therefore the model needs to be extended to incorporate more complex motor tasks. This is not only desirable from a robotic application point of view, but also from the fact that mirror neurons are action-related, as they reside in motor associated cortical areas such as F5 and respond to performance, description and observation of actions (Rizzolatti and Arbib 1998) [14]. For this, we will integrate the language- and motor-sensory related 'pick' action with the more vision relate the more vision related tracking action (Fig 8).

The model of Elshaw and Wermter (2002) [6] and Elshaw, Wermter and Watt (2003) [7] handles an organisation of a variety of actions on a self-organising layer of neurons as an avenue to include a larger number of motor tasks. Fig. 9 shows the plan of a proposed network.

In the envisaged network of Fig 9, mirror neuron properties are expected to evolve among some of the neurons in the top layer. They carry an internal representation \vec{r} of all of the inputs, below. The inputs are from multiple modalities including higher level representations. The vector \vec{l} contains language input information. This can




Fig. 8. The MIRA robot performing the 'pick' action (top) and recognising and tracking an orange with its pan-tilt camera (bottom).

include internal representations from language areas or the goal area of the cell assembly model for Broca/Wernicke areas. $\vec{p}v$ contains the visual perception which includes the identity and perceived location of a target to be grasped. \vec{m} are the motor unit activations including wheels, gripper and pan-tilt camera. $\vec{m}s$ denotes motor sensory unit activations and may also include available idiothetic information such as the rotation angle of the robot. \vec{l} are other internal states such as the goal related value function of the critic used in reinforcement learning.

Thick lines with arrow heads denote the weights. The vertical connections are trained with a sparse coding unsupervised learning scheme similar to the Helmholtz machine which we described for image processing (Fig 6). The inputs are collected from real robotic actions (after exercising with simulated data) which are performed interactively in the environment. The data contain only instantaneous information, i.e. the whole action sequence is not known. Therefore, neurons do not necessarily fire over a sustained period in time as do mirror neurons. However, since \vec{r} is a distributed code, some of the units may specialise to code for longer sequences. The horizontal




Fig. 9. The envisaged associative architecture

recurrent connections (depicted as open circle) are trained as an autoassociator neural network. They are used in a neural activation relaxation procedure which de-noises the representation \vec{r} and may also encourage prolonged firing. As a possible extension, associator recurrent connections may also feed back to the input. This would be particularly interesting for the cortical feed back to the motor units, because of implications for motor control.

VI. DISCUSSION

We have developed biologically inspired solutions for tasks which are needed by a robot that should learn by demonstration and instruction.

The robot sensor inputs to the modular, self-organising network were partitioned in a way that they match the three body areas 'leg', 'head' and 'hand'. The match is intuitive, but equivalents of the robotic sensor readings (like "gripper opening") are likely to be represented at various locations on the cortex, as a visual or motor-sensory perception or distributed in the language system as a "word web" [11]. The network can in principle realise the findings of Pulvermüller et al. (2000) [13] on the processing of action verbs with different clusters representing the specific body parts. The network was able to identify the semantic features from the actual sensor readings for the individual action verb classes that were specific to the appropriate body part. These features were likely to include the degree of move, whether there was an object involved and the type and number of motors used.

The performance of the head, leg and hand self-organising networks are in principle suitable for use in a robot control system based on language instruction. This is because it is likely, based on the clear clustering demonstrated, that the sensor reading input will be accurately represented and mapped to the appropriate network region. As this location is the basis for the association between the action and the word this will contribute to the successful identification of the action and its description.

A recurrent associator network with distributed coding was applied to the visually related part of the task. Such associator networks form the neural basis for multimodal convergence and at the same time can supply a distributed representation across modalities as has been proposed for linguistic structures [12]. Multimodal representations furthermore allow for mirror neuron-like response properties which shall emerge in our application within a biomimetic mirror neuron-based robot, MirrorBot.

Two actions, interactively performed with the environment, shall supply input data to the envisaged mirror neurons. Since reinforcement learning which we used to train these actions is attributed to the basal ganglia, the model extends beyond the cerebral cortex, in a biologically plausible fashion.

VII. CONCLUSIONS

We have described some research toward integrating learning by demonstration and learning by instruction on a neural substrate on a robot. Our approach is not so much on imitating complex behaviour. Rather our focus is on testing mirror neuron concepts and other neurocognitive evidence like the topological arrangement of actions in order to provide a multimodal integration of the robots own actions, as well as visual observation and language instruction. We think that visual observation and language instructions are complementary forms of programming robots in a natural manner to perform and link their performance to their own underlying actions. An associative neural organisation of the internal memory may therefore be advantageous for a robot's learning of visually described actions or verbally instructed actions.

VIII. ACKNOWLEDGEMENTS

This work is part of the MirrorBot project supported by the EU in the FET-IST programme under grant IST-2001-35282.

