SECOND LANGUAGE LEARNING AND ACADEMIC ACHIEVEMENT REVISITED: A CONNECTIONIST DECONSTRUCTION OF THE COGNITIVE CONSEQUENCES OF BILINGUALISM

by

ANNE HAWSON

B.Sc., University of St. Andrews, 1964
M.Sc., University of Western Ontario, 1968

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE STUDIES

(Centre for the Study of Curriculum and Instruction)

We accept this thesis as conforming
to the required standard

THE UNIVERSITY OF BRITISH COLUMBIA

March 1996

~~Department of~~ CENTRE FOR THE STUDY OF CURRICULUM AND INSTRUCTION

The University of British Columbia
Vancouver, Canada

Date April 3, 1996.

ABSTRACT

The original Threshold Hypothesis proposed in 1976 by Cummins as a reasonable explanation of the data gathered to that date on the cognitive consequences of learning a second language has gained widespread acceptance in second language teaching circles, but has been the subject of considerable debate among second language researchers. An examination of the critiques of the hypothesis as well as the responses invoked by these critiques reveals that much of the criticism is well-founded. Given that the Threshold Hypothesis proposes that "cognitive" benefits or deficits can accrue through the learning of a second language, the conceptualization of cognition underpinning the hypothesis is important in understanding its inadequacies. Upon examination of Cummins' work, it is concluded that cognitive processing is considered by Cummins to be a discrete category of brain processing operating at the conceptual level.

A fundamental premise of this thesis is that connectionism, which proposes that cognition is more appropriately described as operating on a sub-conceptual level, and as being intimately intertwined and dependent upon what have been regarded as "non-cognitive" aspects of brain functioning such as emotion, attention, and sensory processing, provides a new and potentially enlightening perspective on cognitive issues. A critique of connectionism as a valid model of learning in the brain suggests that it has considerable validity, but cannot yet describe all types and/or all processes involved in learning. Nevertheless, it is concluded that a model of second language learning built upon the conjectures of connectionism could bring new insights into how second language learning in immersion situations affects brain processing, and thereby provide clues which might explain perceived cognitive benefits and deficits.

The model developed proposes that, under immersion conditions, second language learners may be undergoing an attentional shift in information processing towards the visual system, and away from the auditory. If this hypothesis can be substantiated by empirical tests, it could explain why second language learners in immersion situations are in danger of becoming poor readers, and may provide fertile ground for the development of remedial educational treatments. This is a wholly theoretical study based on literature research.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

Thanks are due to all my committee members, Dr Stephen Carey, Dr Victor Froese, and Dr Peter Reiner, for staying the course with me, and for encouraging me in my efforts. Special thanks go to Dr Carey for all the stimulating discussions as the work was progressing and for allowing me to pursue my own line of thinking, as well as to Dr Reiner for his technical assistance and continual support for this project. I would also like to acknowledge Dr Walter Werner of the Centre for the Study of Curriculum and Instruction without whose willingness to listen to my ideas and present challenging arguments this thesis would have suffered. With respect to financial support, my gratitude goes to Dr John Willinsky, director of the Centre for the Study of Curriculum and Instruction, who, as well as encouraging me in my research, generously found work for me as a research and editorial assistant within the centre, and also to the Department of Language Education which provided funds for me to attend conferences in the United States. I am indebted also to Dr Steve Petersen of Washington University School of Medicine, Dr Ernst Niebur of The Johns Hopkins University, and Dr MaryLou Cheal for their generous advice, and for forwarding me papers on attention.

Last, but by no means least, I would like to thank my family, my daughters Fiona and Kirsten, and my husband Herb, for their patience, forbearance, and moral support throughout a process which has been much more time and energy consuming than had ever been anticipated. I could never have made it without them.

PREFACE

When I think back to what might have first initiated my interest in the topic of second language learning and academic achievement, there are many images that come to mind: my upbringing in Scotland which made me "bilingual" in Scottish dialect (Dundonian, to be exact) and BBC English; being the first in the whole of my extended very working-class family to attend university; emigrating to Canada as a young woman and finding it difficult to understand Canadians and be understood by them, not only linguistically but culturally; living and teaching in Jamaica and Egypt where the stress of being in a completely different culture took its toll on the whole family, both in general health terms and ability to function well mentally; my experiences becoming fluent in French. There are probably many more reasons, but these seem to me to be the most important.

If I try to rationalize why I became fascinated with how the brain works, however, there are only three incidents spread out over a period of twenty years which come to mind. The first came about as a result of my wanting to inform myself about educational theory in general when my children were just entering the school system. Rummaging through the bookshelves in the local library, I came across a small text written by Maria Montessori in 1949 (and translated into English in 1955): Childhood Education. In it I read the following:

> a prejudice has found its way into the adult - the notion that the life of the child can be changed or improved only through teaching. This prejudice impedes the understanding of the fact that the child constructs himself, that he has a *teacher within* himself, and that this inner teacher also follows a program and a technique of education, and that we adults by acknowledging this unknown teacher may enjoy the privilege and good fortune of becoming his assistants and faithful servants, by helping him with our cooperation. (p. 63)

It was the notion of this "teacher within" that caught my imagination, a self-guided teaching mechanism in tune with the needs of the person it inhabits which would initiate whatever learning was appropriate for the particular stage that the person had reached developmentally. How would such a system work? Would it be essentially the same in animals and humans? What would this teacher look like? Could such a mechanism always be correct? Would it only be present in the young, or could it be an on-going process throughout life? Would its effectiveness be largely genetically determined, or would experience be an effective influence? Through my researching of connectionism, I now feel satisfied that I have some insight which might help in answering these questions.

