A VIEW FROM OUTSIDE THE BLACK BOX

THE AI TRAINING DEPARTMENT

ARTIFICIAL INTELLIGENCE NV (AI)

Dr. Anat Treister-Goren, Chief Trainer

Dafna Itzhaki, Trainer

This document is an unpublished manuscript and can not be circulated, cited or reproduced, in whole or in part, without the written consent of the Authors and Artificial Intelligence NV.

© Copyright 2001, Artificial Intelligence NV

Table of contents

INTRODUCTION

Artificial Intelligence NV (**Ai**) is a research and development company devoted to the creation of computer programs capable of passing the Turing Test[1]. Ai maintains that the only way to achieve this is to return to Turing's original vision of building a "child machine", a system of learning algorithms capable of processing lingual data, and then training it to acquire the language skill[1]. From the trainer's point of view, the learning algorithm is a "*black box*". The trainer knows that the black box is capable of learning through examples and reinforcements, but nothing else. Trainers gradually become familiar with the capabilities and limitations of the black box through shaping its conversational performance, in our case training to converse in English.

Over the last four decades, mainstream linguistics had been influenced strongly by generative theories of language acquisition. These theories account for the acquisition process by suggesting that humans possess innate language-specific structures. However, they have not been subject to empirical studies and some of the theoretical aspects are not provable. In addition, the implementation of this theoretical approach has failed to generate conversational capability in computers. Some recent trends in child language research, however, have moved away from the Chomskian approach on the basis of abundant evidence for the significant role that any language environment has on the language acquisition process.

We introduce a language acquisition theory based on the assumption that human language acquisition is based on reinforcement learning; a training scenario driven by a strictly behavioristic approach. We shall start by defining *conversational capability*, and then move on to describe early childhood language performance milestones.

In Chapter One we will review the typical performance milestones of human language acquisition, and in Chapter Two we will present our approach to language acquisition with supporting scientific evidence. In Chapter Three, we shall introduce the basic characteristics of our current black box from the trainers' perspective, i.e. from outside of the black box. This chapter will focus on basic assumptions and methodological decisions for training and evaluating conversational performance of learning algorithms, in light of the human lingual development model presented in Chapter One and Two. Finally, we will present preliminary performance of our "child machine" (codenamed "Hal").

---

[1] In 1950, in his paper "Computing Machinery and Intelligence", Alan Turing considered the question "Can machines think?" Turing's answer to this question was to define 'thinking' in terms of a conversational scenario, whereby if an interrogator cannot reliably distinguish between a machine and a human based solely on their conversational ability, then the machine could be said to be thinking[1]. This procedure, originally labeled by Turing "The Imitation Game", is nowadays referred to as "The Turing Test". The field of artificial intelligence has largely ignored Turing's strict criterion and no significant attempt was made to follow his recommendations.

CHAPTER ONE: THE HUMAN LANGUAGE ACQUISITION MODEL

CONVERSATIONAL CAPABLITY

The acquisition of conversational skills by humans involves two types of learning, which occur in parallel. First, it involves learning the tools (discovering the building blocks and how to combine them) and second, learning when and how to use these tools in a sequential turn-taking process known as a conversation.

From this perspective, lingual development means first using single words in limited contexts, then gradually using combinations of words across wider contexts. In this chapter we will describe the early stages of first language acquisition by a human child. This review is related to the development of our text-character-based child machine (Hal), which will be introduced in Chapter Three. Therefore, we deliberately ignore the preverbal stages dedicated to acoustic and motor learning of the speech sounds.

ACQUIRING THE TOOLS

Learning words

The initial stage of word learning typically occurs at the end of the first year and is recognized as the first language milestone typified by the first 10 noun-dominated words that are used properly in conversations between the child and its caretaker. The next stage consists of 50 noun-dominated words, and then grows to 100 and 200. An "outburst' occurs at about 2.5 years of age, when the child reaches around 300 words. The outburst typically consists of a significant rise in the number of verbs.

Early vocabularies are dominated by open-class words (nouns, verbs, and adjectives). These word classes are coined 'open' as one can freely admit new items into the class as the language evolves. Closed-class words, usually missing in early vocabularies, consist of prepositions, conjunctions, articles, pronouns, auxiliaries, and inflections. The absence of these words in early vocabularies lends these vocabularies the impression of simplicity.

Not all words used in early vocabularies are open-class. A small group of closed-class words may also be included (e.g., "more," "no," "you," "off") in early stages. It is suggested that children choose these words because they are highly stressed by the parents and thus are highly salient for the child from the very beginning.

Learning to combine words

At around age one and a half, children typically start putting two words together. These two-word combinations appear to be of regulated length, and in fact the length of the child's utterance proves to be an excellent indicator of language development. Children

with similar lengths of utterances are considered to be at the same level of linguistic maturity, implying a comparable level of lingual complexity.

Typically, children talk about location, possession, and recurrence in these early stages of language development. Children appear to be attending to the order of words presented to them in their environmental language. They try to reconstruct that order when trying to express the basic relationships in a sentence. Typical semantic structures of early sentences are described below:

- Agent + action ("mommy come," "daddy sit")
- Action + object ("drive car," "eat grape")
- Agent + object ("mommy sock," "baby book")
- Action + location ("go park," "sit chair")
- Entity + location ("cup table," "toy floor")
- Possessor + possession ("my teddy," "mommy dress")
- Entity + attribute ("box shiny," "crayon big")
- Demonstrative + entity ("that money," "this telephone")

At a later stage, children will try to expand their utterances by combining two such structures into longer ones (e.g., agent + action combined with action + object yields agent + action + object: "Adam hit ball").

