Deixis and the Development of Naming in Asynchronously Interacting Connectionist Agents

Jean Baillie **Chrystopher Nehaniv**

Department of Computer Science
University of Hertfordshire
Hatfield, Hertfordshire AL10 9AB
U.K.
email: {j.baillie, c.l.nehaniv}@herts.ac.uk

Abstract

In this paper we describe a series of interaction games in which an elementary 'proto-language' is generated based on innate deictic ability in a community of asynchronously interacting connectionist agents. Deixis—or pointing (and its perception)—combined with interaction and 'speech' generation ability support the emergence and drift of shared names within a family of 'parent' and 'child' agents. An agent may point at a 'referent' agent and then associate any perceived speech with the referent. Perceived speech is either received as a response from an interlocutor or else the agent itself may generate speech, if it points but hears no 'name' in response. Deixis toward a referent followed by perceived speech is associated ('learned') by the agent's recurrent time-delay connectionist artificial neural network establishing a basis for naming. Therefore subsequent deixis with the same referent triggers attempted reproduction of the associated ensuing speech. The interaction may also be 'overheard' by other agents which then may associate the particular deitic acts perceived with the speech heard, resulting in a propagation of naming conventions.

Interaction, via pointing and speaking, is asynchronously scheduled, and in this respect biologically plausible. Issues related to mirror-neurons, attention and deictic gaze, development of vocabularies, as well as synonomy, lexical drift and convergence are also discussed briefly. This illustrates the dynamics and development of deixis-grounded naming systems in (asynchronous) interaction and thus extends the work of Steels on the emergence of vocabularies in communities of agents as well as the work of Billard, Hayes, and Dautenhahn on the grounding and learning vocabularies in robots and agents using connectionist methods.

1 Introduction and background motivation

In this paper we describe a series of experiments in proto-language development, based on the first of the two hypotheses described in (Steels 1996), namely, that 'language is an autonomous adaptive system which forms itself in a self-organizing process'. At the start of the experiment no language exists, though the facilty for deixis and generation of phonemes does exist in a family of connectionist agents.

Controversies on the existence and role of a possible innate language acquisition device (e.g. (Chomsky 1957, Chomsky 1968, Chomsky 1975, Chomsky 1981, Bickerton 1990, Pinker & Bloom 1990) for human language as opposed to language-readiness based on other, more general cognitive capabilities (e.g. (Steels 1996, Hurford, Knight & Studdert 1998, Arbib in press)) motivate the study of the origin and maintenance of language and language-like phenomena in animals (e.g. (Savage-Rumbaugh & Brakke 1996, Herman & Austad 1996, Bickerton 1990, Pepperberg in press)) and artifacts (e.g. (Steels 1996, Steels 1997, Steels 1998*b*, Hurford et al. 1998, Billard & Hayes 1999, Billard & Dautenhahn 1999, Nehaniv 2000)). The various degrees to which non-human animals exhibit various such capacities suggest that many linguistic and related phenomena can be explained in a manner parsimious also with human language capacties, as part of a general cognitive ethology of 'animal minds' (Griffin 1976, Griffin 1992) that does not require treating human abilities as somehow discontinous from the rest of nature.

Various researchers have now attempted to address linguisitic category formation, phonology, syntax, and aspects of semantics in software and robotic simulations in order to study the extent to which simple mechanisms can serve to explain some of the observed natural phenomena (see e.g. (Steels 1996, Hurford et al. 1998)). Some of this body of work is surveyed and evaluated in (Nehaniv

2000), with an eye to distinguishing carefully whether and to what degree solutions to the problem of generating a particular language-like phenomenon have been built-in. For example, no artificial system has so far been shown to develop predication abilities, nor have constituent structures of syntax been shown to emerge without these having been tacitly (and often unknowingly) built-in by its implementers.

Against this background, we study the development and dynamics of the use of naming in the course of asynchronous interaction of a family of agents endowed with simple deixis and 'speech' production/perception capabilities and possessing only simple connectionist cognitive capabilities.