The second incident that headed me in the direction of brain studies, and in particular towards the study of attention, happened while I was teaching in Cairo at an international school. The school was situated close to a railway track which crossed the main road leading up to it. So every morning, I and many others had to cross the tracks at the start of the day. This one morning, I noticed a young girl about 11 years old walking towards the track some distance ahead of me. She looked completely absorbed in her thoughts, walking steadily towards the track by herself, never lifting her head to see where she was going. I became a little anxious about her when I heard a train, headed by an old-fashioned steam engine, approaching the crossing, but the driver was clanging the big bell by his cabin and I felt sure she would take notice eventually. As I watched her and the train drawing closer and closer I suddenly became really alarmed: she wasn't raising her head in spite of the fact that the noise from both train and bell was horrendous, so much so that when I yelled out a warning as loudly as possible I couldn't hear my own voice - and *still* she walked on doggedly. I began to run, sick in the knowledge that I would never reach her before the train did; the driver was leaning precariously out of his cabin waving his hands vigorously and calling to catch her attention - all to no avail. She stopped, finally, when she had one leg raised to cross the rails and the train was about two feet from her leg, jumping back just in time to avoid being run over. When I caught up with her a few seconds later, she looked quite calm though surprised, and I gasped out "Didn't you hear the train?" She replied "No. But I hear it now." Later in school, when I regaled the assistant headmistress with the story she said "You know, I don't know what's wrong with these kids. Just a little while ago, someone else walked into the side of a train in that same area and had his leg broken". For me the questions abounded: How could someone not hear a noise that loud at one moment, and hear it the next? How could someone not see a huge train just a few yards away? What kind of mechanism can "lock out" information so central to survival? Did living in a country in which the language and culture were completely foreign have anything to do with all of this? I now feel that I have the beginnings of answers to these question through my study of attention.

The third experience which guided my search came through teaching French as a foreign language in a private school to grade 7 and 8 students here in Vancouver. Because the children had come from various schools within the area (the school was newly opened at the time), they had all had different instruction in French and so were very varied in their levels of expertise, although they had been streamed into the lower class in this school. One thing they did have in common was that they did not like French, and some of them actively loathed it. Needless to say, discipline in French class was a problem, and the amount of learning that went on there was minimal. One morning at home, I happened to see a children's animated film called Remi on the French channel. It was a story told in many weekly episodes about the life of a young boy of around 10 years old who wandered around France at the turn of the century as an apprentice to a travelling performer and his animals.

The film was very lively, beautiful visually, and accompanied by moving music, but the French was of a very high level, difficult enough for me to understand never mind these students. I videotaped it nevertheless, and showed it to them on the premise that if they liked it I would tape some more for them, if they didn't, then they just had to let me know and that would be that. They loved it in spite of not being able to understand a word of what was said. We developed a pattern for watching the shows: I gave them an overview in English of what they were going to see without giving away the plot completely, and supplemented this with comments during the viewing of the program; we watched the whole half hour episode and then they asked me for the vocabulary that they would need to write a few sentences for homework telling what they had seen. Within a couple of months they had not only become astonishingly keen on doing the writing, but their interest level in all aspects of French had done a quantum leap. As I watched this process, with great relief and delight I might add, it was clear to me that it was the emotional content of the program which had inspired it. The story line was full of great hardships, and even deaths, as well as laughter and joy, all of which was reflected emotionally on the faces of the students as they watched. And once again my mind was caught up in questioning: Why was it that emotional content initiated such a change in interest in learning? How could this ostensibly "non-cognitive" aspect of brain processing be so effective in promoting learning? I now have tentative answers to these questions too thanks to the insight of Antonio Damasio.

Needless to say, perhaps, I have drawn a great deal of satisfaction from obtaining insight into both the hugeness of my questions and how neuroscience might provide enlightening responses to them. Though the researching and writing of this dissertation has meant thousands of hours of work, often mind- and eye-wearying, there have been moments of great joy as the lights went on, illuminating dark corners of my knowledge on how the mind/brain works and lifting my spirits with the appreciation of the paradoxical complexity yet simplicity of its function. It is my hope that other educators will find the results of my searching both interesting and useful.