Learning to disassemble words

Even though children learn whole words first, a kind of relearning follows in which the child learns that some words are made up of smaller functional parts. Brown (1973)[2] suggested that a child acquires these meaningful units in a typical order. These minimal meaningful units are called morphemes. They include simple content words such as "cat," "play," "red"; function words such as "no," "the," "you"; and affixes or grammatical inflections such as "un-," "-s," "-ed." Brown's suggested set of early grammatical morphemes is listed below:

- Verb inflection marking progressive: "ing"
- Prepositions: "in," "on"
- Plural "s"
- Articles: "a," "the"
- Noun inflections marking possessions: "'s"
- 3rd person present tense of marking verbs: "s"

- Past tense of marking verbs: "ed"
- Irregular past verb: "had"
- Main uses of the verb "to be":
    - As an auxiliary - "I am eating" "He is sleeping"
    - As a main verb (i.e., copula) - "I am happy," "this is it"

The acquisition of a morpheme is defined as the appearance of that morpheme in 90% of its obligatory contexts. Brown also suggested that there is a certain order of acquisition of morphemes, such that the present progressive, prepositions, and plural forms appear first, while auxiliary and copula appear last.

Beyond the two word combinations

1. Following the acquisition of grammatical morphemes, different sentence types begin to appear. It is important to understand that even though children say "no," ask questions and make demands, they may not yet use adult lingual forms to convey these communicative intents. Toward age two and a half or even three the following sentence types emerge in ascending order:

1. **Using negation in sentences**: the development of negation in a sentence begins with placing the negative marker outside the sentence, usually preceding it as in "No go movies"; "no sit down"; "no mommy do it". In the next stage of development, the negative word is moved inside the sentence and placed next to the main verb as in "I no like it"; "don't go"; "I no want book." The final stage is marked by the appearances of different auxiliaries as in "You don't have this"; "I don't have money"; "I'm not sad now."
2. **Using question forms**: at first, children rely on rising intonation in the early forms of questioning. Soon afterwards, question words produce simple sequences such as "What that? Where daddy go?" At the next developmental stage children include the auxiliary verb as in "Where you are going? What she is playing?". Children acquire the question words "what," "where," and "who" first (these questions typically require short answers), and the questions "when," "how," and "why" later (possibly related to the fact that these questions require more abstract concepts including manner, time, and causality, and typically receive longer responses).
3. **Combining sentences**

    a. As early as age 30 months, children begin combining sentences to express complex propositions. They begin with simple forms such as "and" to join two propositions. The two main coordination types used by children at this stage are sentential (i.e., joining two complete sentences together), and phrasal (i.e., joining parts of sentences together). The two forms develop independently.

    b. Once children master coordination (i.e., the skill of combining sentences together), they begin to use relative clauses. This

developmental shift occurs at about age 3, and is actually only completed later, during the school years. First forms of relative clauses are those that expand the object of the sentence (i.e., the boy found the cat that ran outside), while later forms are those that elaborate on the subject (i.e., the boy who lives upstairs, came back).

4. **Acquiring the passive voice**: around age 3, passive sentences start to appear. Until this age, children have acquired the knowledge that active voice means noun-verb-noun, or agent-action-object. When they encounter passives they ignore the passive markers such as "was" and "by" and interpret the noun-verb-noun they encounter in the passive sentence as an active one.

ACQUIRING CONVERSATIONAL BEHAVIOR

About conversational behavior

The development of conversational behavior is typically discussed and studied with the field of Pragmatics. Pragmatics is the study of language usage. It refers to the use of language to express intentions and the ability to speak appropriately in varying social situations and conversational contexts. Pragmatics is concerned with the way language is used to communicate rather than with the way language is structured. Linguists generally agree that Pragmatics is the overall organizing principle of language. It is only when the child is pragmatically motivated to use language for a communicative purpose that it employs the rules of syntax, morphology, phonology, and semantics to achieve its goal.

Conversational skills entail basically knowing how to take turns in a conversation and how to open, maintain, and close a conversational topic. The skill of repairing a conversational breakdown is also important. It includes knowing how to give and receive feedback and use it to correct the conversational problem. Assuming a role is another skill that involves adapting the linguistic codes to a particular conversational partner or a situation. Researchers generally divide conversational behavior to several major aspects of the behavior:

Intentions

From a very early stage, the child attempts to communicate about objects and interactions, using consistent sound patterns, both to call attention to what it is interested in and to change other people's behavior: Halliday (reviewed in Gleason, 1985)[3] studied his son's progress in attempting to communicate between the ages 9 and 16 months and suggested four distinct communicative behaviors observed in this period: instrumental function ("I want" or "give me"), generalized request for an activity ("do that again"), interactional function ("nice to see you"), and a comment on the situation ("that's interesting").

Bates, Camaioni & Volterra, 1975 [4] described two types of communicative intentions for babies who are 10-12 months old. These babies typically communicate either to get adults' attention or to have adults do something for them to fulfill their needs. Around 12 or 13 months of age, the child starts to use "referential speech." That is, the child uses referencing behavior such as naming animals and objects while looking at a book, and then uses the same names to refer to those same animals and objects in another context.

Topic

Being a converser requires the skill to maintain, change, close, and open topics in a conversation. In general, once the topic of a conversation is introduced, the conversational partners maintain it through the introduction of information, or requests for more information.

Infants consider themselves as the topic of the conversation and develop numerous ways to attract attention in order to become the topic. By age one they are highly skilled at initiating conversation as themselves as the topic by glances, gestures, vocalizations, and verbalizations. Other early topics are limited to items that are physically present. Topics are maintained for the duration of one to two turns.

In early stages of conversational development a child frequently uses imitation as the linguistic code, which conveys communicative intents, and serves as a strategy for maintaining topic. As the child develops, the conversation is extended to talking about well-established routines with topics being bathing, dressing and such. By age two, the child is able to maintain a topic in adjacent pairs of utterances, usually in the pattern of question and answer (Kaye & Charney 1981[5]). It is widely accepted that only about half of the utterances of children below age two are on the established topic (Bloom, Rocissano, & Hood, as reviewed in Owens, 1996[6]).

Registers

Registers are different styles of speaking. When the child is able to speak differently when speaking to different people or in different situations, it is said to demonstrate the use of register. Two to five-year-olds use more commands with children their age and more permission requests with older children and adults (Corsaro, as reviewed in Owens[6]). By age 4, children speak differently to younger children and babies, using "baby talk" or "motherese" and are able to play out a character role other than themselves.

