

Robot Learning of Holophrases, Words and Proto-Grammar from Simulated Babbling and Physical Interaction

Joe Saunders, Caroline Lyon, Chrystopher L. Nehaniv, Kerstin Dautenhahn, and Frank Förster

Adaptive Systems Research Group

University of Hertfordshire

College Lane, Hatfield, Herts AL10 9AB, United Kingdom

Email: {J.L.Saunders,C.M.Lyon,C.L.Nehaniv,K.Dautenhahn,F.Forster}@herts.ac.uk

Abstract—It is thought that meaning may be grounded in early childhood language learning via the physical and social interaction of the infant with those around him or her, and that the capacity to use words, phrases and their meaning are acquired through shared referential ‘inference’ in pragmatic interactions. In this paper we report on experimental proposals which would allow a robot to carry out language learning in a manner analogous to that in early child development. In order to create appropriate conditions for language learning it would therefore be necessary to expose the robot to similar physical and social contexts. However in the early stages of language learning it is estimated that a 2-year-old child can be exposed to as many as 7,000 utterances per day in varied contextual situations. In this paper we present a series of experiments which we hope will ‘short cut’ holophrase learning in physical robots. This is achieved by moving from, firstly, simulated babbling through mechanisms which will yield basic word or holophrase structures, to, secondly, an interaction environment between a human and a robot where shared ‘intentional’ referencing and the associations between physical, visual and speech modalities can be experienced by the robot. The output of these experiments, combined to yield word or holophrase structures grounded in the robot’s own actions and modalities, would provide scaffolding for further proto-grammatical usage-based learning via interaction with the physical and social environment involving human feedback to bootstrap developing linguistic competencies. These structures would then form the basis for further studies on language acquisition, including the emergence of negation and more complex grammar. This paper describes progress on these proposals to date and planned future studies.

I. INTRODUCTION

In learning to use language to communicate and manipulate the world around them, human children benefit from a positive feedback loop involving individual learning (by interacting with their hands and bodies with objects around them), social learning (via close interaction with parents and others), and gradual acquisition of linguistic competencies. This feedback cycle supports the scaffolding of increasingly complex skill learning and linguistic development giving the child ever greater mastery of its social and physical environment, as well as supporting the development of cognitive and conceptual capabilities that would seem impossible without language. Our work is aimed at realizing this same kind of feedback cycle supporting that scaffolding of behavioural, linguistic and

conceptual competencies in robots. The purposes of doing this are not only to better understand possible mechanisms for such learning in humans, but also to achieve similar competencies in artificial agents and robots (even if they are not acquired by exactly the same routes).

II. FROM BABBLING TO THE ACQUISITION OF WORDS AND PHRASES WITHOUT MEANING

This work is inspired by the observed progress in language acquisition by human infants. Though we do not aim to simulate this development as a whole we investigate certain mechanisms that could play a key role. We make the artificial assumption that we can isolate different developmental paths and examine them separately. The development of the ability to segment a speech stream into words and phrases overlaps with the acquisition of semantic understanding and the recognition of primary language structure, but here we initially just investigate the emergence of stable phonetic forms or strings independent of meaning. Our simulation aims to show how a robot might learn to segment an utterance, an acoustic stream of sounds, that will be the basis for extracting meaningful elements [1, p. 39]. The starting point is taken as the stage at which infants typically produce canonical babbling, in the second half of their first year. This is the age at which infants begin to show they are learning the phonetic characteristics of their own ambient language [2, p. 46], [3, p. 148]. Earlier the ability to distinguish different phonetic contrasts is universal, but sensitivity to some foreign phonemes is then lost or diminished – for instance the capacity of infants learning Japanese to distinguish between /r/ and /l/ as both sounds belong to the same Japanese phoneme [r], while discrimination of native sounds improves [4].

Initially, the infant’s ability to perceive and analyse acoustic signals is much greater than the ability to produce them. Infants typically start to acquire vowel categories of their native language and combine them with certain consonants to produce characteristic repetitious strings of syllables *dadada*, *mamama*, etc. typical of babbling. This is the basis on which the infant will build its language production abilities [5, p. 50] as “[t]he first (pre-symbolic, pre-referential, context-limited)

words produced reflect a match between the child’s babbling patterns and adult patterns produced in a meaningful context” [6, p. 136].

