Aberystwyth University

A developmental approach to the emergence of communication in socially situated embodied agents

Sheldon, Michael; Lee, Mark

Published in:

IEEE 9th International Conference on Development and Learning

DOI:

[10.1109/DEVLRN.2010.5578843](https://doi.org/10.1109/DEVLRN.2010.5578843)

Publication date:

2010

Citation for published version (APA):

Sheldon, M., & Lee, M. (2010). A developmental approach to the emergence of communication in socially situated embodied agents. In *IEEE 9th International Conference on Development and Learning* (pp. 204-210) <https://doi.org/10.1109/DEVLRN.2010.5578843>

General rights

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

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

Take down policy

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

tel: +44 1970 62 2400

email: is@aber.ac.uk

A developmental approach to the emergence of communication in socially situated embodied agents

Michael Sheldon

Department of Computer Science
Aberystwyth University
Aberystwyth, Ceredigion, SY23 3ET, UK
Email: mts08@aber.ac.uk

Mark Lee

Department of Computer Science
Aberystwyth University
Aberystwyth, Ceredigion, SY23 3ET, UK
Email: mhl@aber.ac.uk

Abstract—This paper reports on a developmental approach to the learning of communication in embodied agents, taking inspiration from child development and recent advances in the understanding of the mirror neuron system within the brain. We describe a part of the ROSSI project which focuses upon gestural communication in the form of pointing. We are examining the idea that pointing may be a key step towards simple spoken communication and exploring the internal representations that may be formed during this process.

The possible developmental stages leading to proto-imperative pointing actions in a robotic system are outlined, and how this may be built upon to result in an understanding of two word speech is discussed. The learning mechanism is based around Piagetian schema learning whilst the developmental path follows a mixture of Piagetian and Vygotskian theories.

Index Terms—Language Development, Human-Robot Interaction, Embodied Cognition, Grounding of Knowledge and Representations

I. INTRODUCTION

The developmental approach to robotics in which systems attempt to mimic similar stages of development to a human infant has so far had little application to the possible emergence of communication and symbol grounding in robotic systems.

In this paper we explore the emergence of early gestural communication as a side-effect of sensorimotor robot learning and how this may be used to boot-strap simple linguistic communication for robotic systems. While we take inspiration from infant development we do not claim to be accurately modelling human development.

The success of this approach depends upon the willing co-operation of other social agents to aid the robot in its learning. The robot is not imbued with any innate theory of communication, so if it never experiences communicative acts from other agents it will be unable to learn to communicate itself.

This approach allows symbolic meaning in the form of language to be strongly rooted in the sensorimotor experience of the agent, with the various concepts involved in communication arising out of interaction with the environment and other agents. The same learning framework is used throughout all stages of development (although not all aspects of the framework are used at all stages) allowing more advanced concepts to be grounded in simpler ones from earlier developmental stages.

A. Developmental stages

The following seven stages outline a possible robotic developmental progression leading from a “new born” state to simple linguistic communication. This paper focuses mostly on the stages leading to pointing gestures, with future work extending this to build up to speech. This progression, especially in the latter stages, is based heavily upon that described by Iverson and Goldin-Meadow [10], discussed further in section II-A (communication).

1) *Motor babbling*: In this initial stage the robot has had no prior experience of the world or of its own body. It performs spontaneous motor actions in order to discover the properties of its motor systems and its anatomical constraints.

2) *Motor vision mapping*: The movements learnt in the previous stage are then mapped to the changes they create in the robot’s vision system, this allows it to move its arm to touch (or point towards) an object detected visually. While the focus in this paper is on visual mappings this could equally be applied to other sensor modalities.

3) *Failed grasping leading to pointing*: In attempting to touch objects that lie outside of the work-envelope of the robot it will incidentally perform what looks, to a human observer, like a pointing motion. Through assistance from a human observer, fetching the indicated object for the robot, the robot’s representation of this action moves away from being a direct attempt at manipulating the world towards an attempt at social communication.

