In Proceedings of IASTED Int. Conf. on Modelling, Simulation and Optimization, Gold Coast, Australia, 1996

A Neural Network Simulation of Child Language Development at the One-word Stage

Syed Sibte Raza Abidi
School of Computer Sciences
Universiti Sains Malaysia
11800 Penang, Malaysia.
Tel: +60-4-6577888 ext. 3618
Fax: +60-4-6573335
Email: sraza@cs.usm.my

Keywords: *Neural networks, Language development, Learning (supervised and unsupervised), Modular architecture*

Abstract
Neural networks, or the so-called connectionist networks, provide a basis for studying child language development in that such networks emphasise *learning,* either from observations or from being told. We report a neural network based simulation of aspects of child language development, essentially focusing on child language at the one-word stage. We simulate the learning of associations between the so-called conceptual relations, perceptual entities and functional words, leading to the production of child-like one-word utterances. The simulations are carried out using ACCLAIM - A Connectionist Child LAnguage development and Imitation Model. ACCLAIM is a 'hybrid' neural network architecture, systematically synthesising both supervised and unsupervised learning networks, thus accounting for the diverse nature of inputs to and outputs from a child learning language. The simulations carried out are 'language informed' as realistic child language data has been used for training the neural networks.

1. Introduction

The emergence of *Neural Networks*, a class of *parallel distributed processing* architectures, aims to understand the nature of human intelligence by simulating aspects of human behaviour through a collection of idealised 'neurons'. *Neural networks*, or the so-called *connectionist networks* provide substantial opportunities for simulating various human learning activities in that neural networks have a natural propensity for learning - either supervised (based on instructions) or unsupervised (based on observations), and that the design of these networks simplifies questions related to the representation of linguistic and 'world' knowledge through the use of a network of 'simple' nodes and links.

Child language development is an exemplar of human learning, a development that has an evolving theoretical framework which can do with the objectivity that may be implicit in the methodology of neural network based simulations of human learning. In this regard, neural networks offers mechanisms, such as adaptive learning, generalisation, self-organisation, feature extraction, and pattern-recognition that appear to have a direct relevance towards the simulation of child language development [1,2,3,4].

In this paper we demonstrate how neural networks can be used to simulate, or more accurately trained, to mimic the development of language amongst children during infancy (9-24 month age group). In particular, we would address aspects of the child's language at the one-word stage, which in terms of communicative competence accounts for the child's earliest sentences - the so-called *one-word utterances*. We report a simulation of the learning of a mapping between a set of 'conceptual relations' and 'functional words', which leads to the production of child-like one-word utterances. Our simulation is *language informed* such that it uses realistic child language data taken from the archives collocated by Lois Bloom [5].

2. ACCLAIM-A 'Modular' Neural Network Architecture

We have developed ACCLAIM (A Connectionist Child LAnguage development and Imitation Model) - a 'modular' neural network architecture to simulate aspects of child language development within the age group 9-24 months [6].

Architecturally, ACCLAIM aims to achieve a degree of psychological plausibility: language development aspects which can be construed to be innate development have been simulated by using unsupervised learning regimes (like Kohonen maps and Hebbian connections), whilst environmentally-determined aspects of language development have been simulated by using supervised learning regimes (like backpropagation networks).

ACCLAIM incorporates a 'modular' architecture, whereby individual neural networks are configured in a meaningful manner to realise a 'neural network module'. Within a neural network module the individual neural networks retain their identity and merely interact with each other to provide a more powerful and elaborate response. Each neural network module is then dedicated to simulate a particular aspect of child language development. Using the individual neural networks (shown in figure 1b) as building blocks, we have developed four different modules, where each module can simulate a different aspect of child language development (figure 1a). Extending the modularity approach further, the various modules are then synthesised, to realise the unified architecture of a 'modular' language development model - ACCLAIM (figure 1b).

ACCLAIM has been used to simulate the development of concepts amongst children together with the ostensive naming of these concepts [7]: the *concept memory* and *word lexicon* have been simulated using Kohonen maps and are linked together through a Hebbian connection based *naming connection network.* Backpropagation networks have been used to implement a *conceptual relation network* (for one-word utterances) and a *word-order testing network* (for two-word sentences). Children's evolving semantic performance has been simulated by a *semantic relation network* using a Hebbian connection network. ACCLAIM has been trained on 'realistic' child language data and has learnt to produce child-like one-word and two-word sentences.