INTRODUCTION

Ascribing general cognitive benefits or deficits to a learner or, even worse, to groups of learners is a matter which requires careful reflection. This is especially so as the repercussions of such a designation can grow out of all proportion to the magnitude and validity of the informational basis upon which the judgement was made. It will be argued in this thesis that one highly influential case in point is the Threshold Hypothesis of Jim Cummins, first delineated in 1976, which associates cognitive benefits and deficits with the learning of second languages, particularly under immersion*[1] conditions. However, his hypothesis does not stand alone. In the last three decades in particular, hypotheses espousing the notion that general cognitive benefits and deficits accrue from the learning of second languages have been proffered and countered (see Carey, 1991a, 1991b; Cummins, 1984, 1991; Diaz, 1985; Edelsky et al., 1983; Reynolds, 1991; Rivera, 1984). It is interesting to note with regard to these hypotheses that results obtained in studies examining their validity differ in conclusions according to national boundaries, with Canadian studies tending towards the cognitive benefit side, while American studies lean more towards cognitive deficits (see Carey, 1987, 1991a, 1991b). This observation heightens suspicion that political expediency may be biasing either research questions or

---

[1] Please note: definitions for all starred (*) words are available in the glossary.

interpretation of the results or both[2]. There is certainly no doubt that social factors play a significant role in the development of the phenomena defined as additive* and subtractive* bilingualism, and therefore social factors must in corollary fashion be considered as contributing to the reputed cognitive benefits or deficits incurred by students in bilingual language learning situations.

For second language educators, this constitutes a disturbing conundrum. We ask ourselves: What is there in the nature of cognition*, often considered to be a discrete and central mental capacity, that makes it so susceptible to sociological influences? Can we neutralise such influences or at least employ them for the benefit of our students? Is it possible to reach into the core of the learning experience and come up with some insights which would offer all students, whatever their socioeconomic status, a good learning environment? Couched just below the surface of these questions, however, is one which is the touchstone for all of them: Are attempts to link varying second language performance to general cognitive advantages or disadvantages valid? Though this may seem like a bold question to some, the asking of it is not unprecedented. Allan Reynolds poses essentially the same query in his analysis of the cognitive consequences of bilingualism (1991, p. 159).

---

[2] The different political contexts in the United States and Canada is discussed more fully in section 2.1.

The problem appears rooted in an inability to clearly define cognition, its processes and its manner of functioning despite the fact that it has been the subject of enormous research effort. The processes of cognition are often taken to be quite simply those involved in thinking and learning as opposed to "conation (the action of the will) and affect (the actions of the emotions)" (Sanford, 1985, p. 1), and cognition itself is rarely defined in literature employing the term. Indeed, on being asked how cognition was defined in his recent book Schools for Thought (1993), John Bruer, an eminent cognitive psychologist interested in applying psychological theory to educational issues, replied "in the usual way,"[3] and once again, no formal definition is available in the text. But does everyone who reads Bruer's text understand what he means by the terms "cognition" and "cognitive" functioning? Is it not more reasonable to assume that everyone understands these concepts according to their level of expertise in cognitive theory, and if this is so, are we all talking about the same phenomena?

For anyone interested in finding out the "usual" meaning of these terms, a check into some psychological dictionaries and general texts is both fruitful and surprising. In many psychological texts, cognition is not defined. Even The Encyclopaedic Dictionary of Psychology (Harré & Lamb, 1983) has no separate entry for cognition, and refers to its meaning in a very cursory manner under

[3] I put this question to the author at the West Coast Attention conference in Eugene, Oregon in the spring of 1993.

the reference for "cognitive complexity" as follows: *cognitive complexity*: Reflects a style of thinking (cognition) ... (p. 91)

In those texts which do give a definition, the range of mental phenomena which are grouped under cognition is astonishingly broad, a fact which is brought to our attention by Reber in The Penguin Dictionary of Psychology (1985):

> *cognition*: A broad (almost unspecifiably so) term which has been traditionally used to refer to such activities as thinking, conceiving, reasoning, etc. Most psychologists have used it to refer to any class of mental "behaviors" (using that term very loosely) where the underlying characteristics are of an abstract nature and involve symbolizing, insight, expectancy, complex rule use, imagery, belief, intentionality, problem-solving, and so forth. (p. 129)

It is also interesting to note that the list of phenomena grouped under cognition by Anderson in the Concise Encyclopedia of Psychology (Corsini, 1987) has little overlap with that given by Reber:

> Cognition is a general term or a "generic term used to designate all processes involved in knowing," according to Ernest R. Hilgard cognition comprises all mental activity or states involved in knowing and the mind's functioning, and includes perception, attention, memory, imagery, language

functions, developmental processes, problem solving, and the area of artificial intelligence [sic]. (p. 202)

Symbolizing, insight, expectancy, ... *as well as* perception, attention, memory, imagery, ... ! Armies of neuroscientists and cognitive psychologists, not to mention artificial intelligence researchers, are presently carrying out studies in an effort to gain some understanding of the fine structure of these processes, and they themselves are reluctant to give definitions of some of the phenomena they are investigating due to the incompleteness of their knowledge of influential factors! It seems a mite cavalier on the part of researchers, therefore, to be employing terms like "cognitive benefits" and "cognitive deficits" in descriptive and prescriptive ways. Labels of this type can, unfortunately for the "deficit"-labelled, become self-fulfilling prophesies, in spite of the fact that, given the state of inquiry into the nature of cognition, none of the definitions to date can be wholeheartedly endorsed as being biologically well-founded.