4) *Recognising pointing in others*: Using a goal directed approach based on mirror neuron theory, the robot is then able to learn to reciprocate, providing objects to humans (or other robots) when they are indicated. This allows the structures necessary for a simple give/take conversation to emerge prior to the introduction of language.

5) *Complementary one word speech with pointing*: The robot is then given auditory input (reduced to a text token by speech recognition software) whilst it points at objects, or whilst it sees a human or other robot point at an object. This input is directly related to the object being indicated, for example the word “ball” or “block”.

6) *Supplementary one word speech with pointing*: In this stage the auditory input relates to the action being indicated rather than the object itself. The pointing action has been used

in the preceding stages as a request for the indicated object, so the word “give” becomes associated with this action.

7) *Two word sentences*: Finally the learnt speech can be combined to allow the robot to form and understand two word sentences of the form “give block”, replacing the pointing behaviour from the earlier stages.

II. BACKGROUND AND RELATED WORK

A. Communication

Vygotsky suggested that pointing develops out of a failed grasping behaviour in which the child attempts to reach for an object which is too far away, the parent interprets this as the child pointing at a desired object and as such fetches the object for the child, thus associating a new meaning with the act of reaching for a distant object [27], [14]. Initially all social meaning in this act is inferred entirely by the parent, the infant is making a real attempt to reach the object and failing, but through the actions of the parent the infant comes to associate the same communicative meaning.

This has been classified by many researchers as *proto-imperative pointing* or *ritualised grasping*, used by the child to indicate an object of desire to a nearby adult, and typically emerges at around 10-12 months. On average 3 months [5] after the emergence of proto-imperative pointing the child has also learnt to perform *proto-declarative pointing* which is used to acquire joint attention on an object with an adult.

There is however evidence presented by Masataka [17] to indicate that proto-imperative pointing and proto-declarative pointing may follow different developmental paths, with proto-imperative pointing actually arising out of index finger extension for the purposes of object exploration. The developmental progress previously outlined for robotic systems currently emerges proto-declarative pointing before going on to make the leap towards treating this in a proto-imperative manner in some of the later stages, while this may not be the exact progression experienced by children it provides a simpler mechanism for a robot to learn pointing gestures, albeit in a less rich developmental context. In addition to this Tomasello, et al. [25] show that infants may possess a much deeper social understanding at this stage than previously thought, able to communicate a great deal through pre-linguistic gestures such as pointing.

Butterworth [4] provides various evidence supporting the theory that gesture is nearly universal on the road to further language development.

Iverson and Goldin-Meadow [10] describe the early developmental path of infants learning to communicate verbally. They show that in most cases infants follow a consistent progression from pointing to two word speech, as described in the later stages of the previously outlined developmental progression.

B. Neuroscience

The mirror neuron system was first discovered in the brains of monkeys [7], [22] and later studies showed a similar system at work in the human brain. A mirror neuron is a neuron

which fires both upon the execution of an action and upon the observation of another agent performing the action. Each mirror neuron is paired with a canonical neuron, however the canonical neuron is only activated during the execution of an action and not during its observation. This has prompted speculation that the mirror neuron system may have been crucial in the evolution of language [1].

Tettamanti, et al. [24] show that listening to action related sentences can trigger a mirror neuron response in humans and Kohler, et al. [11] have previously found that a noise related to an action can trigger a response in monkeys. This adds further weight to the idea that the mirror neuron system encodes action content at an abstract level and that this content can be activated auditorily. This suggests that language is strongly linked to the sensorimotor system.

A study by Buccino, et al. [3] suggests that mirror neuron responses only occur for actions that the observer can duplicate. For example humans watching a dog biting will show frontoparietal activity, while they will not when watching a dog bark. This also shows that the mirror neuron system generalises to different species, possibly suggesting that the goal of the action has a much greater effect than the observation of the action itself.