Deixis—or pointing—via gestures or gaze direction, or also by the use of other signals, including words like *this* or *those*, can serve as a substrate for joint attention that grounds the development of linguistic and social abilities in humans as they grow from pre-linguisitic infants into adults (cf. (Scaife & Bruner 1975, Butterworth 1991, Moore & Dunham 1995)). Roboticists have begun to implement the mechanisms of joint attention (such as gaze detection and deictic gaze) as steps in building up the social competencies of these robots (Scassellati 1999) (compare also (Dautenhahn 1994, Dautenhahn 1997) regarding the introduction of social intelligence into agents and robots). We claim that deixis will also serve as a useful substrate component for the implementation and development in a social context of linguistic capabilities of robots and agents.

To a certain extent, this claim is already supported by work of (Steels 1996) illustrating the 'self-organization' of a common vocabularies in a community of agents. His work tacitly uses deixis, but is limited by the use of synchronous, pair-wise interactions in a noise-free software environment. Moreover, Steels' agents appear to be implemented essentially as modifiable look-up tables in a functional programming language. These properties are not very plausible biologically, but we will show that his results nevertheless do generalize, with some modifications, to asynchronously interacting connectionist agents (an abstract class of agents of which biological agents are also instances).

In this paper, we use agents implemented instead in a recurrent time-delay connectionist architecture and do not impose any deterministic or synchronous scheduling protocol on interactions. For instance, an agent may not be listening to its interlocutor when the latter is speaking, and indeed the agent may be generating a new name if its attention was elsewhere (i.e. was not scheduled or was scheduled with a different task than listening). These types of networks have been used successfully in the acquistion of an existing vocabulary by learner robot from a teacher robot (with different sensors and embodiment) in experiments of (Billard & Dautenhahn 1999). Our work differs from this in that our agents start with no preexisting vocabulary, but generate 'speech' output when there is no perceived response to their deictic acts. Even in this case, and without a reinforcement feedback mechanism, we show how the connectionist architecture and deixis support the emergence of shared vocabularies of names for the referents of the deictic acts. An ontogenetic advantage of the development of naming, in addition to the deictic capability of the agent, is in the ability it affords for referencing agents which can no longer be seen and therefore cannot be pointed to. This reference to absent entities can be grounded in deixis and the simple 'cognitive' capabilities of a simple time-delay recurrent neural network.

1.1 The connectionist architecture

DRAMA (Billard & Hayes 1999) is a time-delay recurrent neural network which uses Hebbian update rules and which was designed for dynamic control and learning of autonomous robots. In this paper we use a implementation of the architecture, written in Java and operating in a concurrent (or more accurately, multithreaded) Unix environment, in an application which looks at the development over time of a simple proto-language consisting of shared—or somewhat shared—names.

Connections in DRAMA are associated with two weight parameters: a confidence factor, ω, modelling the frequency of correlated activation of any two units; and a time parameter τ which makes correlations between delayed and simultaneous occurrences of different input patterns. The network is fully recurrent and non-symmetrical and the weight parameters record separately the spatial and temporal features of the input patterns. For a full account of the architecture, see (Billard & Hayes 1999). In the case studies described in (Billard & Hayes 1999, Billard & Dautenhahn 1999) which involved learner and teacher robots, not only was it necessary to make sensor/actuator associations but it was also important that agents 'remember' the delay experienced between the events thus associated. This is exploited here.

2 Overview of the experiments

2.1 The paradigm

We wish to construct a community of interacting DRAMA agents in which individual agents may focus their attention on various others (including themselves) and in which agents may join or leave the community at will. As in human communities, agents may have little control over which other agents are focusing on them—indeed there may be no-one listening at all—and may change their own focus over time. They may also miss a communication, even from the agent on which they are focusing, if their attention was not engaged at the critical moment. Against such a background, we wish to observe the development of naming in the agents through *interaction games* (Steels 1996, Nehaniv 1999) involving deixis and response, and in which the DRAMA agents are themselves referents.