Where does one turn to, then, for a coherent and informative view of cognition and its underlying processes? My quest for an answer to this question led me to the writings of Patricia Churchland, a self-styled neurophilosopher* who questions, as a result of her studies in the field of neuroscience and her discussions with countless researchers from many disciplines whose areas of expertise overlap with that of "cognitive" researchers, whether an "*empirically justified* distinction between the cognitive and the non-cognitive will survive, or if it does, whether it will confirm even a rough approximation of our current hunches" (Churchland, 1986, pp. 151-152). On this view, then, the designations of some processes as cognitive and others as non-cognitive may prove completely groundless from a brain-orientated perspective, just as recent research on the role of emotion in processes such as learning, reasoning, and decision-making has shown carving off "emotional processes" from cognitive ones to be (Damasio, 1994). Churchland believes that to get a handle on cognition and its manner of functioning one must start with the study of some of the phenomena considered fundamental to it, learning and memory (ibid.). I am in agreement with her. My contention is that little will be gained by approaching the problem of what is going on in the brains of second language learners from the level of general cognitive effects on the grounds that the degree of resolution of this level is too low to be of use as an analytical tool. A more incisive instrument for elucidating the issues under study can be found in connectionism*, an approach to the study of cognition which focuses on understanding the plasticity or learning capability of the brain, human as well as non-human. To apply a neurobiologically based perspective on learning to second language learning issues is not so radical a proposal as might be supposed: Jacobs and Schumann (1992) (see also Jacobs, 1988) suggested exactly this approach towards the study of second language acquisition, and many connectionist machine models of language learning have already been constructed. Connectionism, after all, is in essence an information processing methodology of the type espoused by Cummins in his clarification of the nature of language proficiency (1984a, 1984b), and, of course, by McLaughlin and his group (McLaughlin, Rossman, & MacLeod, 1983). What is different, and, in my opinion, superior about connectionism is that it provides a more detailed accounting of the mechanisms by which the brain processes information, an accounting which Edelsky et al. (1983) have identified as missing from Cummins' theoretical explanation of the literature, and the lack of which Cummins himself appears to regret in one of his most recent papers touching on the Threshold Hypothesis (Cummins, 1991, p. 86). My intent, then, is to offer an alternate explanation of the data which gave rise to the formulation of the Threshold Hypothesis as well as of data which dispute or support its premises, one which is based on a connectionist approach to information processing in the brain. And since there appears to be consensus among second language researchers that there is a need for a theoretical perspective which will shed some light on the interaction of social and linguistic variables in the determination of academic achievement, the explanation offered will attempt to clarify how a non-linguistic factor such as the social status of a language might "get inside the head" of a learner and thereby influence how well a language is learned.

Methodology

The problem being addressed in this study is the need identified by certain second language researchers for a more acceptable hypothesis on the cognitive consequences of second language learning in immersion situations than the Threshold Hypothesis proposed by Cummins, and the goal of the study is to provide such an hypothesis based on a connectionist approach to learning.

The major questions formulated at the beginning of the study are as follows:

1. Why did Cummins propose the Threshold Hypothesis?
2. What are the major criticisms aimed at the Threshold Hypothesis and how valid are these?
3. What model of cognition underpins the Threshold Hypothesis and how adequate is this model as a representation of how the brain functions?
4. Why was a connectionist approach to the study of the mind/brain developed?
5. What does connectionism propose about learning and learning processes in the brain, and how valid are connectionist models of cognitive function as these are manifested in models of learning?
6. What would a connectionist hypothesis propose with respect to the cognitive consequences of learning a second language under immersion situations?
7. What evidence is there to support this new hypothesis in

the language research literature?

8. How does this new hypothesis deal with the concerns raised by critics with respect to the adequacy of the explanatory power of the Threshold Hypothesis?

9. What are the implications of this hypothesis for the teaching of students in second language immersion situations?

As the investigation progressed, it became evident that two other questions were highly relevant to the study: What is the nature of attention, and how does it function? The importance of these questions arose out of the importance of attention to the learning process as it is described in connectionism, and it is dealt with in the section discussing the validity of connectionism.