The goal directed nature of mirror neurons is further reinforced by a study by Umiltà, et al. [26] in which the neural response from monkeys was measured when they observed the experimenter grasping an object and when they observed a mimed grasp with no object present. It was found that the mimed grasp produced no response, while the real grasp did. It was also found that if the view of the object was occluded so the final stage of the grasp wasn't visible then some response was still produced, suggesting that the goal was being inferred from the action.

Oztop and Arbib [19] hypothesise that the mirror neuron system may have evolved to provide feedback for visually directed grasping with the social usage being an exaptation¹ occurring when this became applied to the hands of others.

Oztop, Kawato and Arbib [20] provide a computationally guided review of mirror neuron literature and provide box diagrams of a computational model called the MNS model. Bonaiuto, et al. [2] have made attempts to extend this model, creating a more comprehensive version titled MNS2. Small sections of this model have been implemented and tested, but the model as a whole remains largely theoretical.

C. Robotics and artificial intelligence

Drescher [6] suggests a constructivist approach to learning based on Piagetian ideas using the notion of “*schemas*”. Schemas are units of knowledge associating perceptions, actions and predictions. If the environment is perceived to be in a certain state then taking an action associated with this state should cause the environment to change to match the sensor values specified in that schema's prediction.

¹An exaptation being the exploitation of an evolutionary adaptation to serve a different purpose than the one it initially developed for.

In its simplest form a schema consists of a set of pre-conditions, an action and a set of post-conditions (often represented in the form *pre-conditions/action/post-conditions*), providing a basic forward learning model.

Holmes and Isbell [9] extended Drescher's work to enable the use of continuous value sensors (the original implementation was limited to binary sensors). They showed that it was possible to model Partially Observable Markov Decision Processes (POMDPs) via this mechanism.

Guerin [8] has since used this approach in a simple simulated robotic environment, but as yet little work has been performed using this technique on a physical robot.

Perotto, et al. [21] introduce a Constructivist Anticipatory Learning Mechanism (CALM), which makes use of a schema based learning mechanism. The schemas are organised in a tree hierarchy going from most general to most specific, making it possible for the system to fall back on more general solutions if a specific one fails or is unavailable. In contrast to Holmes and Isbell this system took a property based approach to the environment providing a more direct mapping between the environment and the agent's perceptions than a state based environment.

Lee, et al. [13], [12] discuss the use of a Lift Constraint, Act, Saturate (LCAS) loop to artificially constrain the inputs to the robotic system and so reduce the complexity of the learning required at each stage of the system's development. This approach is similar to the scaffolding [15] performed by parents when helping children to learn in that the staged constraints placed upon the system's sensory input provides a framework that guides the robot through its development. Once there is little novel input being found at one stage of learning a constraint is lifted, allowing the system to build upon its knowledge from the previous stage whilst being exposed to a more complex and detailed view of the world.

Marjanovic, et al. [16] introduce a motor-vision mapping system that learns to perform pointing motions towards visual targets. Our system differs from this in that the one presented by Marjanovic has an explicit goal of pointing, while in our system this behaviour emerges as a side effect of other developmental processes occurring at the same time and as a product of social interaction.

Steels, et al. [23] show that the concept formation process of agents must be based on similar sensor input and result in similar conceptual repertoires for communication to develop in a population of agents. It also shows that once a lexical system is in place it can overcome the randomness inherent in verbal communication.

Oudeyer and Kaplan [18] explore the intrinsic motivation of language learning rooted in play and curiosity, using a framework based around Vygotsky's zone of proximal development [28] (although this is termed "*progress niches*" within this system). It shows how an intrinsic motivation system can allow a robot to self-organise its learning process.

III. HARDWARE CONFIGURATION

The hardware that the system is being tested on consists of an Adept manipulator arm mounted on a rigid vertical back-plane. The arm is configured to operate on a two-dimensional manifold above a table upon which objects can be placed for it to interact with, the manifold curves up at the extremities tracing the outer limit of the robot's work envelope allowing for pointing towards distant objects. The arm has a single "finger" as an end effector, which has four touch sensors attached giving directional touch input. This end effector can be used for interacting with objects by touching them and pushing them around the work area and for communicating by pointing at an object.