IX. REFERENCES

- [1] H. Asoh, S. Huyamizu, H. Isao, Y. Motomura, S. Akaho, and T. Matsu, "Socially embedded learning of office-conversant robot jijo-2", in *Proceedings of 1997 International Joint Conference on Artificial Intelligence, Nagoa*.
- [2] A. Billard, "Imitation: a means to enhance learning of a synthetic proto-language in an autonomous robot", in Dautenhahn, K. and Nehaniv, C. (eds), *Imitation in Animals and Artifacts*, Academic Press, pp. 281-311, 2001.
- [3] C. Breazeal and B. Scassellati, "A context-dependent attention system for a social robot", in *Proceedings of the 1999 Sixteenth International Joint Conference on Artificial Intelligence (IJCAI99)*, Stockholm, Sweden, pp. 1146-1151.

- [4] W. Burgard, A.B. Cremers, D. Fox, D. Hähnel, G. Lakemey, D. Schulz, W. Steiner, and S. Thrun, "Experiences with an interactive museum tour-guide robot", *Artificial Intelligence*, Vol. 114, No. 1-2, 2000.
- [5] Y. Demiris and G. Hayes, "Imitation as a dual-route process featuring prediction and learning components: A biologically plausible computational model", in Dautenhahn, K. and Nehaniv, C. (eds), *Imitation in Animals and Artifacts*, MIT Press, 2002.
- [6] M. Elshaw and S. Wermter, "A neurocognitive approach to self-organisation of verb actions", in *Proceedings of 2002 International Joint Conference on Neural Networks, Honolulu, USA*, pp. 24-29.
- [7] M. Elshaw, S. Wermter and P. Watt, "Self-organisation of language instruction for robot action", in *Proceedings of 2003 International Joint Conference on Neural Networks, Oregon, USA*.
- [8] V. Gallese and A. Goldman, "Mirror neurons and the simulation theory of mind-reading", *Trends in Cognitive Science*, Vol. 2, No. 12, 1998, pp. 493-501.
- [9] P. Gaussier, S. Moga, J.P. Banquet and J. Nadel, "Learning and communication in imitation: An autonomous robot perspective", *IEEE Transaction on Systems, Man and Cybernetics, Part A: Systems and Humans*, Vol. 31, No. 5, 2001, pp. 431-444.
- [10] G. Maistros and G. Hayes, "An imitation mechanism for goal-directed actions", in *Proceedings of TIMR 2001 - Towards Intelligent Mobile Robots, Manchester*.
- [11] F. Pulvermüller, *The neuroscience of language: On brain circuits of words and serial order*, Cambridge University Press, 2003.
- [12] F. Pulvermüller, "Words in the brain's language", *Behavioral and Brain Sciences*, Vol. 22, No. 2, 1999, pp. 253-336.
- [13] F. Pulvermüller, M. Hare and F. Hummel, "Neurophysiological distinction of verb categories", *Cognitive Neuroscience*, Vol. 11, No. 12, 2000, pp. 2789-2793.
- [14] G. Rizzolatti and M. Arbib "Language within our grasp", *Trends in Neuroscience*, Vol. 21, No. 5, 1998, pp. 188-194.
- [15] G. Rizzolatti and G. Luppino, "Cortical motor system", *Neuron*, Vol. 31, 2001, pp. 889-901.
- [16] G. Rizzolatti, L. Fogassi and V. Gallese, "Neurophysiological mechanisms underlying the understanding and imitation of action", *Nature Review*, Vol. 2, 2001, pp. 661-670.
- [17] S. Thrun, M. Bennewitz, W. Burgard, F. Dellaert, D. Fox, D. Haehnel, C. Rosenberg, N. Roy, J. Schulte and D. Schulz, "MINERVA: A second generation mobile tour-guide robot", in *Proceedings of the 1999 IEEE International Conference on Robotics and Automation (ICRA'99)*.
- [18] C. Weber and S. Wermter, "Object localisation using laterally connected "What" and "Where" associator networks", in *Proceedings of the 2003 International Conference on Artificial Neural Networks, Istanbul, Turkey*, pp. 813-820.
- [19] C. Weber and S. Wermter, "Robot Docking with Neural Vision and Reinforcement", in *Proceedings of the AI-2003, Twenty-third SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence, Cambridge, UK*, (accepted).
- [20] S. Wermter, J. Austin, D. Willshaw and M. Elshaw, "Towards novel neuroscience-inspired computing", in S. Wermter, J. Austin, and D. Willshaw, (Eds.), *Emergent Neural Computational Architectures based on Neuroscience*, Heidelberg, Germany: Springer-Verlag, 2001, pp. 1-19.
- [21] S. Wermter and M. Elshaw, "Learning robot actions based on self-organising language memory", *Neural Networks*, Vol. 16 pp. 5-6, 2003, pp. 661-669.
- [22] S. Wermter and C. Panchev, "Hybrid preference machines based on inspiration from neuroscience", *Cognitive Systems Research*, Vol. 3, No. 2, 2002, pp. 255-270.