Because it was clear from criticisms of the methodologies of the studies which made up the data base for the formulation of Cummins' Threshold Hypothesis that all of the empirical studies on the cognitive consequences of bilingualism were inadequate in one way or another (see section 1.3.4 and Reynolds, 1991), and suffered from a simplistic application of natural science investigation methods to the complexity of human cognitive functioning, a different orientation to the study of cognitive function, one which would be more sensitive to the ecology of brain function, was sought. The neurocomputational (or connectionist) approach was considered promising in this respect since it is constrained by biological data from all levels of investigation of brain function, and incorporates the notion of the dependence of cognition on environmental circumstances into its very definition (see glossary). Due to the difficulties mentioned previously in designing valid empirical studies relating cognitive consequences to bilingualism, a new empirical study based on past theorizing was thought to be redundant. Before carrying out more empirical studies, therefore, it seemed important to employ a different, more empirically valid data base to develop a new hypothesis on the cognitive consequences of bilingualism which would give a better accounting of the data already available on the cognitive consequences of bilingualism, and would itself be more readily testable empirically. Thus, the research design chosen was non-empirical and comprised the following steps:

1) a detailed meta-analysis of analyses of Cummins' work on the cognitive consequences of bilingualism (with special emphasis given to the Threshold Hypothesis) with the goal of establishing specific concerns as to the validity of this body of work.

2) a detailed analysis of connectionism with the goal of establishing the validity of its premises on how learning occurs in the brain.

3) a synthesis of data from several disciplines on learning in general, and language learning in particular, consistent with a connectionist approach to learning into an hypothesis of the cognitive consequences of second language learning in

immersion situations.

4) a reinterpretation of studies on the cognitive and/or academic consequences of bilingualism based on the new hypothesis to see if support for its premises could be found in the existing literature.

The data chosen for analysis were those considered most pertinent to answering the guiding questions of the study detailed above. Hence, the major disciplines researched were: a) second language learning, especially with regards to its cognitive consequences; b) neuroscientific data on brain function in general, and learning in particular; c) connectionist modelling of learning; d) attention research; and e) psychological studies of language.

In general, these disciplines are regarded as parallel strands of knowledge rather than ones which overlap in their areas of expertise. Perhaps the most unique aspect of this study, therefore, lies in its attempt to weave them together in order to construct a strong and broadly based hypothesis of the cognitive consequences of second language learning. Other original work to be found in this study are the following: i) a meta-analysis of analyses of Cummins' work on the cognitive consequences of bilingualism; ii) an analysis of the validity of connectionism as an approach to the study of learning; iii) the development of a working definition of attention; iv) speculation on how attention is represented implicitly in connectionist models of learning; and v) a reinterpretation of three major sets of data on the cognitive and/or academic consequences of bilingualism, those of Cummins, Collier (1987), and Diaz (1985).

The choice of data and the design of the study were undoubtedly influenced by my early background in the physical sciences which are heavily dependent on analysis. No other approach to the study of learning was considered. My experiences as a foreign language teacher and learner, as well as experiences of living in foreign cultures, however, have been influential in opening my mind to acceptance of the validity of studies based on the methods of naturalistic inquiry. It is interesting to note that, in keeping with the connectionist approach, the Threshold Hypothesis proposed by Cummins was also born out of the scientific research tradition as it is applied in psychological studies. The differences that emerge between the latter hypothesis and the one developed in this study do not stem from differences in research tradition, therefore, but from differences in choices as to which data are considered relevant to the problem investigated, and in the interpretation of certain language research studies.

Because the interaction of attention and learning is of particular interest to me, it may appear that the influence of attention on learning has been over-emphasized. Though I do believe that attention is of the utmost importance to the learning process, there is no doubt in my mind that other factors such as emotion and motivation to learn are also extremely important. Indeed, in the discussion of the nature and functioning of attention, these phenomena are perceived to be determiners of where attention is allocated and thus of what is learned and how well it is learned.

Though this study discusses the implications of a connectionist hypothesis for second language teaching, it does not deal in any depth with first or second language teaching practice. The purpose of discussing the implications of the new hypothesis for second language teaching is merely to indicate how changes in cognitive strategy brought about by immersion second language learning might support certain teaching strategies.

The most important prediction of the connectionist hypothesis of the cognitive consequences of second language learning in immersion situations is that, because language learning is a multimodal sensory process heavily dependent on audition and vision, second language learners in immersion situations will undergo a shift in information processing strategy which will result in an increase in attention allocation to the visual system and a corresponding reduction in attention allocation to the auditory system during the period of low speech comprehension in the target language. I have named this hypothesis the Shift in Information Processing Strategy Hypothesis.

Throughout this dissertation, an effort has been made to maintain consistency of terminology. In the case of terms such as theory, model, hypothesis, paradigm, and heuristic, definitions have been provided in the glossary. On a less formal level, a *theory* is taken as an explanation for a phenomenon. A *model* is a manner of depicting the variables, mechanisms, and constructs of a theory and their interrelationships. In most instances, a model is intended to represent a momentary static view of a dynamic process. When dealing with connectionist models of learning, however, a model often refers to actual machine models. It is hoped that which meaning of model is being used is clear at all times. The usefulness of theories and models depends in large part on their ability to generate predictions or *hypotheses* with respect to the relationships between variables. Thus hypotheses are drawn from theory and should be testable. Evidence which supports any hypothesis creates confidence in the underlying theory. An *heuristic* is also drawn from theory and involves predictions based on theory, but is seen as less structured than an hypothesis. The stricture of testability does not apply to an heuristic whose function is that of stimulating discussion on a topic. The term *paradigm* is used in its broadest meaning of "disciplinary matrix" (see Hoyningen-Huene, 1993, p. 141); that is to say, it refers to everything subject to professional consensus in a given research community. In this sense, it encompasses theory as well as its own narrower meaning of an exemplary problem solution. *Algorithm* is employed strictly in its mathematical sense of a method of solving a problem involving a finite series of steps. An algorithm is therefore based on mathematical theory.

