How Children Learn to Use Language

An Overview of R Narasimhan's Ideas on Child Language Acquisition

Raman Chandrasekar

Raman Chandrasekar has been with Microsoft Research, Redmond, USA. since 1998. His current research interests are at the intersection of information retrieval, natural language processing and machine learning. He received his PhD from TIFR, Bombay.

There have been many theories about how children learn to use language. Professor Narasimhan proposed a theory of child language acquisition based on behavioural pragmatics. In this article we present a simplified version of his theories about how children learn to communicate, to describe, manipulate and explore the world around them from exposure to variety of language utterances and non-verbal inputs such as gestures and pointing. We also discuss the method he used to substantiate his ideas, and briefly present a computational model of the ideas arising from his work.

1. Introduction

What is the process by which children learn to use language? This process takes place all around us, and it happens all the time. It is so normal, you probably don't think about it too much. But yet it is so amazing, so incredible when you actually do stop to think about it. Most children, regardless of their country of origin, race or other factors acquire the language to which they are exposed; they learn by interacting with their parents or primary care givers without any explicit external tutoring. How does this miracle take place?

This is a question that many people have tried to answer. Usually their answers have centered on syntax (the rules governing the ordering of words and other elements in language) and semantics (which is about how we get meaning from language structure). Others have offered psychological explanations based on stimulus-response theories. There have been theories by Noam Chomsky and others about how children are born all wired up to use language, but their ideas do not explain all the data we have on child language acquisition.

Professor R Narasimhan's thesis is that children do not acquire their first language as a set of rules about how to put words in a language together. Instead, he says, they acquire "language behaviour" – they learn how to use language to describe, manipulate or explore the world around them – to communicate with others and express their intentions, to describe things and to make others perform actions. Children learn all this from being exposed to a variety of language utterances and non-verbal inputs such as gestures and pointing. Thus his ideas are based on (behavioral) pragmatics – based on language use.

But how do children learn from these inputs? Narasimhan has tried to explain what we have observed, and has collected additional data to prove his ideas and theories.

In the rest of this article, we will present a simplified, relatively jargon-free view of Narasimhan's ideas of how children acquire language. We will talk about some of his ideas, talk about the scientific method he used to substantiate his ideas, and briefly talk about a computational model of the ideas arising from his work. There is a lot we cannot cover in this article. For those of you interested in learning more about this fascinating field, there is a list of books and articles at the end of this article.

2. Features of Child Language Acquisition

In his book, *Language Behaviour*: *Acquisition and Evolutionary History* ([1], the primary reference I have used for this article), Professor Narasimhan notes down the main features of child language acquisition (pages 41–42):

1. Children acquire language behaviour by living and growing up in a language community and by interacting intimately with the members of that community.
2. Language behaviour is acquired in a relatively short time without any *systematic* tuition.
3. The acquired language behaviour exhibits several well-defined traits of the language community within which it is acquired.

R Narasimhan's thesis is that children do not acquire their first language as a set of rules about how to put words in a language together.

Children acquire language behaviour by living and growing up in a language community and by interacting intimately with the members of that community.

4. The acquisition process goes through discernible developmental stages.

5. At each developmental stage, language behaviour exhibits two fundamental traits: (a) it is productive and (b) it has a generalizing tendency.

6. A new-born child acquires the language behaviour of the community it lives with and grows in, independent of the language background of its parents.

Note especially that children need inputs from a language community around them to learn language[1]. While children do not need regular classes to learn language, they do need exposure to language utterances in "appropriate behavioural contexts".

Narasimhan postulates that language behavior is example driven; he suggests that imitation, rehearsal and analogizing all play critical roles in language acquisition. Children hear the sounds of language and repeat these sounds to learn them. They are also capable of tirelessly practicing verbal behaviour, starting from babbling – rehearsing what they have learnt – in order to acquire such behaviour. Finally, once they have learned some aspects of language, they create new forms by analogy. This could be at the level of sounds, or grammatical modifications to indicate tense, number, etc. This is so well known that there is a standard notion of 'over-generalization', for example when children mistakenly apply –'ed' as the standard past tense marker to an irregular verb. For example the verb to *see*, would be said as *see-ed,* as in the sentence "I see-ed Vanitha yesterday"!