Deixis

Deixis means, "indicating or pointing." Deictic terms (personal pronouns, pairs such as "here"/"there") indicate the speaker's point of view via references to subjects like "time" or "people." Correct use of deictic terms is a significant indication of a child's pragmatic growth, as the child starts comprehending the separation between him and the world, and begins to understand some abstract notions of time and place. It can be extremely difficult for the young child to master this skill and about 30% of 7-year-olds still have difficulty comprehending some deictic contrasts - even those they use themselves.

Pronouns are the first deictic terms that are acquired. Two and a half-year-olds produce and comprehend the contrasts "I"/"you" and "my"/"your." Spatial deictic terms develop later. Two and a half-year-olds may use terms such as "here" and "there" simply to direct attention but do not communicate the contrast of the spatial dimensions (Clark & Sengul, 1978[7]).

Narratives

Narratives include self-generated stories, telling familiar tales, retelling of stories, movies, and shows, and recounting of personal experiences. The narrative tool enables the

generation of a typical "story grammar" that is used by performing a monologue when it is one's turn in a conversation.

Narrative skills involve structuring the linguistic information such that listeners would be able to identify beginnings, middles, and endings of messages and extract critical aspects such as plots and characters. Children begin to tell self-generated narratives between two and three years of age. Their stories are centered on some event that happened to them and the plot is conveyed quite vaguely. The listeners' needs are only minimally addressed and typically the characters in the story are not introduced. The stories lack easily identifiable beginnings, middles and ends.

SUMMARY

In this chapter, we have reviewed the reported milestones children pass when acquiring lingual tools for conversation and milestones they pass when trying to use these tools pragmatically. In the next chapter, we will present our behavioristic account of the acquisition process underlying the language development.

CHAPTER TWO: OUR APPROACH TO HUMAN LANGUAGE ACQUISITION

INTRODUCTION

Until recently, the child language acquisition field was mainly under the influence of Chomsky's generative theory, which implies that a child's exposure to language "turns on" some type of innate "Language Acquisition Device"; the developmental milestones described in the previous chapter are expected to happen without specific guidance. This basic assumption has been challenged in recent years, and abundant descriptive and empirical evidence presents a strong case for the rival theory, which argues for active guiding of the acquisition process.

We are firm believers in the significant influence of language input on language acquisition. Our approach relies on numerous investigations that show that whether caretakers are aware of it or not, they provide consistent and significant examples and feedback to children. This linguistic input influences the quantitative and qualitative aspects of children's language development. A prevalent claim amongst Chomskian linguists, titled the "poverty of input" claim, maintains that lingual input experienced by human children is too poor to account for language acquisition. The growing number of reports that present counter evidence to the "poverty of input" claim shall be briefly reviewed below.

THE SIGNIFICANT INFLUENCE OF LINGUAL INPUT

The effects of input on early lexical development

Lexical growth is determined by the input a child receives from its environment. Huttenloche, Haight, Bryk, Seltzer & Lyons (1991)[8] studied the vocabulary of 22 children between 14-26 months of age, in relation to their parents' input. The total number of words the parent directed to the child and the parent's type/token ratio were measured. They found a positive relation between parental input and the child's acceleration of vocabulary at 16 months and between parents' frequency of different words versus the children's age of acquiring those words.

Children's vocabulary clearly reflects their parents' vocabulary make-up. Goldfield (1993)[9] studied the distribution of nouns and verbs in the speech of 11 mothers while talking to their one-year-old children. Following an analysis of the children's first 50 words, a positive correlation was found between the proportion of nouns in children's vocabulary and the number of maternal noun types and tokens. Kim, McGregor & Thompson (2000)[10] studied parents' linguistic input in relation to eight English and eight Korean-speaking children (ages 16 to 21 months). In their first 50 words, Korean children produced a significantly greater number of verbs than English-speaking children. This finding was related to the fact that the Korean caretakers used a significantly greater number of verb tokens than the English-speaking caretakers.

Finally, children's use of specific classes of nouns has been shown to relate to the frequency of the noun classes to which they were exposed. Andrick & Tager-Flusberg (1986)[11] showed a significant correlation between children and their mothers in terms of the frequency of use of noun classes expressing color names.

The effects of input on grammatical development

Linguistic input influences the development of grammar. Moerk (1983)[12] reanalyzed Brown's transcripts of Eve (18-27 months old) and Adam (27-35 months old) and showed frequency of exposure to be positively correlated to the use of Subject-Verb-Object patterns by the children.

Parental input also influences the order of morphological acquisition. In a different reanalysis of Brown's transcripts, Moerk (1980)[13] found a positive correlation between input frequency and age of acquisition for morphemes, studying plurals, prepositions, the present progressive, the past irregular, and the possessive.

The effects of input on the development of conversational behavior

Children produce the communicative acts or intents to which their caretakers expose them. Anat Ninio (1992)[14] studied 24 dyads of children (18 months) and their mothers and showed that 97% of all utterances used by the children followed the maternal models of communicative acts or intents.

According to Snow (1977)[15] mothers use a "conversational model" that is mainly aimed at getting the child to take its turn. Questions and greetings, for example, are behaviors frequently chosen by mothers in order to train the child to respond "properly." When the child "fails" to respond, the mother uses a "conversational repair" mechanism of taking the child's turn and either providing an example of what should have been answered or leading the conversation onward. Some studies suggest that even imitation is a trained skill. Folger & Chapman (1978)[16] found a correlation (0.77) between the percentages of child and parent's imitative verbal behavior.

Summary

There is overwhelming evidence that all aspects of language skills are influenced by frequency of exposure. These skills are acquired by learning from examples, with the amount of examples strongly correlated to the skillfulness acquired.