2.2 The experimental topology

In the particular experiments reported on here, there are four DRAMA agents involved in the game, numbered 0 through 3, where the numbering enables the deictic phase of the experiment—agents can distinguish between one another. Each agent has a sensor, and an output buffer to which the values of its network's nodes at the end of each processing cycle are written. In these simple experiments an agent focuses on exacly one other agent in the community throughout an interaction game.[1] The number of agents taking part is fixed; agents joining and leaving the community is not implemented here but is left for future development. Agents may or may not perceive deixis and speech from the agent they are focused on, depending on the details of the asynchronous processing.

In each experiment two of the agents engage in a dialogue; we may term these 'parents'. Each parent speaks to, and listens to, the other (that is, the sensor of each is focused on the output buffer of the other). The other two agents are the 'children'; these two listen to one of the parents, deemed to be the mother (that is, each child's sensor is focused on the output buffer of the mother). The children's output is not the focus of any other agent. The children behave as if engaged in the dialogue, that is, they generate referents and sometimes names, but the parents—and indeed the other child—are not listening, though the children are unaware of this (arguably perhaps a not-unrealistic model). The mother is influenced only by the other parent. This arrangement in effect results in that the parents may modify their deixis and speech associations in response to each other, but will not change their speech as a result of the children's utterances. The mother's use of speech and deixis is available for the children, but that of the other parent is only indirectly available via its influence on the mother's speech.

Modifications of the interaction topology, including a dynamically changing one, as well as allowing agents to join and leave the community, are certainly possible, and are interesting to study, but we shall focus on this illustrative case in this paper.

2.2.1 The mechanism

At the outset, each agent may (asynchronously) generate a referent to which it points, and then 'listens' for a response through its sensor. In practice for any given agent there is at least a 50% probability that no other agent is in fact listening, so any response received will not, from an external perspective, be a response in the strict sense at all, though the agent itself has no knowledge of that. Agents know which other agent they are listening to but not who is listening to them. If no response is received within a (predetermined) number of cycles, the agent might itself generate a response, though with a low probability. Intuitively, we can think of agents as 'asking' one another for the name, or some property, of the referent. They point at a referent and listen for a response. If a name is heard then it may be learned; if not, then one might be generated and output, and (possibly) learned by listening agents. The communication is permitted to continue for a fixed number of cycles (per agent: running processes are non-deterministically interleaved) after which all nodes are reset to zero and the process starts again with a newly generated referent. The state of each agent—that is, the values of the connection weights—is sampled for each agent at regular intervals during the agent's development and stored.

2.2.2 Retrieval

Retrieval is effected in a non-concurrent environment; each agent is treated as a lingusitic informant and interviewed separately from the others, with no further learning permitted during this phase. It is prompted with a referent and its responses recorded over time. The interrogation process is repeated for each referent. Sample developmental traces showing the evolution of the protolanguage at the same regular intervals as mentioned above are recorded (as in Tables 1, 2, and 3 below).

3 Event detection

The sensor consists of an event-detector which reads data from a designated agent's output buffer. An event is deemed to have occured if the value of an input data signal exceeds the previous value of that signal (stored in the sensor) by a predefined threshold. Although the incoming data is real-valued, as indeed is the threshold value, the actual values transmitted to the DRAMA network are binary—a DRAMA input node takes the value 1 if an event has been detected and 0 otherwise. Hence DRAMA inputs are always binary. This event-detection mechanism differs slightly from that described in (Billard & Hayes 1999); in addition, the scheduling algorithm employed by the operating system kernel may also have a bearing on event detection.

3.1 Asynchronous random interaction

In contrast with (Steels 1996), where processing is sequential, agents here are operating in a multi-threaded real-time environment, and so operating system scheduling issues form an integral part of the communication. It is a feature of the architecture that the values of DRAMA nodes decay by a fixed factor over time (that is, over a number of processing cycles, eventually dying altogether) in the absence of any event that 'refreshes' them. If a *read*

[1] It would also be desirable to study such games with agent focus dynamically changing within an experiment.