The vision system consists of an AVT Stingray F-046C firewire camera, which provides a resolution of 780x580 at up to 61 frames per second. This is mounted on a pan tilt platform above the arm looking down on the work space.

This hardware setup can be seen in figure 1.

Fig. 1. The current hardware configuration.

Use of this learning framework in the context of a more complex system, involving many more degrees of freedom, is discussed briefly in section VII (future work).

IV. THE SOFTWARE FRAMEWORK

The system consists of two main components, the schema memory and the developmental controller. The developmental controller determines the goal of the system based on the current excitation level and motivation, as well as handling the reduction of complexity in sensory input based upon the current learning stage. These sensor values are then passed to the schema memory along with either an action or a desired goal state, the resulting schema(s) are then executed and the results stored for use in judging their suitability for future tasks.

A. Schema learning framework

A similar approach to Drescher's schema learning is used to achieve the desired learning behaviour, albeit with a number of modifications from Drescher's original design to make the technique more applicable to robotics. Unlike Drescher's binary system or Holmes and Isbell's continuous value system the schema framework makes use of discrete sensor values

made up of sensorimotor fields which reduce the complexity of the sensor input and motor output.

While a very symbolic schema representation has been chosen here a neural implementation should give similar results, however we believe that a more explicit symbolic representation lends itself to easier analysis of the resulting generated internal structures.

B. Schema chaining

The linking of pre-conditions and post-conditions from different schemas (“*schema chaining*”) creates a traversable network representing different world states and the actions required to move between these states, as illustrated in figure 2. Without schema chaining the robot’s interest in unreachable objects would decrease as it failed to reach them. Schema chaining allows for cases in which the feedback of an action isn’t instantaneous to still be recognised as being useful. Thus making the entire series of actions required to point at an object, wait for another agent to move the object then touch the object interesting to the robot, despite the reward (touching the object) being at the end of the chain of actions.

Fig. 2. A high level example of schema chaining, allowing the robot to gain access to an object that would otherwise be outside of its reach through communication with another agent.

C. Tracking of Schema Probabilities

The schema framework keeps track of the observed probabilities of the post-conditions of each schema, allowing it to predict the most likely outcome. The storing of probabilities for the likelihood of individual items, instead of the probability of the schema as a whole being successful (as proposed by Drescher) allows the system to select the best action for achieving its target goal, regardless of the likelihood of less interesting side-effects of the action. For example it makes little sense for the system to care how likely it is that a particular block is moved when the goal of the action is just to move the arm to a specific location, the movement of a block (or lack thereof) is merely an uninteresting side-effect in the context of this particular goal.

D. Schema Generalisation

The system periodically attempts to generalise its existing schemas, the specific schemas from which these generalisations arise are retained and when an attempted action does not meet the expected outcome from a general schema a new specific schema is created, allowing future attempts at refining the generalisation with the added information from the failed tests. When performing an action a specific schema is preferred over a generalised schema if one exists that matches.

For example when seeing a number of specific schemas along the lines of *[object in field 5] / [move arm to field 5] / [object in field 5, finger in field 5, touching]* the system will generate a general schema of the form *[object in field \$x] / [move arm to field \$x] / [object in field \$x, finger in field \$x, touching]*.

E. Mirror Neuron Influenced Schema Learning

To enable the schema system to mimic the behaviour of the human/primate mirror neuron system it is split up in to two distinct classes, traditional schemas, with pre-condition, action and post-condition components, and “*perceptual schemas*” which lack an action component and are used for observing the actions of another agent. These classes are linked together when a perceptual schema and a traditional schema have matching post-conditions allowing the observation of other agents performing an action to be associated with the observer’s motor schema for that same action.