In the acquisition of linguistic capabilities, e.g. mastery of word-usage frames and case markings, there are analogies with holistic utterances hypothesized as components of protolanguage in early evolutionary stages [7], [8].

From the start we will set up a dialogue, with turn taking, between a Linguistically Enabled Synthetic Agent (LESA) and a human, or simulated human, teacher. We make certain assumptions and then see whether they suffice for aspects of autonomous development of language abilities in the LESA. The assumptions include:

- LESA has the intention to communicate.
- Communicative ability is learnt via interaction with a teacher.
- Perception and production of speech are based on simulated ‘mirror neuron’ type structures - i.e., the same elements reflect components of perceived speech and generate synthesized speech.

A. Experiment I. Phonemics and Holophrasis Acquisition

For these preliminary investigations LESA will be a software agent. The teacher will be partly an actual human, partly a simulated human so that we can increase the quantity of input. From the start we aim to eventually have LESA “understand” a co-operative but naive human – as evidenced by speaking and acting appropriately in context. Therefore, we do not just want to work with stylized sentences, but with the vagaries of natural language.

Using the SCRIBE corpus of phonetically rich English [9] a store of phonemes, diphones and demi-syllables will be implemented, simulating a component of LESA’s phonemic knowledge. Other sources could be used for different languages. The structure of the syllable as perceived by infants will be modelled [2, p. 27]. This is analogous to the result of an infant constantly hearing English speech around him or her, not necessarily infant-directed speech (IDS). At this initial stage in our work the structure of memory and knowledge fusion processes are simplistic.

For the case of English, the LESA will start by combining certain consonants and vowels into syllables, using the range of phoneme combinations observed in infants acquiring English [10, p. 149], [11, p. 176], with appropriate weighted chances of being used. LESA will generate output strings of simulated babble in variable length utterances. The teacher will produce short English language utterances, typical of the length found in motherese as in the CHILDES database [12], which are converted and encoded into phoneme strings. The phonemic alphabet includes a start/end-of-utterance indicator.

To begin with there is no correspondence between utterances: interaction between LESA and its teacher is confined to turn taking, and to updating the weight attached to phonemes, diphones, demi-syllables in LESA’s store of phonemic knowledge. Then learning to segment an utterance and extract elements begins. Each time LESA ‘hears’ a phoneme string it

makes an attempt to segment it, using mechanisms described below. These segments constitute new components of the store of perceived phonemic and ‘pre-lexical’ knowledge, and are available for LESA to use in future productions. With a simulated teacher, LESA’s output is checked against a phonetic dictionary to see if it constitutes a proper word or contains a proper word along with other phonemes, or contains more than one word. If it does the teacher will respond positively, when it first occurs, and the segment, a string of phonemes, will be stored. This learnt item can then be appropriately weighted and produced as part of a future utterance.

With a real human teacher, whose speech would be converted into strings of phonemes, LESA’s output would be synthesized speech, and the human would make a subjective judgement.

Taking serial learning of short segments as target basic building blocks is consistent with neuroscientific evidence [5], [13] and empirical linguistic observation [14].

B. Segmentation Mechanisms

In making an attempt to segment a stream of sounds we will take into account the following characteristics of the process in humans.¹ At this stage IDS (infant directed speech) is typically made up of words with few syllables [12].

Salient factors are start and end of utterance markers, of which the latter is the most important. Mothers typically place words on which they want the infant to focus at the end of an utterance. If this word is a noun in English this is usually a grammatical form, but, for instance, in Turkish it leads to ungrammatical productions. In spite of this the practice is widely observed, suggesting that at this stage word segmentation is more important than correct syntax [16].

Another factor is the frequency of certain phonemic combinations. Rare or unobserved sequences tend to indicate a word boundary. Infants of about 11 months have been shown to recognize the phonotactic constraints indicating word boundaries [17].

We will also take account of observed phenomena such as the prevalence in English of the unit known as a “minimal word” containing a binary foot, composed either of two syllables (e.g. CVCV, where C:consonant and V:vowel) or two moras (e.g. CVC or CVV).

We will investigate methods of integrating these factors to find the placement of word boundaries, via the use of simple, single layer neural networks and information-theoretic methods. Discrete sequential data processed as bigrams or trigrams is typically linearly separable [18].