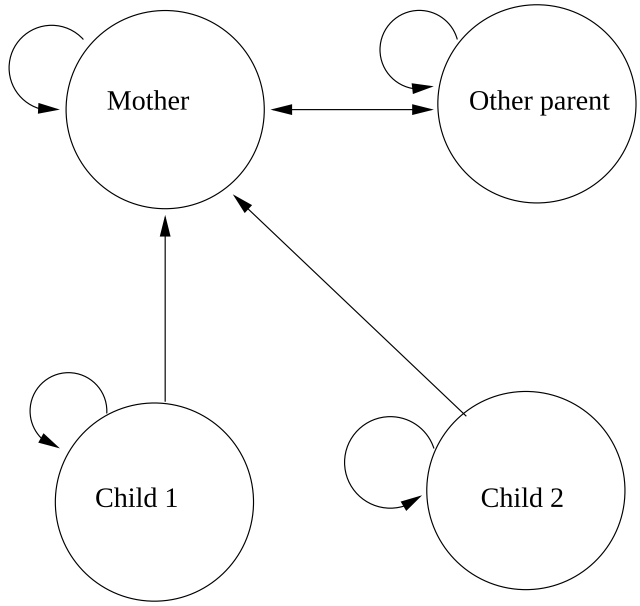

Figure 1: Schematic diagram showing agent communications. An arrow (→) should be read as ‘attends to’.

operation of an agent’s output buffer by a particular sensor is not scheduled for a number of processing cycles, a data value may well have decayed to below the point where an event is deemed to have occurred and so may not be detected when its reading is eventually scheduled. No attempt has been made to influence the scheduling algorithm or the time-slice in the Unix operating system or the Java Virtual Machine. The arbitrary nature of the scheduling is regarded as part of the process and enhances biological plausibility.

4 Assumptions

In the application described in this paper each sensor and its associated DRAMA agent consists of 14 network nodes (though of course it is not necessary for the sensor to have the same number of nodes as the network, and also there may be many sensors in an application). Notionally, the nodes are classified in two ways: nodes 0—3 are deictic, and nodes 4—13 are phonetic output. Events on the output nodes may be mapped onto the elements of a phonetic array; the phonetic ‘utterances’ are output when a corresponding node ‘lights up’, that is, takes the value 1. The resulting successive sounds are referred to as ‘proto-words’, and as ‘names’ if they are associated with deictic acts. Agents have no feedback as to whether linguistic behavior is “successful” (in contrast to (Steels 1996)). There is no positive or negative reinforcement; the recurrent neural network merely associates perceived deixis to utterances that it senses as occurring proximate in time.

4.1 Deixis

Each agent has a built-in ability to point at any other agent (or itself). The interactive phase for the DRAMA agents works as follows: a referent is randomly generated within the agent, that is to say, one, and only one, of the deictic nodes is set to 1. For example, agent 3 is pointed to when the third node is set. Any agent attending to that agent perceives the pointing as deixis to the appropriate referent agent.

4.2 Speech Output

As stated above, each parent agent listens to one other agent and the output from that agent is taken in via the listening DRAMA agent’s sensor, where events may be detected. This may or may not happen immediately, depending on when the process is next scheduled. If in response to deixis (that is, while the referent node still has a positive value) events are detected on the output nodes, then the association between the referent and the output is learned. If there is no response within a predetermined number of cycles, a response from the agent itself may be generated, since each agent perceives its own deixis.