Experiments with monkeys have suggested that their mirror neuron system is largely dependent on the goal of an action [7], [26], in that mirror neurons will fire when the monkey observes the experimenter grasping an object, but will not fire, or only fire very weakly, when they observe a “*pantomime gesture*”, in which the experimenter performs the same action but without an object present. The system mimics this behaviour by using post-conditions as the linking mechanism between the different schema classes, for example the visual input of watching another agent move an object can be associated with the motor behaviour performed when the observer is themselves moving an object.

In addition to providing a mechanism for recognising the actions of other agents, this also provides part of the framework necessary for spoken language. This will allow the linking of auditory observations to actions and other observations, helping the system move from a “*motor-meaning*” based representation to a “*symbolic-meaning*” which is one of the key differences between Piaget’s stage 3 and stage 4 infant.

F. Developmental controller

A control program has been developed that makes use of the schema framework. The control system has two different modes of operation, a “*play*” mode, in which it randomly executes schemas based on their predicted excitement and reward levels and a “*task*” mode, in which it can perform more goal directed actions.

The resolution of the robot’s inputs are reduced by the control program to speed learning. The possible joint configurations are reduced to typically 200-300 combinations depending on the robot configuration (referred to as the “*motor fields*”). This is achieved by limiting the robot to the use of two joints, each moving in 10 degree increments, a third joint becomes accessible when the robot is at the outermost limit of its two joint work envelope allowing the end-effector to be moved outwards tracing a vertical arc to allow for pointing. The visual system is similarly divided in to circular fields each with a 10 pixel radius, referred to as “*visual fields*”, a specific type of sensorimotor field. A new visual field is created each

time the robot observes an object outside of its current fields, with the centre point of this object forming the centre of the field. The fields are initially discovered by the robot exercising its previously learned motor schemas and observing its end-effector entering the different visual fields [12].

The controller implements a Lift-Constraint, Act, Saturate (“LCAS”) [13] based approach to staged learning. Additional constraints are added to the robot’s sensory input, these constraints are lifted as the robot becomes habituated to its current level of development. The point at which these constraints should be lifted is determined by the system’s excitation level. This excitation level is also used to decide which actions to perform next, for example an action which would cause a new schema to be created would be considered more exciting than the execution of an existing schema. Whether an action is executed or not depends on whether or not it is above a certain threshold below the global excitation level. This means that if most of the actions the robot is performing are creating new schemas then it is unlikely to execute any existing schemas, however once there are fewer new schemas to discover it will begin to re-activate existing schemas with a preference for those which have had the fewest activations. Figure 3 shows the system reaching a plateau during the motor babbling stage, once this has been reached the constraint on the vision system is lifted and the robot begins to map visual input to its existing schemas, the number of schemas does not begin to rise again until the robot’s environment is made more complicated through the introduction of wooden blocks for it to manipulate. In addition to novelty-triggered excitation the robot also receives a reward for successfully touching an object (making such actions more exciting). This biases it towards actions that may result in contact when an object is present, this helps to speed learning by focusing the robot’s attention on actions more likely to lead towards the desired behaviour.

Fig. 3. As time progresses the robot discovers less novel situations requiring new schemas and instead exercises existing schemas to test their reliability. Once the developmental controller detects the saturation of this stage of development the complexity of the input to the system is increased and the schemas can be further refined.

V. REPRESENTATIVE SCHEMAS

A. Motor babbling

Initially very basic schemas are created with no context, representing only actions.

Pre-conditions	Action	Post-conditions
	<i>Motor action</i> joint1=0.69 joint2=0.87 joint3=0	

B. Motor vision mapping

Later the most basic visual result of these actions (the end effector appearing in a different field) are added as post-conditions.

Pre-conditions	Action	Post-conditions
	<i>Motor action</i> joint1=0.69 joint2=0.87 joint3=0	End effector in field 7

C. Touching objects

Next the robot is given a few examples of touching objects in different positions.