III. ACQUISITION OF MEANING THROUGH MEDIATED PHYSICAL INTERACTION

Part of our research focuses on whether it is possible to associate speech and gestural actions of a human with

¹In this current work we do not investigate the role of prosody and intonation, which play a significant role in speech segmentation and enable infants to recognize some organized forms of speech from their earliest months [15].

action, visual, proprioceptive, and auditory perceptions of a robot in order to derive ‘meaning’ for the robot associated to perceived speech patterns.² In carrying out this research we take a ‘usage-based’ view on language acquisition following Tomasello [1] and Bloom [21] and take inspiration from the constructivist work of researchers such as Steels on the emergence of various linguistic capacities in agent communities [22], and also Roy and Pentland [23], Roy [24], and Yu and Ballard [25] in attempting to link perceived speech with object and action perception. One of the problems faced in an associative approach is that it requires the exploitation of statistical regularities between speech and perception. This implies that many learning episodes would be required. However, although infants are exposed to a high number of ambient speech events and direct feedback learning via their carers (being exposed to as many as 7,000 utterances per day), they actually learn new words with very few presentations. According to Bloom [21] co-occurrence of seeing an object and hearing the object name does not occur regularly enough to allow an associative mechanism to differentiate the object (e.g. 50% of the time: see ‘milk’, hear ‘milk’ - 50% of the time hear ‘milk’, see ‘cat’ - but ‘cat’ is never associated with ‘milk’). Similarly carers rarely provide feedback on actions (Bloom gives the example of the carer arriving home and saying ‘Hello baby, whatya been doing today’ rather than ‘Hello baby, I’m closing the door’). In order to achieve learning in few episodes the infant must have its learning experiences biased in some way. This is thought to occur via intentional reference. Thus for example 9-month old babies will follow the gaze of their mother, follow her pointing gestures and monitor her emotional states. By 1 year old, the infant points on their own and observes the adult’s gaze whilst checking if they have changed the adult’s attention. If they fail to capture the adult’s attention they will alternate between gazing at the object and the adult until they succeed in getting the attention of the adult onto the object. The utterances of the adult together with reinforcement via forms of affective feedback (possibly prosodic features in the adult response) and success/failure of shared intentional reference for the infant whilst situated in context allow enough bias for fast learning to take place. Within this shared context the action modalities of the child are also associated with simple verbal speech patterns of the adult (‘see the doggy’, ‘hold the bottle’). In our experiments learning will thus rely on the action modalities of the robot together with the robot’s ability to share context and referential ‘intent’ with a human teacher.

A. Software Platform and Architecture

Physical instantiations of a LESA can take humanoid form, such as the iCub [26] or Kaspar2 [27] robots. In carrying out these studies we have modified our existing social learning architecture (ROSSUM [28], [29]), originally implemented on wheeled mobile robots, to the Kaspar2 humanoid robotic platform (see figure 1). The system has been used to learn

²Following Wittgenstein [19], any derivation of ‘meaning’ must ultimately be evidenced by appropriate embodied action in language games (cf. [20]).

scaffolded behaviours via directed learning from a human teacher [28] (but not previously linguistic behaviour). We have added additional functionality for face, motion, and colour detection together with a facility for recording the phoneme sequences made by the human teacher. These new modalities are in addition to existing object detection and proprioceptive feedback modalities already present in the architecture (see III-B below).



Fig. 1. *Learning via Social Interaction.* The humanoid robot Kaspar2 is taught how to hold a patterned box.

ROSSUM [29] allows experiments to be carried out where, via the processes of self-imitation and observational imitation, the humanoid robot learns through experiences grounded³ on its own visual, auditory and sensorimotor feedback the relevant interaction modalities presented by a human tutor. Initially learning would be based on interactions between the human tutor and the robot whereby the tutor reveals to the robot, via speech, deixis, gesture and reference, the various affordances that are available and the relevant effectivities that can be used to exploit these affordances. Here *speaking* is regarded as a particular type of *gesturing*, i.e. motor activity for manipulating the physical or social environment (see [31], [32]). The combination of self- and observational imitation would allow the robot to take both a first and third person perspective in making these discoveries and associate the gestural components with these discoveries.

B. Robot

The first physical experiment due to be carried out will use the University of Hertfordshire’s Kaspar2 robot. This is a minimally expressive small humanoid robot with 8 degrees of freedom in the head, and 5 degrees of freedom in each arm. Proprioceptive feedback from each arm is available both when powered (and maintaining a position) and when unpowered. The person interacting with Kaspar2 when unpowered can physically manipulate its arms (for example to make the robot reach for an object) and the proprioceptive actuator readings can be continuously recorded.