A newly generated name takes the following form: two of the agent’s own output nodes are successively set to 1, that is, one of the consonant nodes 4—8 and one of the vowel nodes 9—13, ensuring (one might assume) that at retrieval time one consonant and one vowel are generated (such is the arrangement of the phonetic array onto which these nodes are mapped, though this is arbitrary). The consonants are b, p, f, k, t and the vowels a, e, i, o, u. After a fixed number of (asynchronous) processing cycles each agent’s nodes are reset to 0 and for several more

	Deixis/speech output association			
Agent	0	1	2	3
0†	-, -, fe	-, -, fe	-, ku, ku	-, fktaiu, fktaiu
1‡	-, -, -,	fe, fe, fe	-, bo, bkou	-, -, -
2	-, -, -,	fe, fe, fe,	-, bo, bkou	-, -, -
3	-, -, -	fe, fe, fe	-, bo, bkou	-, -, -

Table 1: Game 1

	Deixis/speech output association			
Agent	0	1	2	3
0	-, -, be, bfae, bfae	-, tu, tu, tu, btaeou	-, fi, fi, bfai, bfai	-, -, -, -, ke
1‡	-, -, be, be, bfae	-, tu, tu, tu, btaeou	-, fi, fi, bfai, bfai	-, -, bo, bo, kbeo
2†	-, be, fa, fa, be	pu, tu, tu, ba, bfkaeo	-, fi, fi, bfai, bfai	-, bo, bo, fa, keo
3	-, -, beo, bfaeo, bfao	-, tu, tue, tue, tau	-, fi, fi, bfai, bfai	-, -, be, be, ke

Table 2: Game 2

	Deixis/speech output association			
Agent	0	1	2	3
0	-, -, fi, fi, fi	-, -, -, eee, peo	-, po, po, po, po	-, po, bi, bi, bi
1‡	-, -, fi, fi, fi	-, te, te, te, peiou	-, po, po, po, po	-, -, bi, bi, bi
2	-, fi, fi, fi, t	-, te, kteu, pkteiuo, pteiu	to, pu, tpu, tpiu, pteiu	bi, bi, bi, bpi, bpi
3†	-, -, fi, fi, fi	-, te, te, te, pteo	-, -, po, po, po	-, btiu, btiu, btiu, btiu

Table 3: Game 3

cycles the network operates on inputs of patterns of 0's, flushing the system, before the process is repeated with a newly generated referent.

Interaction continues over several hundreds of cycles and the values of the two sets of connections for each agent are sampled and stored every 200 cycles. At the end of the interaction period, retrieval tests are run on stand-alone agents, created by loading the final state of the agents (that is, the connections) from the interaction sessions. The output from three such retrieval experiments is shown in tables 1—3 below. The number of the agent being tested is in the left hand column and the number of the referent along the top row. The output in the body of the tables shows the evolution of the proto-language over time every 200 cycles in each of three experimental interaction games[2].

5 Discussion of Results and Conclusions

From the tables some convergence in the proto-language can be seen in each experiment. Although syllables are generated in consonant-vowel pairs, they have temporal extent and are not necessarily spoken in that order. In the data, we see the emergence of more complex articulations than just simple consonant-vowel pairs, although this had not been intended. Moreover, we also observe lexical drift, slow change of the phonetic face of the proto-words, and replacement of one word by another in all or part of the community. Semantic drift cannot be observed in these experiments since the deictic mechanisms does not permit any ambiguity of referent, but we would expect it with less tight grounding (cf. (Steels 1998*a*) where ranges of sensor values are used instead of deixis). Similar phenomena of convergence and lexical drift and phonetic change as seen in the above tables are also well-known from historical linguistics of natural language (see, for example, (Jeffers & Lehiste 1979) for a good introduction).

There are a number of significant differences between these experiments and those described in (Steels 1996). Steels' interactions are synchronous and ordered; his attentional mechanisms are built-in and never fail; and successful or failed communicative acts are reinforced or yield error feedback (although the apparently external criteria of success seem not to be made explicit in his publications). Here in our work, we have asynchronous non-deterministic scheduling creating potential delays in response (possibly even failure to respond) from listening agents. There are no positive or negative feedback signals. The asynchronous scheduling of the agent actions and perceptions is very much reminiscent of phenomena related to attention and distraction. If an agent is attending to the speech of an interlocutor, or only attending to part of it in the temporal course of articulation, then what is learned from that speech is highly dependent on the portion that is actually attended to.