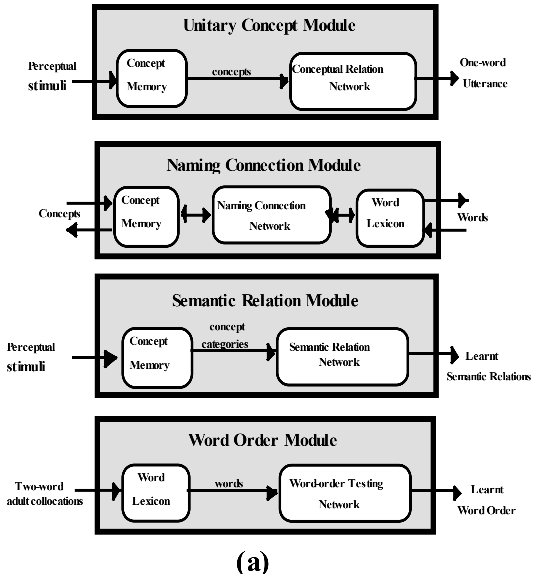

Figure 1: (a) Four neural network modules each simulating some aspect of child language development; (b) The modular architecture of ACCLAIM

3. An Account of Children’s One-Word Language

What are children's first words? and indeed what do they mean? The search for an answer to this question has resulted in an ongoing investigation spanning several years with contributions from a large number of researchers. Lois Bloom’s [5] characterisation of children's' early words is of relevance here as it is developmental in spirit. Bloom argues that context and behaviour play an important role in the child's speech; the child makes reference to objects which can be seen, heard, felt or otherwise perceived. Children therefore seem to acquire words that encode their existing concepts about the environment or to encode the concepts that are still developing.

Bloom divides children's first words into two categories or 'word-forms' - 'substantive words' and 'functional words'. Substantive word-forms are object specific and they make reference to classes of objects and events based on their perceptual features. Functional word-forms make reference across these perceptually distinct classes of objects and events based on relational notions like 'disappearance', 'recurrence', 'upness', 'cessation' and some others. Typical functional words noted by Bloom are ‘more', 'there', 'up', 'uh oh', 'away', 'gone', 'no' and 'stop'. Bloom argues that sensori-motor development aided with perceptual experiences, then, enables the child to represent the regularities in his/her environment in terms of functional words that encode certain 'conceptual relations -- the underlying cognitive structures (concepts or thoughts) that represent the relations between persons, objects and events in the world. (the so-called ‘perceptual entities’). For instance, objects are acted upon, can exist, cease to exist and recur; people do things and they appear and disappear and so on. Bloom suggests that these are the kinds of conceptual relations children talk about at the one-word stage. For instance, after finishing a glass of milk, the child might ask for more milk by merely saying the functional word *more*.
From Bloom's discussions we have collected 25 conceptual relations, 3 perceptual entities and 14 functional words. This information is now systematically organised in table 1.

Conceptual Relation	Perceptual Entity (ies)	Functional Word
Disappearance	object, people	*gone, away*
Request, Recurrence	object, event	*more*
Existence	object, people	*this, there*
Existence	event	*uh oh*
Non-occurrence, Failure, Rejection	event	*no*
Non-existence	object, people	*no*
Greeting	people	*Person's name (Mama, Dada, etc.)*
Cessation	event	*stop*
Pointing (to draw attention)	object, people	*there*
Actor	object, event	*Person's name*
State (large size)	object	*big*
State (small size)	object	*small*
State	people	*dirty*
Upness (Location)	object	*up*
Downness (Location)	object	*down*
Substantive	object	*Object's name*
Substantive	people	*Person's name*

Table 1: Conceptual relations with their corresponding functional words. Also shown are the perceptual entities related with a conceptual relation

4. Simulation of language development at the one-word stage

From the above discussion it appears that at the one-word stage the child expresses his/her 'communicative intention' regarding a perceptual entity by uttering a functional word which encodes a conceptual relation. One can then argue that the task of learning conceptual relations, which leads to the production of one-word utterances, actually involves the learning of two associations: (a) an association between conceptual relations and corresponding 'functional' words. The learning of such an association ensures that whenever a conceptual relation is intended to be expressed the corresponding word is produced; and (b) an association between conceptual relations and perceptual entities, which determines the relationship between conceptual relations and perceptual entities.

From a neural network perspective it appears that the type of learning involved in the simulation of children's one-word language is *associative* in nature. Guided by the above learning constraints, for our simulation we use a Backpropagation (BP) network [8], which loosely speaking allows the establishment of arbitrary, non-linear relationships or mapping between input and output patterns.

4.1. Specification of the *Conceptual Relations Network*

Within ACCLAIM, the conceptual relation (CR) network is used to simulate aspects of one-word language development. The input layer of the CR network comprises 25 units, and encodes in a localist manner 22 conceptual relations and 3 perceptual entities. The output layer of the CR network comprises 18 units and again in a localist manner encodes 15 typical children's one-word utterances (words) and 3 perceptual entities. The hidden layer comprises 5 units. Table 2 gives the organisation of the input and output layers of the CR network.