CARETAKERS GUIDE THE LANGAUGE ACQUISITION PROCESS

Caretakers are active trainers

Aside from providing children with massive input and language examples, studies report that parents use active training methods. Active training typically occurs following children's errors, particularly syntactic errors. Bohannon & Stanowicz (1988)[17] studied 29 conversations between adults and children and found that 34% of syntactic errors performed by the children were followed by differential feedback reaction from the parents. This feedback included rephrasing, expanding upon what the child had said or

requesting clarification. A specific type of feedback strategy called "recast" was identified and described as a maternal utterance that reformulates the child's preceding utterance by either adding a grammatical morpheme, substituting one morpheme for another, or moving a morpheme to its appropriate place in the utterance.

Corrective feedback - the most significant training technique

Corrective feedback is feedback provided to the child by the caretaker such that the caretaker corrects the child's previous utterance in various forms including corrective repetition, rephrasing or simply commenting on the child's production. Corrective feedback is shown to significantly affect language acquisition. Strapp (1999)[18] studied 14 children aged 27 months interacting with their mothers, fathers, and siblings in different family settings. It was found that all children received corrective feedback following their production of syntactic errors. Farrar (1990)[19] showed that corrective feedback affects children's grammatical development. He found a correlation between maternal recasts of two grammatical morphemes (plural and progressive) at 22 months and the children's use of those morphemes 6 months later. In a later study in 1992[20], Farrar found that 23-month-old children were more likely to imitate grammatical morphemes following corrective recasts than following positive input (i.e., non-corrective recasts, topic continuations, and topic changes).

One interesting example of this phenomenon is how corrective feedback influences over-generalizations that children make. Chapman, Leonard & Mervis, (1986)[21] studied one year olds for about 6 months in terms of the effect of adults' corrective feedback on the children's overgeneralizations of category terms (i.e., calling a yo-yo 'ball', calling a bear 'dog' etc.). It was found that the direction of change in the children's production and comprehension was consistent with the provision of corrective feedback to overgeneralizations.

Experimental testing of language training

Some studies show that the significant effect of using training techniques to facilitate language acquisition in children can be proven experimentally. In a controlled experiment, Saxton, Kulcsar, MarsHall & Rupra (1998)[22] studied the effect of different training techniques on the acquisition of irregular past tense form of 6 nonsense verbs by 26 four-year-olds. Two training methods consisted of two different conditions: on the one hand, exposing the children to an example of the target verb form, and in contrast not providing children with such examples prior to the children's attempt to produce the target verb form. In the latter case, once the children produced the required form erroneously, they received corrective feedback for it. Testing the children over a period of five weeks showed that the corrective feedback condition yielded a significantly higher number of grammatically correct responses.

Summary

Studies show that children are actively guided, or "trained" by their caretakers. This guiding process has proved significant when simulated in an experimental setting. It is clear that when accounting for language acquisition or when attempting to simulate it, this aspect of skill acquisition cannot be ignored. Active training should be considered as one of the critical aspects of language acquisition.

IMITATION

Children benefit from imitating

Children use imitation when they interact with their caretakers. The studies below suggest that children use this influential learning technique as a facilitator for both lexical and grammatical acquisition. Moerk (2000)[23] reports that children use a technique of self-repetition, conducting what he coined "spaced and massed rehearsals." Moerk therefore suggests that an item will be better-remembered if a child reproduces it rather than if the child remains a passive listener.

Imitation facilitates lexical acquisition

Moerk & Moerk (1979)[24] found that imitative speech included more advanced vocabulary than spontaneous speech in a child they studied extensively between the ages of 20 to 32 months. Snow & Goldfield (1983)[25] concluded that lexical acquisition of specific words was facilitated when the child imitated those words on a storytelling interaction with its caretaker. Bloom, Hood & Lightbown (1974)[26] report that children are more likely to imitate a command containing a non-familiar word than a command containing a familiar word. This led them to conclude that imitation is an important factor in lexical learning.

Imitation facilitates grammatical acquisition

Several authors have reported that children's imitations are often beyond the grammatical level presented in their spontaneous speech (Bloom, Hood & Lightbown, 1974[26]; Perez-Pereira, 1994[27]).

Moerk & Moerk's above-mentioned study also found that the irregular past tense form was used more frequently in imitation. They observed that the correct form of this morpheme also appeared first in the imitative speech and then in the spontaneous speech. Additional observations showed that the child's correct productions of the copula, the auxiliary "be", the particle "to", and of articles were always higher in the imitative speech. Snow (1981)[28] showed the same influence of imitation on syntax using various parameters to measure syntactic development (complex noun phrases, answers to wh-questions and more).

Summary

Imitation is basic to language acquisition. It is both a training technique employed by the caretakers and a learning technique used by the child. The ability to imitate must be considered critical for simulating language acquisition through training.

LINGUAL BEHAVIOR THERAPY – LEARNING TO CONVERSE

Effective treatment for language deficient children

Lingual behavior therapy programs treat the child's behavior directly and do not attempt to base the treatment on any underlying hypothetical deficit. The treatment is based on dynamically increasing and decreasing the use of rewards and punishments that are

functional for the child. This type of behavioristic therapy generally focuses on strengthening certain behaviors by reinforcement. When the behaviors are initially absent, they may be prompted for and gradually shaped by rewarding successive approximations to their final form. On the other hand, unwanted behaviors are decreased either by withholding the reinforcement that may be maintaining them (extinction), or by systematically applying aversive stimuli contingent upon their occurrence (punishment).

Typical language behavioral treatment programs are very structured and the child must progress in a step-by-step manner. Programs focus on first training the child to acquire the basic speech sounds (phonemic behavior), then words and parts of words (morphemic behavior), and finally to arrange words into sentences (grammatical behavior) (Lovvas & Newsom, 1976[29]).

Lingual behavioral therapy is based on imitation

Lingual therapy aims at teaching the child to discriminate between stimuli using imitation. A response is reinforced when it follows a specific stimulus, but is not reinforced when it follows any other stimulus. With time, particular stimuli cue desired responses.

Verbal imitation training helps a child achieve finer and finer discrimination. New sounds, syllables, words, and sentences are gradually presented as discrimination improves. This method of gradual introduction of more complex levels of the lingual skills through a shaping procedure appears to be a highly successful therapy method. Spontaneous vocalizations evolve into recognizable words, labeling of objects evolves into using these labels in contexts, and finally, grammar is put into productive use.