Even without an explicit feedback mechanism reinforcing any 'successful' communitative acts many phenomena of community linguistic ontogeny are apparent.

[2] In the table † indicates the mother agent and ‡ the other parent. A '-' indicates no response.

It is an interesting research question to characterize exactly how much reinforcement is really necessary for achieving various aspects of linguistic phenomena. Certainly notions of meaningful information connecting language to utility and value for individual agents will be necessary to achieve more sophisticated linguistic systems (cf. (Nehaniv 1999)).

The connectionist architecture used has properties very close to those exhibited by mirror neurons, which may have played a role in the evolution of human language readiness (Rizzolatti & Arbib 1998, Arbib in press). Mirror neurons in the pre-motor cortex of monkeys fire when an act or action affordance is perceived as well as when it is used. This happens whether the actor is the animal itself or another who is being observed by the animal (see references cited for more details). Similarly in this implementation of the DRAMA architecture, agents make no distinction between a name that has been heard and one that was generated by the agent itself; in both cases the same neurons fire. Also, in observing or generating deictic acts the same neurons fire. An agent may set its own deictic or phonetic nodes, and their activation is perceived by its sensor and so the sensor perceives and acts in the same manner regardless of who generated an event, both in pointing and in speaking.

In the experiments described here there are always two child agents which are not heared by others in the dialogue (though they behave as though they are, both in pointing and in generating names). The children listen to the parents and learn their language but do not contribute to its evolution in the community, since the parents do not listen to their children but only to each other (and not always then). However, a child may nevertheless autonomously generate its own names and use them; this may possibly conflict with the child's learning of names used by the parent, e.g. resulting in a name that is the combination of a word used by the parent and a word it generated itself if both are perceived in the interval following deixis toward the referent of these two words (synonomy leading to a new name). Similar situations can occur if the speech of the parents overlaps and thus influences their names for given referents (including possibly different ones).

This work illustrates the dynamics and development of deixis-grounded naming systems in (asynchronous) interaction and thus extends the work of Steels (Steels 1996, Steels 1998*a*) on the emergence of vocabularies in communities of agents – with some attentional phenomena and observations as just discussed. By adding the mechanism of name generation, this work also extends the work of Billard, Hayes, and Dautenhahn (Billard & Dautenhahn 1999, Billard & Hayes 1999) on the grounding and learning of vocabularies in robots and agents using connectionist methods, albeit in a simple simulation. While their work used an explicit pre-existing proto-language, ours does not but illustrates how such simple proto-language can arise with very simple agents in very simple interactions. Their work, like ours, also did not employ any reinforcement mechanisms.

The results reported here illustrate the use of deixis to ground the genesis and growth of developing vocabularies, and the mechanisms could be carried over to more complex agent communities and to robotic platforms without difficulty.

References

Arbib, M. A. (in press), The mirror system, imiation, and the evolution of language, *in* K. Dautenhahn & C. Nehaniv, eds, 'Imitation in Animals and Artifacts', MIT Press.

Bickerton, D. (1990), *Language and Species*, Chicago.

Billard, A. & Dautenhahn, K. (1999), 'Experiments in learning by imitation - grounding and use of communication in robotic agents', *Adaptive Behavior*.

Billard, A. & Hayes, G. (1999), 'DRAMA, a Connectionist Architecture for Control and Learning in Autonomous Robots', *Adaptive Behavior* **7**(1), 35–63.

Butterworth, G. (1991), The Ontogeny and Phylogeny of Joint Visual Attention, *in* A. Whiten, ed., 'Natural Theories of Mind', Blackwell.

Chomsky, N. (1957), *Syntactic Structures*, Mouton Publishers.

Chomsky, N. (1968), *Language and Mind*, Harcourt, Brace and World.