While earlier the child imitates words, phrases or sentences, later the child uses these as templates or schemas to generate new sentences, in the form of descriptions, commands or questions. This increasing language ability helps the child interact and learn even more.

How do children acquire language behavior, or the ability to interpret and generate language? To understand this, consider some field work done on language acquisition by Narasimhan and his associate Raghunathan Vaidyanathan.

[1] This idea is explored in two films from the '70s. Francois Truffat's 1970 film "L'Enfant Sauvage" (or "The Wild Child"), based on Jean Itard's *Rapport fait à son excellence le Ministre de l'intérieur sur les nouveaux développemens et état actuel du sauvage de l'Aveyron (*a report to the (French) Minister of the Interior about new developments and the current state of the wild child of Aveyron), portrays a French child who was found in a forest in 1798, who could not walk, talk, read or write. Similarly, Werner Herzog's 1974 film "Jeder für sich und Gott gegen alle" (*Every Man for Himself and God against All*, also known as *The Enigma of Kaspar Hauser*) is based upon the true and intriguing story of Kaspar Hauser. Kaspar appears in Nurenburg in 1828, barely able to speak or walk; it turns out later he was locked up in a dungeon since birth, without any language inputs.

3. Collecting Data to Test these Theories

In the early 1980s, Narasimhan and Vaidyanathan collected data on how parents interact with children [2]. They examined the very early stages of language acquisition of a girl named Vanitha. Vanitha was then the only child of a family in Bombay (now Mumbai) whose mother tongue was Tamil. Approximately every two weeks (with some gaps when Vanitha was away), Vaidyanathan visited the family on the same day of the week, at the same time. He collected data from the time Vanitha was nine months old to the time when she was 33 months old. During each visit, about 45–60 minutes of audio recordings were made inside the house while Vanitha's parents were looking after and playing with the child. The recordings were supplemented with written notes containing relevant contextual information about the non-verbal environment, the agents and objects referred to during interaction.

All this speech data was transcribed. The Tamil segments were translated into English and merged with the notes. The 24 sessions of parent-child interaction fills 135 word-processed pages, and each session has been divided into logical episodes. This data – called the Vanitha corpus – was contributed to CHILDES [3], an international database of child language acquisition data, and is till today the only available data in this area covering an Indian language! Others besides Narasimhan and Vaidyanathan have used this data to explore aspects of child language acquisition.

As you go through this corpus, you can see how Vanitha's language behaviour slowly becomes a more dominant part of her interaction with others. In early episodes Vanitha draws attention to items by pointing to them. Later she starts naming objects. Also, initially she answers questions with actions, and later answers verbally. For example consider the following episodes between Vanitha and her father.

[Note: In these examples, compared to what is available in the original, I have used a slightly simpler way of transcribing the Tamil segments.]

Narasimhan and Vaidyanathan examined the very early stages of language acquisition of a girl named Vanitha.

In early episodes Vanitha draws attention to items by pointing to them. Later she starts naming objects. Also, initially she answers questions with actions, and later answers verbally.

As Vanitha acquires more language, she starts to express desires, wants and intentions. She also learns to find out more about the world around her.

Episode 13.5: Answering with Actions/Miming

Vanitha is 18 months old. Her father shows her an illustration of a hair-oil bottle.

Father:	*idu enna*	What's this?
	ennai	Oil
	enna pannuva ennaye?	What'll you do with oil?

Vanitha rubs her head with her hand.

Episode 20.23: Answering with words

Vanitha is now 25 months old. She points to an illustration in a newspaper.

Vanitha:	*Saampu paattiyaa?*	(Did you) see shampoo?
Father:	'Shampoo' *enna pannuva?*	What'll you do with shampoo?
Vanitha:	*tala tappipaa*	Rub on the head.

As Vanitha acquires more language, she starts to express desires, wants and intentions. She also learns to find out more about the world around her. These ideas are shown in the interactions below:

Episode 18.18: Finding out about the state of the world

Vanitha is now 21 months old. She goes to her Mother.

| Vanitha: | *tuuttam taa* | Give (me) water (to drink). |
| | *amma* 'brush' *enge?* | Where's mother's brush? |

Episode 19.1: Learning about objects

Vanitha is now 25 months old. She looks at the recording microphone (that the observer is using).