³The idea of ‘grounding’ refers to relating the meaning of symbols to embodied sensorimotor experience [30].

Kaspar2 employs two video cameras for its eyes. These are used to obtain images which are processed [33] to yield an additional set of modalities. These include object/pattern recognition, object colour, face detection and motion detection. A number of objects are available for Kaspar2 to interact with. These objects are detected using the ARToolKit system [34]. The objects are pre-learned using ARToolKit, thus Kaspar2 can detect these objects and recognize that they are individual entities in the world, however no other meaning is attached to them. We justify the use of this simplifying step in these early experiments as, firstly, it eliminates the need for a complex vision processing modality and secondly, and more importantly, it reflects a ‘whole object’ bias found both in children and adults (see Bloom, Chapter 4 [21]).



Fig. 2. *Sharing Reference with a Teacher in Context.* Kaspar2 interacts with the teacher, where both the robot and the human share attention on the coloured objects, a case of rudimentary shared intentional reference.

The robot also records the directed speech made by the human. The speech pattern of the human is then processed to yield a stream of phonemes using a system such as the University of Bielefeld ESMERALDA speech recognition system [35].

C. Experiment II. Grounding Speech - Action and Object Learning

The following experiment is currently under development (experiment II). This will have a human teacher/carer interacting with Kaspar2. The human and robot will sit on opposite sides of a desk (see figure 2). Kaspar2 will initially focus on the human’s face. If the human stays still Kaspar2 will eventually ‘get bored’ and start to look around. If the human changes his/her face direction (to look at an object) Kaspar2 will broadly follow the line of the human’s gaze. During this interaction the human will be encouraged to ‘talk’ to Kaspar2 and move Kaspar2’s arms and hands to push, pull and touch the objects whilst simultaneously explaining the action to Kaspar2. All of Kaspar2 modalities will be actively recorded during this period together with the phoneme sequences generated from the analysis of the human speech.

Subsequently the recorded data is analysed firstly to look for clusters of association occurring between the phonemic

strings and the ‘experiences’ (i.e. sensorimotor readings over a time window [36]) in the robot’s own modalities. We intend to analyse the data using variants of Crutchfield-Renyi information distance measure [37] as this has shown to effectively associate sensory data from differing modalities in our previous studies [36], [38], [39]). Subsequently, we will additionally use the database of phoneme strings generated by the experiment in section II above (experiment I). We expect that by using this data the clustering of associations will be biased to the holophrases that the robot is already ‘aware of’ and to which it will subsequently associate to particular classes of objects, actions, or interactions.

D. Assumptions

In carrying out the physical experiments above we make a number of key assumptions. Firstly, that the LESA, embodied as Kaspar2 in this case, is motivated by novelty - thus if the LESA sees a new object or a human face it will continue looking at it. Kaspar2’s focus will eventually habituate, after which it will move semi-randomly (randomly but within fixed limits) until it finds a new focus of attention. If it sees the same object again Kaspar2 will fixate on it only if a sufficient time has passed. This ‘boredom’ threshold is currently set at ~50 seconds. Secondly, that the robot is motivated to share the same attention space as the human trainer. Thus if the human decides to look away from Kaspar2 the robot will try and focus in the same general area as the human. We achieve this last step by initially allowing the robot to focus on the human face. If the human head moves a motion detection system will provide a global movement vector, essentially yielding a single integer value describing the angle of movement of the human head. Kaspar2 will then move appropriately with its own head to broadly focus on the same area that the human is looking at.

E. Experiment III: Speaking and Interacting

The output from analysis of the data from experiment II should yield a set of clusters which directly associate the phonic strings, word sequences, or holophrases with a grounded representation of the robot’s perceptions. We will then use these clusters to re-run experiment II. This time however the robot will ‘talk’ by pronouncing the associated phoneme strings of any cluster recognized as similar to the current situations (e.g. familiar objects/actions). This phonetic output will be via speech synthesis (using, e.g., the eSpeak system [40]). The aim here is to analyse the effect of human feedback on the robot’s learning and on scaffolding further linguistic competencies. Thus if the human hears the robot say, for example, ‘push circle’ when pushing a circle object, we would expect the human to provide some reinforcement signal back to the robot, for example ‘yes, Kaspar pushes the circle!’ or ‘No, that’s a square!’. This would then allow further analysis to yield valuable information on the nature of the reinforcement and also of negation signals. This allows us to study the emergence of various forms of negation [41] through the mechanisms of communicative social interaction;

indeed, negation has been hypothesized to have been an extremely important qualifier in the emergence of symbolic representation capabilities.