Other terms employed in specific ways throughout this dissertation are listed in the glossary.

General plan of the dissertation

Chapter one examines Cummins' Threshold Hypothesis, detailing first how it came about, whether it is correctly classified as an hypothesis, and investigating the validity of critiques levelled at its premises. The aim here is to display the complexity of the concepts behind the hypothesis as well as the variety and extent of the criticism it has received at the hands of other second language researchers, and also to identify specific nodes of concern among critics that any alternate theory of second language learning under immersion conditions must try to resolve. Since the concept of cognition implicit to the hypothesis is of particular importance to this thesis, it is discussed in some depth.

In *chapter two*, connectionism and its premises are investigated and their validity examined in like manner. Following a historical perspective which sets the scene for an investigation of the premises of connectionism relevant to this discussion, its correct classification is sought, and then a detailed rendering of its biological underpinnings, particularly with respect to how the brain is organized and learns, is given. Since the attention system of the brain is seen to be of importance to learning in the brain and since its nature and functioning are not yet fully characterized, its nature is also examined in depth in order to identify nodes of consensus among researchers. Finally, the nature and function of connectionist machine models of learning is examined as a means of providing empirical support for the validity of the concepts upon which they were constructed, and a discussion of how attention might be represented in these models of learning is undertaken. The rationale for entering into great detail in the presentation of the underpinnings of connectionism and attention is, firstly, to make clear the multiplicity of levels at which any aspect of brain functioning can be studied and still be relevant to learning as it is understood by educators, and secondly, to develop an adequate conceptual basis for the construction of a connectionist model of second language learning in immersion situations.

Chapter three is devoted to the development of a connectionist hypothesis of second language learning in immersion situations based on the premises set out in chapter two. Immersion situations are focused upon specifically due to the fact that the controversy surrounding the cognitive consequences of bilingualism centres around immersion learning situations of various sorts. The hypothesis offered is founded on a general model of learning in the brain, and models of first and second language learning, all of which are elaborated as a prelude to the development of the hypothesis. This hypothesis is then employed to show how its explanatory power can give greater insight into the differential academic achievement results for minority* and majority* group second language learners, and finally its relevance and implications for second language education are discussed.

The concluding chapter, *chapter four*, re-examines Cummins' Threshold Hypothesis on the basis of this connectionist vision of second language learning in immersion situations with the intent of elucidating three areas of concern that were identified in the original analysis of the literature on the Threshold Hypothesis, the model of cognition underpinning it, "deficit" hypothesis charges, and the reality and nature of linguistic thresholds.

Chapter 1. A CRITIQUE OF CUMMINS' THRESHOLD HYPOTHESIS

1.1 Historical Perspective

From time immemorial, the learning of a language other than the mother tongue of the learner has proved to be a daunting task, and, to this day, it remains one of the most difficult feats of learning that can be accomplished. This is supposing, of course, that the language is to be well-learned; that is, learned such that the learner can use it, not only as an effective means of oral communication in the multiplicity of different communicative situations that one encounters within a society, but also read and write it to a level adequate to achieve well academically in an educational culture which employs the language as its primary mode of instruction.

In many parts of the world where multilingual communities can be encountered, there is no necessity for this kind of in-depth knowledge of more than one language. In the past, particularly in colonial times, the custom of the ruling power was to impose an "official language" as *the* language of instruction while other languages, though indigenous and therefore more commonly employed in everyday life, were relegated to second class status as academic vehicles. Such was the case in Hong Kong, where English was decreed the language of instruction though Cantonese was the language most commonly spoken by the people. There, however, as elsewhere in the developing world, policies are rapidly changing and, in general, greater status is now being accorded to indigenous languages, a reflection of the upsurge of pride in native language and culture that swept the planet in the sixties and seventies. Today, that enthusiasm has somewhat abated with the realization on the part of the developing countries that the ability to speak at least one of the international languages of power is an absolute necessity for economic survival in a world that is becoming increasingly interdependent economically. (For recent updates on the global situation see Carey, 1993a, 1993b.)

In North America, the situation developed somewhat differently. Essentially, the colonizing powers dismissed the languages of the First Nations as being unimportant on a national scale, and proceeded on the assumption that the language of the colonizing nation or, in the case of Canada, nations, would become the only legitimate language(s) for the purposes of carrying out the country's business. This attitude, which mirrored those of the countries in Europe from which they hailed, was fairly standard throughout the history of the developed world, and it is interesting to note that, in many of the countries in which "colonizing" languages originated, there is no official language policy since it was never deemed necessary to institute one. It was incumbent on immigrants to the "New World", then, to learn the majority language of their adopted country in order to be assured of acquiring employment and generally fitting in to the society, and, more important still, to ensure that their children were brought up speaking the majority language. Although it was accepted that parents might want to teach their children their heritage language, this was not encouraged in general as it was perceived to be a major source of potential academic problems for the child (see Cummins, 1984a, chap. 5, for a comprehensive review of the relationships that were drawn between academic achievement and bilingualism). The attitude, both educational and societal, towards speakers of "foreign languages", therefore, was unabashedly one of linguistic and cultural assimilation (see Taylor, 1991, pp. 2-5).