Pre-conditions	Action	Post-conditions
Object in field 3	<i>Motor action</i> joint1=0.23 joint2=0.43 joint3=0	Object in field 3 End effector in field 3 Touching

Once a number of examples along these lines have been viewed this gets generalised, to give a schema which represents touching an object in any position on the work surface.

Pre-conditions	Action	Post-conditions
Object in field \$x\$	<i>Target action</i> End effector in field \$x\$	Object in field \$x\$ End effector in field \$x\$ Touching

In these generalised schemas we see the use of “target actions” replacing direct motor actions, rather than causing a direct change in the robot’s configuration they represent a target set of post-conditions that should be achieved (which is a subset of the post-conditions of the main schema). This allows the generalisation to occur across the pre-conditions, post-conditions and action with consistent variables.

D. Pointing counter-examples

In the case of pointing the system attempts to execute the above generalised touching schema but fails, generating a specific counter-example. Specific schemas are always preferred over generalised schemas if both fulfil the same conditions. This allows the system to learn where its generalisation fails and create schemas that work in those

situations.

Pre-conditions	Action	Post-conditions
Object in field 203	Target action End effector in field 203	Object in field 56 End effector in field 203

Due to the random, non-contiguous nature of the visual fields the current system must learn each pointing location individually. In future experiments the system will be tested with a predefined contiguous visual space allowing the generation of a second generalised schema representing the pointing space with a conditional observation such as " $x > 134$ ", where field 134 marks the shift from touching to pointing.

VI. POINTING MECHANISM

The controller takes the system through two learning stages to create a mapping between the motor system and the vision system. This mapping allows the robot to move its end effector in to a desired visual field, which can then be used for allowing it to interact with objects (both by physically touching them and moving them around itself and by pointing at them for communication).

The first stage of this process is akin to Piaget's first stage infant, the robot goes through a period of "*motor babbling*", where it exercises all possible joint configurations and creates schemas representing these actions. It receives no feedback from these actions, merely generating a base set of schemas that abstract higher level schemas away from explicit joint commands, allowing them to instead refer to existing schemas as their action components.

In the second stage the vision system is made available to the robot and it begins to associate visual context with the existing motor schemas. This is similar to hand fixation in an infant. This stage is visualised in figure 4. The robot executes the purely motor based schemas it has learnt in the previous stage and forms a new visual field whenever it sees its end effector outside of any existing fields, it then adds this as a new post-condition to the executed schema. The end effector is detected via the vision system, potentially it will add any changes in visible objects as post-conditions, however at this stage of the robots learning no other objects are presented to it.

Fig. 4. A visualisation of the visual fields, part-way through their discovery.

The system operates primarily on the X-Y plane using 2 degrees of freedom, illustrated in figure 5(a). To enable the

robot to point at objects outside of its work envelope it is able to slightly lift its end effector when at the furthest extent of its normal range of motion, shown in figure 5(b). Both of these planes are accessible to the robot throughout all stages of learning, so it first learns to position its end effector in the 'pointing' plane prior to any objects being introduced for it to point at as part of its random motor babbling and vision mapping stages.

It is important to note that this is not giving the robot a full 3D representation of the space it occupies as the robot still effectively lacks accurate depth perception, however for the purposes of this experiment that is unimportant and may indeed be congruous with a child's perception at this stage. If similar operations were performed on a system with more degrees of freedom the same outcome should be possible, with the added benefit that the system would be able to point to objects within its work envelope without touching them. We only constrain the system to 2 DoF to greatly simplify the lower level motor learning tasks.

Fig. 5. (a) Direct manipulation of objects is only possible in the X-Y plane. (b) The robot's motion is extended in to the Z' plane allowing it to pointing to distant objects. This is a simplified example, rather than having two distinct vertical and horizontal planes the system operates on a manifold that curves up at the extremities tracing the outside of the robot's work envelope.