IV. DISCUSSION, QUESTIONS, AND CONCLUSION

Insights of [19] and [42], and more constructively [22], [43], suggest to understand signalling and linguistic behaviour, one needs to take into account usage in their pragmatic embodied social context. The learning of communicative signalling and linguistic systems (at the ontogenic, diachronic, and evolutionary levels) are moreover shaped, not only by details of perception and embodiment (e.g. [44]), but also by details of transmission, sources of error and variability, as well as feedback and repair mechanisms (e.g. [7], [22], [45]).

We have outlined mechanisms for experiments whereby Linguistically Enabled Synthetic Agents (LESAs), are expected to exhibit (1) reinforcing holophrasis and learning of word-level parsing, (2) the grounding of words and lexicon usage frames in action and object learning via physical interactions, and (3) the bootstrapping of simple usage-based proto-grammatical structure via human scaffolding and feedback.

The overall approach is to understand *constructively* what mechanisms could be responsible for the ontogeny of linguistic competencies. That is, for such a constructive theory of language to be successful it is necessary to build an instantiation that exhibits the phenomenon to be explained, and, moreover, different constructive mechanisms could be assessed against each other by comparing what they actually generate. Most preferably these constructivist evaluation test-beds must involve learning in embodied social interactions with humans and physical interactions with rest of the LESA's environment.

Open and challenging research questions in this area include: (1) to what extent can the methods be scaled for human-like acquisition of linguistic abilities?, (2) to what extent does the achievement of one of these stages support the next in autonomous robot learning in social interaction with humans?, (3) are the methods applicable for a human language, say at the level of 3-year old child, or only proto-grammatical approximants thereof?, (4) what 'cognitive' capabilities are necessary or sufficient for recruitment in the development of human-like linguistic competencies?, (5) is it necessary to build in universal mechanisms for categorization and generalization, propositional logic, predication, compositional syntax, etc., or can these emerge from more elementary processes (such as Hebbian learning, 'chunking', sequential processing and locality principles, etc.) or more general cognitive capacities (such as perspective taking; action hierarchies; expectation, prospection and refusal)?, (6) how can different types of linguistic negation be acquired by a LESA?, (7) to what extent are these mechanisms for the development of linguistic abilities universal, i.e. applicable for any given target natural language? (8) what are appropriate semiotic frameworks for pragmatic acquisition of language usage (e.g. fluid construction grammar [46], dynamic syntax [47], etc.)?, (9) to what extent are purported explanations consistent not only with individual

ontogeny of linguistic capabilities but also with diachronic (transmission) and evolutionary (phylogenetic) considerations?

The road ahead to understanding language emergence is long and complex, but constructive approaches offer new means and criteria for validating explanations as we progress on this path.

ACKNOWLEDGEMENT

The work described in this paper was conducted within the EU Integrated Project ITalk ("Integration and Transfer of Action and Language in Robots") funded by the European Commission under contract number FP7-214668.