Vanitha:	*ennadu?*	What?
	ennadu idu?	What's this?
Observer:	'mike'	
Vanitha:	*im*	[not a word]
Observer:	'mike'	
Vanitha:	'radio'	

Note here that Vanitha calls the microphone a radio, perhaps because she has never seen a hand-held compact recorder with built-in microphone that resembles a radio receiver.

Narasimhan notes that each episode falls into one out of a few categories such as:

- Eating/feeding
- Bathing/washing
- Dressing
- Other caretaking
- Play
- Reading/writing
- Story telling/listening

He also points out that within each episode, the language used and the semantics of that language are quite specific and that the child's perception of the environment is synchronized with the corresponding language expressions that denote it. All this makes it easy for the child to learn from these interactions.

The parents' (or caregivers') utterances to the child are typically simple and well-formed. They tend to repeat whole chunks of what they say, but repeated utterances are not always identical copies. Sometimes, they are variants with the same meaning, or referring to a similar context; sometimes the words or phrases are ordered differently. There is also a sing-song cadence to their utterances, and often a special way in which they are said. Some researchers call this speech of parents and caregivers child-directed speech. The term 'motherese' has also been applied to this.

Like other proponents of child-directed speech/motherese, Narasimhan argues that the systematic use of such patterns in the utterances is very important for language acquisition. When you have speech chunks of different sizes, repeated utterances help the child segment longer chunks. Substitutions help the child recognize which parts of the utterances are interchangeable, which parts mean roughly the same things, etc.

The language used and the semantics of that language are quite specific and that the child's perception of the environment is synchronized with the corresponding language expressions that denote it.

When you have speech chunks of different sizes, repeated utterances help the child segment longer chunks.

Parent helps their children learn initially by pointing and using the names of things, then by using onomatopoeic or baby-talk versions of words (e.g., *mum-mum* to indicate food or eating; *bow-wow* or *baw-waw* for a dog) to relate to the situations or items on hand, and later connecting these baby-talk forms to actual names of the situations or items. Parents also explicitly correct mistakes or wrong use of names or labels. However, corrections are provided only when the child has attained a certain cognitive stage. Prior to that, parents just deny the child's mistakes and that too only indirectly. For instance, when a child misnames a cat as a dog, the parents might deny this, asking the child "Adu naya?" ("Is that a dog?"), instead of "Adu nay illa" ("That is not a dog").

In other studies, it has been noted that parents, when asking questions, typically provide the answers as well, as in "What is the colour of this pencil? Is it red or black?" This shows the child what questions go with what answers, and analogously, what kinds of questions go with what kind of answers.

Thus parents constantly provide additional information that helps children digest and assimilate what they say.

Consider the following interaction to see how children could acquire language behaviour tokens like "your head" and language behaviour schema for sentences like "Shake your head".

Episode 3.8: Learning the 'shake' schema

Vanitha is 11 months old.

Mother: *aattu*	Shake.
mandaya aattu	Shake (your) head.
innum konjam mandaya aattu	Shake (your) head some more.

Assume that Vanitha has learnt what *shaking* is about, because of explicit actions shown by her parents. The second statement above shows that you can apply the action of 'shaking' to an object 'head'. This can be represented as

Action: 'Shake'
Object: Can be 'head'

The third statement shows you can add a qualifier ('how much') you can do the shaking – in this case 'some more'!

Action: 'Shake'
Object: Can be 'head'
Qualifier: some more [optional]

You can imagine that the parent can introduce negation, by using a sentence like "No, don't shake your head."

Action: 'Shake'
Object: Can be 'head'
Qualifier: some more [optional]
Negation: [optional]

Now if the parent says something different, like "Now shake your leg", the child's schema for 'shake' could change to something like:

Action: 'Shake'
Object: Can be 'head' or 'leg'
Qualifier: some more [optional]
Negation: [optional]

Thus the child learns new utterances and schemas by adapting schemas from one usage in one episode to other analogous episodes, and by replacing variables or language behaviour tokens with similar other tokens. However, there may be conditions about which kinds of language behaviour tokens are interchangeable. So, for instance, you can substitute *leg* or *hand* for *head*, and still make sense. But you cannot say "Shake yellow".