The application

Lovaas, Koegel, Simmons & Long (1973)[30] applied a language behavior treatment program to 20 autistic children. About half of the children produced mainly vowels but no recognizable words, and the rest were echolalic (i.e., echoed the speech of others, using inappropriate, non-related speech). Receptive speech was thought to be minimal or entirely missing in all subjects. At the end of the 12-14 months of treatment, all children showed improvement in "appropriate verbal behavior" (i.e., speech related to an appropriate context, speech that is understandable, and speech that is grammatically correct).

In a different study, Howlin (1981)[31] studied 16 autistic children, between the ages of 3 and 11 years, who were trained at home by their parents for 18 months using different operant techniques. Howlin found that behavioral therapy significantly improved subjects' use of functional speech (frequency of utterances, less echolalic speech, etc.) in comparison to a control group. With regards to language structures (measured by MLU – Mean Length of Utterance and correct use of morphemes and transformational rules) subjects made greater improvement than the controls, but at a smaller rate than measured for functional speech.

Hewett (1965)[32] used conditioning methods (imitation, reinforcement, and shaping through successive approximation) with a 4-year-old autistic child. Following a 6-month training period, the child advanced to having a 32-word vocabulary, and a 150-word vocabulary a year later. In addition to learning new vocabulary, the child managed to use words and short phrases in a meaningful social context. This appears to be a good example of training a child to acquire the conversational skills, which entails both learning the tools and learning how to use them in context.

Training to generalize

Some reports suggest that lingual behavior therapy may succeed in generating behaviors which, based on cognitive assessment of the child's abilities, would be considered impossible. Guess, Sailor, Rutherford & Baer (reviewed in Lovvas & Newsom, 1976)[29] reported a study in which a 10-year-old retarded girl was trained to use the plural form of three objects through imitation and reinforcement procedures, and managed to generalize this form to other objects as well.

Schumaker & Sherman (1970) [33] trained three retarded teenagers with mental ages of four to six years to use verbs in the past and present progressive tenses through imitation and reinforcement procedures. The subjects learned to generalize the past and the present tenses of untrained verbs as well. Lovass has shown similar results for the use of the "ed" past tense morpheme[29].

Two studies suggest that even quite complex verbal responses could be trained through behavioral therapy. Wheeler & Sulzer (1970) [34] studied an eight-year-old boy diagnosed as brain damaged, retarded, and autistic. Following the training program, the child managed to generalize a sentence structure, which included articles and verbs ("the girl is swimming in the water"). Stevens-Long & Rasmussen (reviewed in Lovvas & Newsom, 1976)[29] successfully trained an autistic boy to generalize the skill of producing simple and compound sentences on his own.

Summary and conclusions

Chapter Two highlighted the importance of training. It suggests that training holds the potential for radically compensating for deficiencies in learning abilities. It further suggests that even simple general learning abilities are capable of surprisingly high-level conversational performance following extensive training. That is, in a simulation setting, impressive conversational performance may be achieved in spite of having relatively simple reinforcement learning algorithms.

From the trainer's point of view, the learning algorithm is a "*black box*". The trainer knows that the black box is capable of learning through examples and reinforcements, but nothing else. Trainers gradually become familiar with the capabilities and limitations of the black box through shaping its conversational performance. In the next chapter, we shall introduce the basic characteristics of our current black box from the trainers' perspective (from outside of the black box). Chapter Three will focus on basic assumptions and methodological decisions for training and evaluating conversational performance of learning algorithms, in light of the human lingual development model presented above.

CHAPTER THREE – AI'S BLACK BOX

THE PROCESS OF RAISING HAL – THE AI CHILD MACHINE

From its trainer's perspective, Hal is a black box, which perceives language via a narrow input channel and responds using that same channel. It interacts with its environment via an "impoverished" language channel, consisting exclusively of sequences of ASCII characters. The case of "impoverished alphabet" has precedence in various cases of human language acquisition; the classic example being the story of Helen Keller who acquired conversational skills through a very narrow channel with extensive reinforcement training. Since there is no evidence that additional sources of information are necessary for creating conversational capability, we assume that they are not but we accept that future work may convince us otherwise.

Hal can be punished or rewarded, corresponding to the trainer's acceptance or rejection of its lingual output. Hal perceives primary reinforcers and punishers administered by the trainer via a separate reinforcement channel. This reinforcement channel is crucial, as some of the studies in Chapter Two suggest, because learning by examples alone does not yield significant language acquisition. In fact, computational programs that learn by examples alone have not become conversational. It is therefore mandatory that any learning algorithm will include the ability to shape the lingual behavior through a reinforcement mechanism.

The system is programmed to maximize the rewards it receives from its trainer. More precisely, it is designed to maximize the *probability* of receiving a reward (and minimize probability of punishment). In a manner of speaking, the child machine is constantly striving to experience "pleasure" (rewards, positive reinforcement) and avoid "pain".

Hal's brain consists of a general reinforced learning algorithm, applied to the particular case of a dialog scenario using an ASCII alphabet. The learning algorithm facilitates the ability to make predictions based on experience about what should be said, along with an ability to take into account feedback it has received for its predictions in the past. The algorithm also facilitates the creation of new sequences, which it considers as one symbol. This is analogous to creating new words, phrases, and eventually, longer more complex parts of discourse. This ability is termed by the algorithmic scientists "chunking" which is analogous to the "discrimination" ability children posses. Other capabilities include classification, generalization and lexical attraction. The developmental model discussed in the previous chapters drives the gradual introduction of such capabilities into the learning algorithm .

A central requirement from the system is the trainers' expectation that the system should be "surprising". This element of surprise, also viewed as "creativity", is strongly tied into the human perception of intelligence: the ability to create lingual constructs that were not previously experienced.