REFERENCES

- [1] M. Tomasello, *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Harvard University Press, 2003.
- [2] B. de Boisson-Bardies, *How Language Comes to Children*. MIT, 1999.
- [3] P.-Y. Oudeyer, *Self-organization in the Evolution of Speech*. Oxford University Press, 2006.
- [4] J. F. Werker and R. C. Tees, "Cross-language Speech Perception: Evidence for Perceptual Reorganization during the First Year of Life," *Infant Behavior and Development*, vol. 7, pp. 49–63, 1984.
- [5] F. Pulvermuller, *The Neuroscience of Language*. Cambridge University Press, 2002.
- [6] M. M. Vihman and R. A. Depaolis, "The Role of Mimesis in Infant Language Development: Evidence for Phylogeny?" in *The Evolutionary Emergence of Language*, C. Knight, M. Studdert-Kennedy, and J. R. Hurford, Eds. Cambridge University Press, 2000.
- [7] A. Wray, "Protolanguage as a holistic system for social interaction," *Language and Communication*, vol. 18, pp. 47–67, 1998.
- [8] —, " 'Needs only' Analysis in Linguistic Ontogeny and Phylogeny," pp. 53–70, in [48].
- [9] SCRIBE, "Spoken Corpus Recordings In British English," <http://www.phon.ucl.ac.uk/resource/scribe/scribe-manual.htm>, 2004, [last visited 25 July 2008].
- [10] P. F. MacNeilage and B. L. Davis, "Evolution of Speech: the Relation between Ontogeny and Phylogeny," in *The Evolutionary Emergence of Language*, C. Knight, M. Studdert-Kennedy, and J. R. Hurford, Eds. Cambridge University Press, 2000.
- [11] K. Demuth, "The Prosodic Structure of Early Words," in *Signal to Syntax*, J. Morgan and K. Demuth, Eds. Lawrence Erlbaum, 1996.
- [12] CHILDES, <http://chilides.psy.cmu.edu/>, 2008, [last visited 25 July 2008].
- [13] V. Braitenberg, D. Heck, and F. Sultan, "The detection and generation of sequences as a key to cerebellar function: experiments and theory," *Behavioral and Brain Sciences*, vol. 20, pp. 229–245, 1997.
- [14] C. Lyon and C. L. Nehaniv, "Developing agents that can speak with humans: pointers from the evolution of language," School of Computer Science, University of Hertfordshire, Tech. Rep. 479, June 2008. [Online]. Available: <http://homepages.feis.herts.ac.uk/~comrcml>
- [15] J. Morgan and K. Demuth, "Signal to syntax: an overview," in *Signal to Syntax*, J. Morgan and K. Demuth, Eds. Lawrence Erlbaum, 1996.
- [16] R. N. Aslin, J. Z. Woodward, N. P. LaMendola, and T. G. Bever, "Models of word segmentation in fluent maternal speech to infants," in *Signal to Syntax*, J. Morgan and K. Demuth, Eds. Lawrence Erlbaum, 1996.
- [17] J. Myers, P. W. Juszyk, D. G. Kemler-Nelson, J. Charles-Luce, A. Woodward, and K. Hirsch-Pasek, "Infants' sensitivity to word boundaries in fluent speech," *Journal of Child Language*, vol. 23, pp. 1–30, 1996.
- [18] C. Lyon and R. J. Frank, "Using Single Layer Networks for Discrete, Sequential Data: An Example from Natural Language Processing," *Neural Computing and Applications*, vol. 5, no. 4, pp. 196–214, 1997.
- [19] L. Wittgenstein, *Philosophical Investigations (Philosophische Untersuchungen)* – German with English translation by G.E.M. Anscombe, 3rd ed. Basil Blackwell, 1968, (first published 1953).
- [20] C. L. Nehaniv, "Meaning for observers and agents," in *IEEE International Symposium on Intelligent Control/Intelligent Systems and Semiotics (ISIC/ISAS'99)*, 1999, pp. 435–440.
- [21] P. Bloom, *How Children Learn the Meaning of Words*. MIT Press, 2002.
- [22] L. Steels, "The origins of syntax in visually grounded robotic agents," *Artificial Intelligence*, vol. 103, no. 1-2, pp. 133–156, 1998.
- [23] D. Roy and A. Pentland, "Learning words from sights and sounds: A computational model," *Cognitive Science*, vol. 26, pp. 113–146, 2002.