As you add more language behaviour tokens and schemas, we can see how the child can learn to understand and say something like "Bring a newspaper" or "Bring a newspaper to read". Extrapolating further, we can add some *query-tokens* to be able to ask who, why, what, etc., kinds of questions. So now you can imagine the child (now) learning to say "Why should I shake my leg?"

The child learns new utterances and schemas by adapting schemas from one usage in one episode to other analogous episodes, and by replacing variables or language behaviour tokens with similar other tokens.

Some readers may recognize that the representation shares similarities with case grammar [4].

Narasimhan and Vaidyanathan have explored several examples in the Vanitha corpus to demonstrate specific language behaviour learning. Vaidyanathan has also developed models of how children learn interrogatives and negation [5,6].

There are many other questions we have not answered about how children acquire the schemas and tokens and how they interpret them, and how they connect non-verbal inputs with such schemas and tokens. The interested reader will find more details in the references at the end of this article.

4. A Computational Model of Language Acquisition

Can the language acquisition model described here be computationally viable, in realistic situations? Can we write a program that will take the utterances given to children, and make the computer understand, learn and behave just like we say children would, given the same input?

V Sembugamoorthy, another student of Narasimhan's, implemented a system named PLAS (short for Paradigmatic Language Acquisition System) and showed that PLAS could interpret situational aspects described to it [7]. For example, you could tell the system that "There is a red book on the table", and it would understand the situation and the relationships between entities, building up a model similar to what we would expect a child to build up if Narasimhan's ideas are viable. The system can be tested with new texts. Thus Narasimhan's ideas provide a computational model of first language acquisition.

Summary

We started with the question of how children acquire language, and listed features of child language acquisition. We outlined Professor Narasimhan's ideas suggesting a model of language acquisition. We then described how the Vanitha corpus was collected, and how it supports his ideas. Finally we made a brief

mention of the PLAS work that provides a computational model for these ideas.

With all this, Professor Narasimhan and his students take us closer to understanding language acquisition. But there is still a lot of work to be done to extend this and similar computational modeling frameworks.

Acknowledgements

Address for Correspondence
Raman Chandrasekar
Microsoft Corporation
One Microsoft Way
Redmond,
WA 98052, USA
Email:
Raman.Chandrasekar@
microsoft.com

Suggested Reading

[1] R Narasimhan, Language *Behaviour: Acquisition and Evolutionary History*, Sage Publications, New Delhi, 1998.

[2] R Narasimhan and R Vaidyanathan, *Language Behaviour Interaction between a Child and her Parents: An Extended Corpus*, TIFR, Bombay, 1984.

[3] B MacWhinney and C E Snow, The Child Language Data Exchange System, *Journal of Child Language*, Vol. 12, pp. 271–296, 1985.

[4] C Fillmore, *The case for case*, In E Bach and R Harms (editors), Universals in Linguistic Theory, pp.1–90. Holt, Rhinehart and Winston, New York, 1968.

[5] R Vaidyanathan, Development of forms and functions of interrogatives in children: a longitudinal study in Tamil, *Journal of Child Language*, Vol. 15, pp. 533–549, 1988.

[6] R Vaidyanathan, Development of forms and functions of negation in the early stages of language acquisition: a study in Tamil, *Journal of Child Language*, Vol. 18, pp. 51–66, 1991.

[7] V Sembugamoorthy PLAS, A Paradigmatic Language Acquisition System: An Overview, *Proc. IJCAI-6*, Tokyo, pp.788–90, 1979.

[8] R Chandrasekar and S Ramani, Interactive Communication of Sentential Structure and Content: An Alternative Approach to Man-machine Communication, *International Journal of Man Machine Studies*, Vol. 30, No. 2, pp. 121–148, February 1989.

[9] R Chandrasekar, C Doran and B Srinivas. Motivations and Methods for Text Simplification, *Proc. COLING-96*, pp.1041–1044, 1996.

[10] R Narasimhan, *Modelling Language Behaviour*, Springer-Verlag, Heidelberg, 1981.