INPUT LAYER		INPUT LAYER		OUTPUT LAYER	
Unit	**Conceptual Relations**	**Unit**	**Conceptual Relations**	**Unit**	**One Word Utterances**
1	Disappearance (O, P)	20	Downness (O)	1	GONE
2	Disappearance (O, P)	21	Substantive (O)	2	AWAY
3	Request (O, E)	22	Substantive (P)	3	MORE
4	Recurrence (O, E)	**Perceptual Entities**		4	THIS
5	Existence (O, P)	23	OBJECT (O)	5	THERE
6	Existence (O, P)	24	PEOPLE (P)	6	UH OH
7	Existence (E)	25	EVENT (E)	7	NO
8	Non-occurrence (E)			8	PERSON'S NAME
9	Failure (E)			9	STOP
10	Rejection (E)			10	BIG
11	Non-existence (O, P)			11	SMALL
12	Greeting (P)			12	DIRTY
13	Cessation (E)			13	UP
14	Pointing (O, P)			14	DOWN
15	Actor (O, E)			15	OBJECT'S NAME
16	State-Large size (O)			**Perceptual Entities**	
17	State-Small size (O)			16	OBJECT (O)
18	State (P)			17	PEOPLE (P)
19	Upness (O)			18	EVENT (E)

Table 2: Organisation of the input and output layers of the CR network

4.2. Nature of the Training Stimuli

The training stimuli used for this simulation depicts a situation in the child's experience which represents the child's communicative ‘intention’ in terms of a conceptual relation and a perceptual entity. The output response is the linguistic manifestation of the underlying intention, i.e., a one-word utterance. Consider a situation (see table 3) when the child is reported requesting for some juice. Here, one can assume that the underlying communicative ‘intention’ is the conceptual relation *recurrence* of a perceptual entity *objects* (*juice)*. The child’s output is the functional word *'more'*, expressing his/her request for more juice.

	Input Stimuli		**Verbal Output**
Situation	**Conceptual Relation**	**Perceived Entity**	**Allison's one-word utterance**
(Mother pours herself juice) (Allison picking up empty cup) (Allison putting her cup aside)	recurrence	object (juice)	*more*

Table 3: An exemplar situation involving the child - Allison (taken from Bloom's data).

Bloom's [5] archives contain numerous such situations together with the child's response. We have prepared 32 training patterns from Bloom's work (the ad hoc list given in table 1).

4.3. The Learning Sequence

Learning involves the gradual development of an association between the conceptual relation and the corresponding desired one-word utterance. In a BP network, learning is an error minimisation process which spans over number of iterations. In each iteration, a training pattern is randomly selected from the ensemble of training patterns to be learnt. Learning ceases when the error between the desired output and the produced output is below a specified *error threshold*. For our simulation 5110 iterations were required to learn the training stimuli.

5. The Simulation Results

5.1. Associations between *conceptual relations* and *words*

Table 4 shows the performance of the learnt CR network to produce the learnt experimental data. The test input patterns comprise a conceptual relation and a perceptual entity. The most highly active unit in the output layer is regarded as the one-word utterance produced in response to the input pattern. We also consider other output units that have high activation levels as this provides us an insight to the global knowledge learnt by the CR network. Table 4 indicates that the CR network has successfully learnt to map the input pattern (conceptual relation and perceptual entity) to the relevant function word according to the relationships given in table 1.

Input Pattern	Output Produced	Activation Level of RU	Other Highly Active Units
Request + Object	***More***	**0.949**	Big (0.074), Object name (0.027), Away (0.024), Up (0.021)
Greeting + People	***Person name***	**0.953**	Dirty (0.056), Away (0.027), Stop (0.020), There (0.015)
Cessation + Event	***Stop***	**0.851**	More (0.071), Dirty (0.034), Uh-oh (0.026), Small (0.022)

Table 4: Result of the simulation of the learning of a mapping between conceptual relations and functional words

5.2. Explicating inherent associations between *conceptual relations* and *perceptual entities*

The training patterns encoded an implicit relationship between conceptual relations and perceptual entities, for instance the *owner* of an object belongs to the perceptual entity *people*.

Input Pattern	Output Produced	Activation Level of RU	Other Highly Active Units
Existence	***Event***	**0.573**	People (0.105), Object (0.014)
Owner	***People***	**0.293**	Event (0.140), Object (0.021)
State	***People***	**0.968**	Event (0.266), Object (0.001)
Recurrence	***Event***	**0.303**	Object (0.252), People (0.005)

Table 5: Results showing the learning of an association between conceptual relations and perceptual entities

In table 5 we show that the CR network has indirectly learnt the possible relationships between conceptual relations and perceptual entities, i.e., what conceptual relations can co-occur with which perceptual entity(ies). This information is made explicit by presenting as input just a conceptual relation and noting the perceptual entity activated in response.