When first learning to point the robot views an object and moves its end effector to occupy the same visual field, using a generalised form of the schemas it has developed to allow it to touch objects (and so receive a reward). However in this case the schema does not successfully result in contact with the end effector, instead it results (from the perspective of a human observer) in a pointing motion towards the object. The robot is receiving no reward when failing to touch the object, however in the event that a human observer assists the robot by moving the object closer it leads to a chain of events which finally results in the robot touching the object and so being rewarded.

A. Morphological implications

This approach raises certain morphological implications. For a pointing gesture that a human would recognise to emerge from this technique the robot in question must itself have a roughly humanoid anatomy. Specifically it requires the robot's vision system to be positioned above the arm system looking out in the direction of action. Additionally for the pointing to

appear accurate the vertical distance between the vision system and the arm should not be too great.

All current testing has been performed with humans with prior knowledge that what they are about to view is intended as a pointing gesture, it might be interesting to investigate the effects this gesture has on people who do not already know what to look for. The anthropomorphic characteristics of the robot in question might play as large a part in this as the quality of the gesture itself. However for now this is outside the scope of the current investigation.

VII. FURTHER WORK

This paper deals primarily with the initial emergence of pointing behaviour and the stages preceding it. We are continuing with the later stages in the developmental progression, including the recognition of pointing from other agents and the transition to linguistic communication. Work on these aspects is ongoing.

We have implemented a neurally-inspired reaching/grasping model for a 7 DoF tactile sensing robot hand (Schunk GmbH & Co.) as part of the ROSSI project. The schema system is in the process of being integrated with this so that a wider range of possible actions and gestures may be investigated. In this configuration rather than dealing with the vision and motor system directly the schema system talks to an affordance based memory which processes object features and determines the appropriate joint configuration for grasping them, meaning the schema system can continue to operate at a fairly high, symbolic level while the affordance memory deals with the low level joint configuration in more detail. This system also has the capacity to recognise human hand positions via a data glove, which provides an ability to imitate humans and will allow us to determine more accurately when a human is pointing at an object. Schema learning adds a capacity for temporal reasoning and goal directed behaviour that is lacking in the current affordance based grasp system.

There is also the potential for further work focusing upon one robot having learnt this process with a human teacher and then going on to teach a second robot in a similar manner. This could be further extended to look at the implications on a larger population of robots and how social meanings might adapt due to slight changes in the teaching process from one robot to another, following a similar methodology to Steels, et al. [23].

In the current system the robot has no mechanism for perceiving the presence of another agent as there is assumed to always be a human present. If this facility were to be added in the future it would allow the robot to discover in which scenarios social acts are likely to be successful.

VIII. ACKNOWLEDGEMENTS

This work is funded by EC-FP7, ICT – 216125 (ROSSI).