Chomsky, N. (1975), *Reflections on Language*, Pantheon Books.

Chomsky, N. (1981), *Lectures on Government and Binding*, Foris Publications.

Dautenhahn, K. (1994), Trying to Imitate – A Step Towards Releasing Robots from Social Isolation, *in* 'Proc. From Perception to Action Conference (Lausanne, Switzerland, Sept 7-9, 1994)', IEEE Computer Society Press, pp. 290–301.

Dautenhahn, K. (1997), 'I could be you — the phenomenological dimension of social understanding', *Cybernetics and Systems* **25**(8), 418–453.

Griffin, D. R. (1976), *The Question of Animal Awaresness: Evolutionary Continuity of Mental Experience*, The Rockefeller University Press.

Griffin, D. R. (1992), *Animal Minds*, Chicago.

Herman, L. M. & Austad, S. N. (1996), Knowledge acquisition and asymmetry between language comprehension and production: Dolphins and apes as general models for animals, *in* M. Bekoff & D. Jamieson, eds, 'Readings in Animal Cognition', MIT Press, pp. 289–306.

Hurford, J., Knight, C. & Studdert, M., eds (1998), *Approaches to the Eovlution of Language: Social and Cognitive Bases*, Cambridge University Press.

Jeffers, R. J. & Lehiste, I. (1979), *Principles and Methods for Historical Linguistics*, MIT Press.

Moore, C. & Dunham, P. J., eds (1995), *Joint Attention: Its Origins and Role in Development*, Erlbaum.

Nehaniv, C. L. (1999), Meaning for Observers and Agents, *in* 'IEEE International Symposium on Intelligent Control/Intelligent Systems and Semiotics', ISIC/ISAS '99, pp. 435–440.

Nehaniv, C. L. (2000), The Making of Meaning in Societies: Semiotic and Information-Theoretic Background to the Evolution of Communication, *in* 'AISB Symposium: Starting from Society: the application of social analogies to computational systems', Association for Artificial Intelligence and the Simulation of Behaviour.

Pepperberg, I. M. (in press), Allospecific referential speech acquisition in grey parrots (*Psittacus erithacus*): Evidence for multiple levels of avian vocal imitation, *in* K. Dautenhahn & C. Nehaniv, eds, 'Imitation in Animals and Artifacts', MIT Press.

Pinker, S. & Bloom, P. (1990), 'Natural language and natural selection', *Brain and Behavioural Sciences* **13**, 707–784.

Rizzolatti, G. & Arbib, M. A. (1998), 'Language within our grasp', *Trends in Neurosciences* **21**(5), 188–194.

Savage-Rumbaugh, S. & Brakke, K. E. (1996), Animal Language: Methodological and interpretative issues, *in* M. Bekoff & D. Jamieson, eds, 'Readings in Animal Cognition', MIT Press, pp. 269–288.

Scaife, M. & Bruner, J. (1975), 'The capacty fo joint visual attention in the infant', *Nature* **253**, 265–266.

Scassellati, B. (1999), Imitation and Mechanisms of Joint Attention: A Developmental Structure for Building Social Skills on a Humanoid Robot, *in* C. L. Nehaniv, ed., 'Computation for Metaphors, Analogy, and Agents', Vol. 1562, Springer Verlag, pp. 176–195.

Steels, L. (1996), 'A self-organizing spatial vocabulary', *Artificial Life* **2**(4), 315–332.

Steels, L. (1997), 'The Synthetic Modelling of Language Origins', *Evolution of Communication* **1**(1).

Steels, L. (1998*a*), 'The Origins of Syntax in Visually Grounded Robotic Agents', *Artificial Intelligence* **103**, 1–24.

Steels, L. (1998*b*), Synthesising the origins of language and meaning using co-evolution, self-organization and level formation, *in* J. Hurford & M. Studdert-Kennedy, eds, 'Approaches to the Evolution of Language', Cambridge University Press.