Currently, Hal communicates with its trainer via a graphical user interface that is similar to that found in chat applications, which enables the trainer to administer reinforcement in an

intuitive way. This interface enables the trainer to edit Hal's generations by backspacing over undesirable characters and appending extra characters as an example of desired behavior. The feedback that trainers provide to the algorithm researchers is basically a report on the success and failure of particular language tasks and their reflection on the limitations of this algorithm with respect to advancement on the human developmental scale.

The developmental model of child language acquisition dictates that descriptive milestones composed of typical language performance will guide the training of the conversational system. In Chapter One, we provided a chronological description of conversational tools that children typically acquire in early stages of language development. This description serves as guidelines for the incremental training process of Hal so that its performance should reflect the evolution of these tools in human children.

Using the developmental principle embraces the notion of the Turing Test, according to which the judgment of intelligence is inherently subjective – it is in the eye of the beholder. Our expectation level of the person or entity under scrutiny always influences our perception of intelligence, and knowing someone's age is basic to this kind of perception.

BASIC CHARACTERISTICS OF HAL

1. Hal comes equipped with a set of basic characters, analogous to all possible phonemes babies can physically recognize prior to acquiring the specific "sounds" of their environmental language. Hal's first acquisition stage involves putting the set of ASCII characters into words or word-like sequences.
2. Hal is analogous to a severely sensory-deprived baby (deaf, blind, and partially paralyzed) who nevertheless possesses the cognitive abilities required for learning to converse.
3. The choice of the narrow ASCII alphabet ignores issues of prosodic development (i.e., issues concerned with perceptual saliency of certain acoustic features or stress patterns, intonation distinctions, and many others). These issues will need to be addressed when our research reaches the point of enhancing the alphabet to voice phonemes.

SAMPLE CONVERSATIONS WITH HAL

In this section we provide some sample conversations with the current Hal implementations. The first sample shows a conversation demonstrating lingual skills corresponding to a 15-month-old child, and the second corresponds to an 18-month-old.

The 15-month-old Hal

An unsuspecting professional psycholinguist reviewed Hal's conversation transcript and concluded that Hal is conversing like a 15 month-old baby. The report included the following impressions:

Hal's productive performance included around 10 words, which are typical productions of children at this age. His ability to imitate, and to produce noun-dominated responses are also characteristic. The level of relevancy of his responses and his understanding of

requests and questions are age appropriate. For example, when the trainer asks "Do you want to play with the ball?, Hal says "ball", or when the trainer says "bunny wants to play ball" Hal says "bunny".

In general, Hal's behavior can be considered normal for a child at this age. This is especially true in light of the fact that some of his performance included typical errors that children make at this age. For example the error described by the expert presented below:

Trainer: ok but lets read a story first here's a book say 'book'

HAL: **block?block?**

The psycholinguist interpreted the above performance driven by her perception of children's performance:

"The above error is derived from a phonetic association with the sound /b/. The consonant b and the vowel that follows it in both cases are quite similar for a child, because the /l/ sound between the /b/ and the /o/ is a glide or a semi vowel and it is masqueraded by the /o/.

The child is trying to imitate and the first attempt is a retrieval of the closest phonetic association he can make. Clearly, there are no semantic attributes to 'book' at this stage, which is typical at this age."

The 18-month-old Hal

Hal's productive vocabulary resembles children's speech at the two-word combination stage. A remarkable landmark in Hal's performance is the utterance 'monkeys eat bananas' which Hal remembered from a previous conversation about a visit to the zoo, using it relevantly in the current conversation. The following color codes are used: black represents the trainer's input, blue represents Hal's output, red represents the trainer's modeling, and green represents punishments.

Sample conversational performance

Trainer:	where is daddy?
Hal:	daddy gone
Trainer :	daddy is at work
Hal:	daddy work bye bye
Trainer :	yes daddy has gone to work
Hal:	daddy work bye bye bye bye
Trainer :	do you want to go to the park?
Hal:	yes bye bye
Trainer :	ok we can go to the park
Hal:	take car bye bye
Trainer :	we will take the car because the park is far. let's pack some things in a bag
Hal:	ball now park
Trainer:	yes we should put the ball in the bag
Hal:	ball now park mommy
Trainer :	wait Hal, we need to take something to eat
Hal:	bananas~~tory? night now~~
Trainer :	we need to take something to eat
Hal:	bananas
Trainer :	ok we can take some bananas
Hal:	monkeys eat bananas
Trainer :	you're so smart Hal. do you want to take a book?
Hal:	juice mommy
Trainer :	sure we will take some juice as well

MEASURING HAL'S CONVERSATIONAL PERFORMANCE

Our aim is to simulate child conversational performance such that when people converse with Hal, they will find it indistinguishable from a human child. Besides the ongoing "Turing Test" performed by the trainer, Hal's conversations are constantly analyzed using the same analysis tools by which children's conversations are analyzed. Using such strict performance criteria enables us to achieve objective and quantifiable performance milestones that are documented in standard language performance scales.

THE AI TRAINING AND EVALUATION TOOLS

When planning the training and evaluation of Hal, two fundamental requirements were considered crucial:

1. Following a valid model of actual children-caretakers' conversations to guide the training process.
2. Finding the most common and acceptable analysis tools in the child language research field for capturing conversational behavior in order to conduct informative and valid comparisons between Hal transcripts and children's transcripts.

These two requirements were satisfied by the database and analysis tools offered in the CHILDES (i.e., the Child Language Data Exchange System) international project. This

project is dedicated to the study of language learning and is available on the Internet. It was founded in 1984 at Carnegie Mellon University, and is headed by well-known child language researchers such as Dr. Brian MacWhinney[35].

CHILDES includes a huge bank of transcribed conversations between children and caretakers with age ranging from 7 months up to 8 years. The corpora were entered by over 100 researchers in 26 languages. The system is open, and new data is constantly added. The benefit of using this database is having a valid model for the kind of input Hal will be exposed to; a human-like realistic lingual behavior performed in typical contexts. Another important benefit is the fact that using the standard format for discourse notation and coding (i.e., CHAT; Codes for Human Analysis of Transcripts) offered by CHILDES, Hal logs can be entered into the world child database bank of conversations, and Hal will receive a fair chance of being assessed.