REFERENCES

- [1] M. Arbib. The mirror system, imitation, and the evolution of language. *Imitation in Animals and Artifacts*, pages 229–280, 2002.
- [2] J. Bonaiuto, E. Rosta, and M. Arbib. Extending the mirror neuron system model. I. *Biological Cybernetics*, 96(1):9–38, 2007.
- [3] G. Buccino, F. Lui, N. Canessa, I. Patteri, G. Lagravinese, F. Benuzzi, C. A. Porro, and G. Rizzolatti. Neural circuits involved in the recognition of actions performed by nonconspecifics: An fMRI study. *Journal of Cognitive Neuroscience*, 16(1):114–126, 2004.
- [4] G. Butterworth. Pointing is the royal road to language for babies. *Pointing: Where language, culture, and cognition meet*, pages 9–33, 2003.
- [5] L. Camaioni, P. Perucchini, F. Muratori, and A. Milone. Brief report: A longitudinal examination of the communicative gestures deficit in young children with autism. *Journal of autism and developmental disorders*, 27(6):715–725, 1997.
- [6] G. L. Drescher. *Made-Up Minds: A Constructivist Approach to Artificial Intelligence*. MIT Press, 1991.
- [7] V. Gallese, L. Fadiga, L. Fogassi, and G. Rizzolatti. Action recognition in the premotor cortex. *Brain*, 119(2):593, 1996.
- [8] F. Guerin and D. McKenzie. A Piagetian model of early sensorimotor development. *Proceedings from Epigenetic Robotics 2008*, 2008.
- [9] M. P. Holmes and C. L. Isbell Jr. Schema Learning: Experience-Based Construction of Predictive Action Models. *Advances in Neural Information Processing Systems*, 17:585–592, 2005.
- [10] J. M. Iverson and S. Goldin-Meadow. Gesture Paves the Way for Language Development. *Psychological Science*, 16(5):367–371, 2005.
- [11] E. Kohler, C. Keysers, M. Umiltà, L. Fogassi, V. Gallese, and G. Rizzolatti. Hearing sounds, understanding actions: action representation in mirror neurons. *Science*, 297(5582):846–848, 2002.
- [12] M. H. Lee, Q. Meng, and F. Chao. Developmental learning for autonomous robots. *Robotics and Autonomous Systems*, 55(9):750–759, 2007.
- [13] M. H. Lee, Q. Meng, and F. Chao. Staged competence learning in developmental robotics. *Adaptive Behavior*, 15(3):241, 2007.
- [14] J. Lindblom and T. Ziemke. Social Situatedness of Natural and Artificial Intelligence: Vygotsky and Beyond. *Adaptive Behavior*, 11(2):79, 2003.
- [15] M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini. Developmental robotics: a survey. *Connection Science*, 15(4):151–190, 2003.
- [16] M. Marjanovic, B. Scassellati, and M. Williamson. Self-taught visually guided pointing for a humanoid robot. In *From Animals to Animals 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 35–44, 1996.
- [17] N. Masataka. From index-finger extension to index-finger pointing: Ontogenesis of pointing in preverbal infants. *Pointing: Where language, culture, and cognition meet*, pages 69–84, 2003.
- [18] P. Y. Oudeyer and F. Kaplan. Discovering communication. *Connection Science*, 18(2):189–206, 2006.
- [19] E. Oztot and M. Arbib. Schema design and implementation of the grasp-related mirror neuron system. *Biological Cybernetics*, 87(2):116–140, 2002.
- [20] E. Oztot, M. Kawato, and M. Arbib. Mirror neurons and imitation: A computationally guided review. *Neural Networks*, 19(3):254–271, 2006.
- [21] F. Perotto, J. Buisson, and L. Alvares. Constructivist anticipatory learning mechanism (CALM) - dealing with partially deterministic and partially observable environments. *Proceedings of Epigenetic Robotics 2007*, 2007.
- [22] G. Rizzolatti, L. Fadiga, V. Gallese, and L. Fogassi. Premotor cortex and the recognition of motor actions. *Cognitive brain research*, 3(2):131–141, 1996.
- [23] L. Steels, F. Kaplan, A. McIntyre, and J. Van Looveren. Crucial factors in the origins of word-meaning. *The Transition to Language*, 2(1):4–2, 2002.
- [24] M. Tettamanti, G. Buccino, M.C. Saccuman, V. Gallese, M. Danna, P. Scifo, F. Fazio, G. Rizzolatti, S. F. Cappa, and D. Perani. Listening to Action-related Sentences Activates Fronto-parietal Motor Circuits. *Journal of Cognitive Neuroscience*, 17(2):273–281, 2005.
- [25] M. Tomasello, M. Carpenter, and U. Liszkowski. A new look at infant pointing. *Child Development*, 78(3):705–722, 2007.
- [26] M. Umiltà, E. Kohler, V. Gallese, L. Fogassi, L. Fadiga, C. Keysers, and G. Rizzolatti. I Know What You Are Doing A Neurophysiological Study. *Neuron*, 31(1):155–165, 2001.
- [27] L. S. Vygotsky. *Mind In Society*. page 56, 1978.
- [28] L. S. Vygotsky. *Mind In Society*. pages 84–91, 1978.