In spite of the fact that many of the immigrants came from nations with long and impressive histories, they appear to have been generally viewed as "wanting". If nothing else, they were wanting knowledge of the language of their new country, and of its culture. But they also tended overwhelmingly to be wanting financially, having often emerged from the lower echelons of the socioeconomic scale in their homeland, and, as a direct result of their low societal status, they were usually wanting in education. This situation has changed in more recent times, immigrants with university degrees and considerable wealth now making up a significant proportion of the immigrant population.

The policies which determined how immigrant children, and later, children of language minority communities, would be treated within the school system, however, were based on studies carried out in the first half of the twentieth century, many of which were methodologically flawed, and incorporated the built in biases detailed above of educators of the time (see Cziko, 1992, and Reynolds, 1991, for recent reviews of some of the problems that still plague the evaluation of bilingual education). The level of academic achievement of these groups of children was consistently found to be lower than the monolingual norm, a situation which was said to be the consequence of their lack of knowledge of the *lingua communis*, and even more generally, of bilingualism itself. Bilingualism, it was said, could lead to disadvantages in, for example, "intellectual development", "educational progress", and "personal development" (Cummins, 1984a, pp. 102-103). As Cummins brings out admirably, some of these statements were extreme. He cites Jensen's 1962 review of more than 200 studies to illustrate the extent to which the rhetoric had become overblown: "He [the bilingual] may become schizophrenic, for most bilinguals, according to Christophersen ..., feel a 'pull in opposite directions which threatens the unity of their personality'" (Jensen, 1962, p. 136).

Though most of these early studies have since been discredited, they were powerful in that they were consistent with the pervasive perception of second language learners in society at large, a perception which was based on first hand observation. As pointed out by Carey (1991a, pp. 340-341), even today "immigrant populations and other minority ethnic groups are usually academically disadvantaged while they are immersed in mastering their second language," a state of affairs with which "the average person is all too familiar." What was insidious about this situation, however, was that the poor academic achievement of these groups was equated with poor cognitive ability. This led to an inordinate representation of minority language students in "learning disabled" classrooms (Cummins, 1984a, pp. 1-2), an inequity which had not been adequately addressed in the United States as recently as 1983 when Oritz and Yates found Hispanic students in Texas overrepresented by a factor of 300 per cent in the learning disabilities category.

Though some more reflective educators called attention to anomalies in this relationship between bilingualism and low academic achievement, to case studies of highly successful groups (see Hamers, 1991, pp. 127-128), for example, or to the higher than normal academic achievement of whole sub-cultures such as the Asian-American and Asian-Canadian communities (see Cummins, 1984a, pp. 96-100), in democratic countries such as Canada and the United States where equality of opportunity was advocated and thought to be already in operation within the education systems, the causal link between deficiency in language ability and poor academic achievement must have seemed the only logical explanation.

In spite of the overall similarities in the policies of these two countries, however, there were important differences due to the fact that Canada had been colonized by two nations, the British and the French (see Carey, 1991a, for a recent update on the development of governmental policies in the United States and Canada). Though English eventually became the dominant language in Canada in terms of proliferation and political power, where the Francophone population was extensive enough, communities had education in their native tongue, French, with English being taught as a second language, while the English speaking community received second language instruction in French. Other linguistic minorities were taught in one of the two official languages largely according to geographical location, and received instruction in the other as a second language. To say this is, of course, to paint the picture only in very broad strokes. No society is homogeneous enough to avoid exceptions to the general trends, children who were fluently bilingual in French and English before starting school, for instance, or children who were allophones. How these were dealt with is unclear, but it seems likely that they were treated according to either their parents' wishes or societal norms. Even today, the mother tongue of a child is considered to be that first spoken at home, or worse still, the first language of their parents, in spite of the fact that many of these children either do not speak that language at all or are more fluent in the majority language of the community due to the prevalence of the latter in their environment (Bain and Yu, 1987; Carey, 1991a). Carey appears to consider this stance an infringement of the rights of minority children. Speaking of the situation of Francophone children in Canada he states:

> [T]he [Canadian] Constitution and article 23 of the Charter of Rights limit the eligibility of official minority children for these minority language programs to those for whom a parent's first language was French. This bases these minority education initiatives not on the linguistic ability of the children but on their parents' childhood language experience. In this sense, the criterion for admissibility is political and an affirmative action to reduce or reverse linguistic assimilation and cultural erosion. (Carey, 1991a, pp. 345, 346)

Now as in the past, therefore, political imperatives are firmly in control of linguistic policy.