- [24] D. Roy, "Grounding words in perception and action: Computational insights," *Trends in Cognitive Sciences*, vol. 9, no. 8, pp. 389–396, Aug. 2005.
- [25] C. Yu and D. Ballard, "A multimodal learning interface for grounding spoken language in sensorimotor experience," *ACM Transactions Applied Perception*, vol. 1, pp. 57–80, 2004.
- [26] iCub, "RobotCub – An Open Framework for Research in Embodied Cognition," <http://www.robotcub.org/>, 2004.
- [27] M. Blow, K. Dautenhahn, A. Appleby, C. L. Nehaniv, and D. Lee, "The art of designing robot faces - dimensions for human-robot interaction," in *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction (HRI 2006), Salt Lake City, Utah, USA, March 2-3, 2006*, M. A. Goodrich, A. C. Schultz, and D. J. Brummer, Eds. ACM, 2006, pp. 321–322, see also <http://kaspar.feis.herts.ac.uk/>.
- [28] J. Saunders, C. L. Nehaniv, and K. Dautenhahn, "Teaching robots by moulding behavior and scaffolding the environment," in *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction (HRI 2006), Salt Lake City, Utah, USA, March 2-3, 2006*, M. A. Goodrich, A. C. Schultz, and D. J. Brummer, Eds. ACM, 2006, pp. 118–125.
- [29] J. Saunders, C. L. Nehaniv, K. Dautenhahn, and A. Alissandrakis, "Self-imitation and environmental scaffolding for robot teaching," *International Journal of Advanced Robotic Systems*, vol. 4, no. 1, pp. 109–124, March 2007, special issue supplement on Human-Robot Interaction. [Online]. Available: <http://www.ars-journal.com/International-Journal-of-Advanced-Robotic-Systems/Volume-4/ISSN-1729-8806-4115.pdf>
- [30] S. Harnad, "The symbol grounding problem," *Physica D.*, vol. 42, pp. 335–346, 1990.
- [31] C. L. Nehaniv, K. Dautenhahn, J. Kubacki, M. Haegele, C. Parlitz, and R. Alami, "A methodological approach relating the classification of gesture to identification of human intent in the context of human-robot interaction," in *14th IEEE International Workshop on Robot and Human Interactive Communication (Ro-Man 2005)*, 2005, pp. 371–377.
- [32] N. Otero, C. L. Nehaniv, D. S. Syrdal, and K. Dautenhahn, "Naturally occurring gestures in a human-robot interaction teaching scenario," *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, vol. 9, no. 3, 2008 (in press).
- [33] openCV, <http://opencvlibrary.sourceforge.net>, 2006, [last visited 30 June 2008].
- [34] ARTToolkit, <http://www.hitl.washington.edu/artoolkit>, 2003, [last visited on 30 June 2008].
- [35] ESMEERALDA, "An environment for statistical model estimation and recognition on arbitrary linear data arrays," University of Bielefeld, Applied Computer Science, <https://www.techfak.uni-bielefeld.de/ags/ai/projects/ESMEERALDA/index.html>, 2005, [last visited 31 July 2008].
- [36] N. A. Mirza, "Grounded sensorimotor interaction histories for ontogenetic development in robots," Ph.D. dissertation, Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire, 2008.
- [37] J. P. Crutchfield, "Information and its Metric," in *Nonlinear Structures in Physical Systems – Pattern Formation, Chaos and Waves*, L. Lam and H. C. Morris, Eds. Springer Verlag, 1990, pp. 119–130.
- [38] L. Å. Olsson, "Information self-structuring for developmental robotics: Organization, adaptation, and integration," Ph.D. dissertation, Adaptive Systems Research Group, School of Computer Science, University of Hertfordshire, 2006.
- [39] N. A. Mirza, C. L. Nehaniv, K. Dautenhahn, and R. te Boekhorst, "Grounded sensorimotor interaction histories in an information theoretic metric space for robot ontogeny," *Adaptive Behavior*, vol. 15, no. 2, pp. 167–187, 2007.
- [40] eSpeak, <http://espeak.sourceforge.net/>, 2007, [last visited 31 July 2008].
- [41] C. L. Nehaniv, C. Lyon, and A. Cangelosi, "Current Work and Open Problems: A Road-Map for Research into the Emergence of Communication and Language," pp. 1–27, in [48].
- [42] R. G. Millikan, *Varieties of Meaning*. MIT Press, 2004.
- [43] L. Steels, "The Recruitment Theory of Language," pp. 129–150, in [48].
- [44] A. Cangelosi and D. Parisi, "The emergence of a 'language' in an evolving population of neural networks," *Connection Science*, vol. 10, no. 2, pp. 83–97, 1998.
- [45] K. Smith, H. Brighton, and S. Kirby, "Complex systems in language evolution: the cultural emergence of compositional structure," *Advances in Complex Systems*, vol. 6, no. 4, pp. 537–558, 2003.
- [46] L. Steels and P. Wellens, "How grammar emerges to dampen combinatorial search in parsing," in *Symbol Grounding and Beyond*, P. Vogt, Y. Sugita, E. Tuci, and C. Nehaniv, Eds., vol. 4211 of Lecture Notes in Computer Science. Springer, 2006, pp. 76–88.
- [47] R. M. Kempson, W. Meyer-Viol, and D. M. Gabbay, *Dynamic Syntax: The Flow of Language Understanding*. Blackwell, 2001.
- [48] C. Lyon, C. L. Nehaniv, and A. Cangelosi, Eds., *Emergence of Communication and Language*. Springer, 2007.