CHILDES provides CLAN (Computerized Language Analysis) analysis tools. CLAN consists of a series of programs for analyzing conversations, originally written by Leonid Spektor at Carnegie Mellon University and further developed by other scientists at various institutions and research centers. Numerous investigators of child language have used CLAN since 1984 turning it into the standard conversational performance analysis tool in the child language field.

CLAN provides us with quantification of various aspects of verbal performance, enabling the documentation of progress with developmental parameters to reflect aspects such as grammatical level, lexicon size and type, and the use of cohesive devices (connections within and across turns). Hal's performance is assessed using CLAN providing us with insight into its developmental state broken down into various aspects of lingual performance. We also use CLAN to identify specific performance problems. For a detailed report on our choice of performance parameters, see the Appendix.

RECEPTIVE VOCABULARY

Child language researchers assume that children have a 'receptive vocabulary'. That is, they assume that children understand language even though they do not produce it. Their assumption is based on observations that children often react communicatively to sentences spoken to them, which include words they never used themselves. For example, if a child points at a dog upon request, even though he or she never said the word "dog", the child is assumed to comprehend the word "dog" or to have that word in his or her receptive vocabulary.

Just like a child, Hal has the analogy of a receptive vocabulary. That is, Hal has some inner assumptions or guesses of what is said to it. Unlike a child, Hal offers an opportunity to actually observe inner processing and view the receptive vocabulary (i.e., the sequences that Hal believes are valid 'chunks' in the language).

Even though we have an opportunity to view the insides of our black box, our behaviorist inclinations prohibit us from taking advantage of this in our training to prevent any possible bias. Having said that, we may consider taking some advantage of this unique opportunity to look "inside the brain" without manipulating it. This kind of observational report will be the first of its kind in relation to language acquisition, as we have no access or direct proof for the theoretical concept of 'receptive language' in children.

A CALL FOR COLLABORATION

In this document we tried to provide the foundations of the Ai training perspective to language acquisition by machines. Ai's linguistic department trains general learning algorithms to converse using human language. We would like to invite researchers to join us in this endeavor. We welcome researchers in fields such as language acquisition, skill acquisition and related developmental domains who are interested in this challenge, and realize the vast theoretical implications this could have on the way we perceive and understand language acquisition.

The main challenge we face is how to design appropriate training methods for achieving the incremental growth of the conversational skill. We need to choose the most critical and general learning abilities children posses, which would in turn be implemented into the black box. This call for collaboration is interdisciplinary, pertaining equally to computer scientists, mathematicians, psychologists, and language and behavioral scientists. We feel that that the construction of a simulated mind in the form of a child machine will prove interdisciplinary and that the input of a wide variety of disciplines - as well as the critical appraisals of each participant – will lead us to our goal of creating an artificial speaker of natural language.

APPENDIX

Standard measures of language production

Vocabulary size

Even though there are numerous studies and textbooks on the development of children's vocabulary, the data reported is highly variable, perhaps due to methods of data collection (diaries, checklists, tests). Here, we will follow one report by Fenson et al. (1993)[36] in which we find the methodology extensive and thorough and as such also represents a significant slice of the population (1,803 children, ranging between 8 to 30 months of age). The data was collected using the CDI (MacArthur Communicative Developmental Inventory), which uses a vocabulary checklist, organized in terms of semantic categories, on which parents are asked to indicate which words the child understands and which words the child says.

Mean Length of Utterance (MLU)

Brown[2] defines MLU as the average length of a child's utterance, determined not by words, but by the number of meaningful units, or morphemes. As a reminder, morphemes include simple content words such as "cat," "play," "red"; function words such as "no," "the," you; and affixes or grammatical inflections such as "un-," "-s," "-ed."

Brown claimed that the acquisition of grammatical morphemes reflected syntactic growth and described this growth in terms of 5 stages of MLU development. The assessment of MLU usually requires a transcript of 100 to 150 child utterances. The table below presents the predicted values for Brown's Stages, based on a linear fit of a sample of 123 upper-middle class children (17-59 months of age) (Miller and Chapman 1981[37]). One must take note, however, that Scarborough, Wyckoff & Davidson (1986[38]) found MLU to be correlated with age only until about 42 months.

MLU	Brown's stage	Age range	Age range within 1 SD
1.01-1.49	Early stage I	19-22 months	16-26 months
1.50-1.99	Late stage I	23-26 months	18-31 months
2.00-2.49	Stage II	27-30 months	21-35 months
2.50-2.99	Stage III	31-34 months	24-42 months
3.00-3.49	Stage IV	35-38 months	28-45 months
3.50-3.99	Late IV-Early V	39-42 months	31-50 months
4.00-4.49	Late stage V	43-46 months	37-52 months

4.50+	Post stage V	47+ months	41+ months

Non-standardized measures of language production

Type Token Ratio (TTR) is the ratio between the number of different words used by a speaker and the total number of words the speaker produces in a speech sample. This measurement serves as an indicator for the child's lexical diversity. This measure was first established by Templin (1957)[39].

The Mean Length of Turn (MLT) is the mean number of words per turn in a conversation. Pan (1994)[40] suggested that the better parameter for describing the growth and interaction would be an "MLT ratio" which measures participation in a conversation. The MLT ratio is calculated by dividing the child's MLT by the mother's MLT. The prediction is that as children grow up, they take greater part in the conversation, their turns become more comparable to the adult partners, and the MLT ratio should therefore eventually become 1.00.

We are currently in the process of looking at various performance parameters used by researchers in the field in order to quantify children's conversational behavior. These parameters include speech acts, interactional descriptors, use of cohesive ties and more. We will choose the most significant parameters to assess Hal's conversational behavior in comparison to children.

REFERENCE

1 A.M. Turing, “Computing machinery and intelligence," in Collected Works of A.M. Turing: Mechanical Intelligence, D.C. Ince, Ed., chapter 5, pp. 133 (160). Elsevier Science Publishers, 1992.