5.3. Generalisation affects

Child theorists have argued that children have the ability to generalise in novel situations, i.e., to produce an appropriate response based on previously learnt knowledge. In table 7 we demonstrate the generalisation capabilities of The CR network when presented with novel input situations. The results of these simulations would indicate that the CR network can handle novel situations by generalising to produce an appropriate response based on prior learnt knowledge.

Input Pattern	Output Produced	Activation Level of RU	Other Highly Active Units
Disappearance + Event	***Gone***	**0.710**	No (0.234), Stop (0.061)
Existence + Event	***This***	**0.344**	No (0.023), Person name (0.021)
Recurrence + People	***More***	**0.404**	Away (0.206), Gone (0.029)

Table 7: Generalisation capabilities of the CR network in novel situations

6. Simulation of the production of child-like *one-word utterances*

Simulation of the production of one-word utterances is achieved by presenting to the CR network some real-life situations from Bloom's child language data, and noting the one-word utterance produced by the CR network. Table 8, illustrates some real-life situation taken from Bloom [5] and the one-word utterances produced by the child - *Allison* and the CR network.

Real-Life Situation	Child's Response	CR network's input pattern	CR network's response
(M pointing to chair) What is this?	***chair***	Pointing + Object	***obj name (chair)***
(M pours self juice) (A picking up empty cup) (A putting her cup aside)	***more***	Request + Object	***more***
(M pours juice; A drinks juice, looks into empty cup. M taking cup) Where's the juice?	***gone***	Disappearance + Object	***gone***
(A pushing car) (car stops) (A walks toward car; to urge to move) (A pushes car; the car stops)	***hmmm*** ***stop*** ***more*** ***stop***	Cessation + Event Request + Event Cessation + Event	***stop*** ***more*** ***stop***
(A pointing up at the microphone) What, darling? (A sliding off chair) You wanna get down? (A gets down from chair) Down	***man*** ***Mommy*** ***down*** ***down/ there***	Pointing + Object Actor + Object Downness (location) + Person Downness + Person /Pointing + Event	***there*** ***Person name (Mum)*** ***down*** ***down/ there***

Table 8: Results of the simulation of the production of one-word utterances. The letters A and M in the situations refer to the child - Allison and her Mother, respectively.

A comparison of one-word utterances produced by the CR network and the child - Allison provides a qualitative measure of the efficacy of the CR network. It may be noted that in each situation the CR network produces a response similar to the child's (Allison) one-word utterance. The generalisation capabilities of the conceptual relation network are demonstrated when producing the correct response for the last situation: the input patterns *downness* + *people* and *pointing* + *event* are both novel for the CR network. However, in both these cases the CR network generalised to produce correct one-word utterances. We would like to point out that, the CR network is not restricted to just the above situations, rather it can produce appropriate one-word utterances for many other situations.

7. Conclusions

The fact that a number of researchers are interested in neural network based simulations of human learning behaviour, leads us to believe that the simulation carried out by ACCLAIM has made a contribution to this interesting area of research. By using 'realistic' child language data we have demonstrated how neural networks can be used to both operationalise extant child language corpora and systematise psycholinguistic theories [1].

We have shown that in a neural network simulation context, development of language at the one-word stage is guided by a single, homogenous learning mechanism in which regular real-life situation and their exceptions can be made to coexist successfully by exploiting the fact that the CR network is able to simultaneously learn individual mappings between conceptual relations and function words and also abstract general regularities in the input situations.

References

[1] Bechtel, W. & Abrahamsen, A. *Connectionism and the Mind*. (Oxford: Basil Blackwell, 1991).

[2] McClelland, J. *PDP: Implications for* Cognition and Development. R. Morris (Ed.) Parallel Distributed Processing: Implications for Psychology and Neurobiology, (Oxford: Clarendon Press., 1989)

[3] Levine, D. *Introduction to Neural and Cognitive Modelling*. (Hillsdale: Lawrence Earlbaum Associates, 1991).

[4] Seidenberg, M. "*Connectionist Models and Cognitive Theory",* Psychological Science, 4, 1993, 228-235

[5] Bloom, L. *One Word At A Time*. (Paris: Mouton, 1973).

[6] Abidi, S.S.R. *A Connectionist Simulation: Towards a Model of Child Language Development*. Doctoral Thesis. (Guildford: University of Surrey, England, 1994).

[7] Abidi, S.S.R. & Ahmad, K. *Child Language Development: A Connectionist Simulation of the Evolving Concept Memory*. M. Aldridge (Ed.) *Child language.* (Clevedon: Multilingual Matters Ltd. 1996)

[8] Rumelhart, D., McClelland, J. & The PDP Research Group *Parallel Distributed Processing, Vol. I & II.* (Cambridge: MIT Press, 1986).