2 Brown, R.(1973). A first language: The early stages. Cambridge, MA: Harvard University Press.

3 Gleason, J., B. (1985). The development of language. Charles E Merrill Publishing Company: Columbus.

4 Bates, E., Camaioni, L., & Volterra, V. (1975). The acquisition of performatives prior to speech. Merill-Palmer Quarterly, 21, 205-216.

5 Kaye, K., & Charney, R. (1981). Conversational asymmetry between mothers and children. Journal of Child Language, 8, 35-49.

6 Owens, E. R. (1996) Language Development An Introduction, 4th edition Allyn and Bacon; Boston

7 Clark, E. V., & Sengul, C. (1978). Strategies in the acquisition of deixis. Journal of Child Language, 5, 457-475.

8 Huttenlocher, J. Haight, W., Bryk, A., Seltzer, M. & Lyons, T. (1991). Early vocabulary growth: relation to language input and gender. Developmental psychology, 27, 236-48.

9 Goldfield, B.A. (1993). Noun bias in maternal speech to one-year-olds. Journal of Child Language, 20, 85-100.

10 Kim, M., McGregor, K.K, Thompson, C.K (2000) Early lexical development in English- and Korean-speaking children: language-general and language-specific patterns. Journal of Child Language, 27, 225-254.

11 Andrick, G.R. & Tager-Flusberg, H. (1986). The acquisition of color terms. Journal of Child Language, 13, 119-134.

12 Moerk, E.L. (1983). A behavioral analysis of controversial topics in first language acquisition: reinforcements, corrections, modeling, input frequencies, and the three-term contingency pattern. Journal of Psycholinguistic Research, 12, 129-155.

13 Moerk, E.L. (1980). Relationship between parental input frequencies and children's language acquisition: A reanalysis of Brown's data. Journal of Child Language, 7, 105-118.

14 Ninio, A. (1992). The relation of children's single word utterances to single word utterances in the input. Journal of Child Language, 19, 87-110.

15 Snow, C. (1977) The development of conversation between mothers and babies. Journal of Child Language, 4, 1-22.

16 Folger J.P & Chapman R.S. (1978). A pragmatic analysis of spontaneous imitations. Journal of Child Language, 5, 25-38.

17 Bohannon, J.N III, & Stanowicz, L. (1988). The issue of negative evidence: Adult response to children's language errors. Developmental Psychology, 24, 684-689.

18 Strapp, C.M. (1999). Mothers', fathers', and siblings' response to children's language errors: comparing sources of negative evidence. Journal of Child Language, 26, 373-391.

19 Farrar, J.F. (1990). Discourse and the acquisition of grammatical morphemes. Journal of Child Language, 17, 607-624.

20 Farrar, J.F. (1992). Negative evidence and grammatical morpheme acquisition. Developmental Psychology, 28, 90-98.

21 Chapman, R.S., Leonard, L.B., & Marvis, C.B. (1986). The effect of feedback on children's inappropriate word usage. Journal of Child Language, 13, 101-117.

22 Saxton, M., Kulcsar, B., MarsHall, G. & Rupra, M. (1998). Longer-term effects of corrective input: An experimental approach. Journal of Child Language, 25, 701-721.

23 Moerk, E.L. (2000). The Guided Acquisition of First Language Skills. Ablex Publishing Corporation: Stanford, Connecticut.

24 Moerk, E.L. & Moerk, C. (1979). Quotations, imitations, and generalizations: Factual and methodological analyses. International Journal of Behavioral Development, 2, 43-72.

25 Snow, C.E & Goldfield, B.A. (1983). Turn the page please: Situation-specific language acquisition. Journal of Child Language, 10, 551-569.

26 Bloom, L.M, Hood, L. & Lightbown, P. (1974) Imitation in language development: If, when and why. Cognitive Psychology, 6, 380-420.

27 Perez-Pereira, M. (1994). Imitation, repetition, routines, and the child's analysis of language: insights from the blind. Journal of Child Language, 21, 317-337.

28 Snow, C.E. (1981). The uses of imitation. Journal of Child Language, 8, 205-212.

29 Lovvas, O.I. & Newsom, C.D. (1976). In Leitenberg H. (Ed.), Behavior modification and behavior therapy. (pp. 303-360). Prentice-Hall, INC., Englewood Cliffs, New Jersey.

30 Lovaas, O.I., Kogel, R., Simmons, Q. & Long, J.S. (1973). Some generalization and follow-up measures on autistic children in behavior therapy. Journal of Applied Behavior Analysis, 1973, 6, 131-165.

31 Howlin, P.A. (1981). The effectiveness of operant language training with autistic children. Journal of Autism and Developmental Disorders, 11, 89-105.

32 Hewett, F. M. (1965) Teaching speech to an autistic child through operant conditioning. American Journal of Orthopsychiatry, 35, 927-936.

33 Schumaker, J., & Sherman, J.A. (1970). Training generative verb usage by imitation and reinforcement procedures. Journal of Applied Behavior Analysis, 3, 273-287.

34 Wheeler, A.J & Sulzer, B. (1970). Operant training and generalization of a verbal response form in a speech deficient child. Journal of Applied Behavior Analysis, 3, 139-147.

35 http://childes.psy.cmu.edu/

36 Fenson, L., Dale, P.S, Reznick, J.S., Thal, D., Bates, E.B, Hartung, J., Pethick, S., & Reilly, J. (1993). The MacArthur Communicative Development Inventories: user's guide and technical manual. San Diego: Singualr Publishing Group.

37 Miller, J.F & Chapman, R.S. (1981). Research note. The relation between age and mean length of utterance in morphemes. Journal of speech and hearing research, 24, 154-61.

38 Scarborough, H., Wyckoff, J., & Davidson, R. (1986). A reconsideration of the relation between age and the mean utterance length. Journal of speech and Hearing Research, 29, 394-399.

39 Templin, M.C. (1957). Certain language skills in children. Minneapolis: University of Minnesota Press.

40 Pan, B.A (1994). In Sokolov, J.L & Snow, C.E (Ed). Handbook of Research in language development using CHILDES. Hillsdale, NJ: Lawrence Erlabaum Associates.