Though the type of formalised instruction given at school at the time rarely led to a high level of bilingualism, it was far more common to find what would now be termed "dominant" bilinguals , those with a higher ability in one language (normally their first) but who were reasonably fluent in the second (Lambert, 1990, pp. 202-203), in the Francophone communities, particularly outside Quebec, than in the Anglophone communities, even those inside Quebec. The reasons for this were evident: the preferential rules set up to protect the rights of English speakers in the society, and the necessity of learning English for economic survival[1]. The English language community was, after all, far more numerous and widely-sown throughout the provinces and, not surprisingly therefore, had higher status. The weight of numbers meant little in Quebec, however, a predominantly French speaking province, in which English was the language of power until the 1960s, at which time the rising tide of the Quebec sovereignty movement began to exert its influence on language policy making. Once again, this was, and still is, a case in which political imperatives dominate the language scene. The upshot of this highly political situation with respect to language was that Francophone educators were very familiar with the academic and "cognitive" effects of bilingualism, having had much more extensive experience than their English speaking counterparts with all levels of bilingual ability. Small wonder, then, that the first serious challenge to the prevailing logic on these matters came from this quarter (see Reynolds, 1991, p. 147).

Wallace Lambert was a prime mover in the effort to make sense of the cognitive effects, if any, of bilingualism. In a recent review

[1] It is interesting to note that this relating of bilingualism to economic advantages has taken on an intriguing twist. Taylor (1991) reports that, in a recent survey, speakers of "heritage languages" in Detroit were still of the opinion that "bilingualism was associated with economic advancement in terms of school performance and job opportunity. Thus, *maintaining the heritage language* [italics added] was linked directly to bread and butter issues, not only to the quality of social relationships" (p.10). Taylor notes that the reasons for this perceived link are unclear, but may be related to socio-economic pressures brought on by exclusion from the dominant group.

of the issues (Lambert, 1990), he comments:

> We [at McGill University] started with concerns about the measurement of proficiency in two languages, but because our research group comprised mainly wildly divergent thinkers, we continually drifted off the proficiency issue to what we saw as inextricably related matters, usually social-psychology ones.... What surprises me is that the issues persist, even though the research approaches evolve and more and more information is accumulated. (pp. 201-202)

In 1962, a landmark paper by Peal and Lambert questioned the validity of the findings to that date on the relationship between bilingualism and intelligence, a study which "has had a profound influence on the fields of psycholinguistics and bilingual education" (Reynolds, 1991, p. 146), and later, along with Tucker, Lambert established the first language immersion programme in the public sector in Canada, another highly influential study which became known as the St. Lambert experiment (Lambert and Tucker, 1972). It was initiated at the urging of

> a group of English-speaking parents in Quebec whose common concern was that their children become highly proficient in French ... Their concern about their children's abilities in French reflected the political and economic realities of their environment - that French was being increasingly emphasized as

> the language of work, and indeed, has since become legalized as the official *langue de travail* in Quebec. (Swain & Lapkin, 1982, p. 1)

The question Lambert and Tucker addressed was whether being bilingual is costly in terms of intellectual development. The answer they discovered was that it was not; the very select group of students which they chose to participate in this special programme were found, according to the instruments used to measure these capabilities, to have suffered no negative linguistic or intellectual consequences as a result of the fluent bilingualism which they developed:

> After five years, we are satisfied that the Experimental program has resulted in no native language or subject matter (i.e., arithmetic) deficit or retardation of any sort, nor is there any cognitive retardation attributable to participation in the program. In fact, the Experimental pupils appear to be able to read, write, speak, understand, and use English as competently as youngsters instructed in the conventional manner via English. During the same period of time and with no apparent personal or academic costs, the children have developed a competence in reading, writing, speaking, and understanding French that English pupils following a traditional French-as-a-Second-Language program for the same number of years could never match. (Lambert & Tucker, 1972,

p.152)

As Carey (1991a) points out, results such as these were more than unusual. Indeed, it was "a precedent in educational history to suggest that it is possible to have an education via a second or non-dominant language without encountering a cost academically in terms of the content being mastered" (p. 336). It is of interest that, at this time in Canadian history, the movement to promote bilingualism and biculturalism initiated by the federal government's concern over rising dissatisfaction with respect to linguistic and cultural status within the Francophone community was well underway, hence the conclusions of the study were undoubtedly "politically correct" for the era. Though critiques of the experiment questioned the generalizability of the findings of the immersion programme due to the unique characteristics of the student population, the strong parental support, the high socio-economic status of the group, the dedication of the teaching staff, research design problems, and the screening of the language ability of the participating students, to mention a few of the more serious objections (see MacNab, 1979), the mind-set of the language teaching community had been shaken to its core, and research in both Canada and the United States focused around testing the scope of the validity of these results. Commenting on the importance of Lambert's contribution to the study of the cognitive consequences of bilingualism, Reynolds (1991) notes that, "since the Peal and Lambert monograph, there have been over seven dozen studies, reviews, and book chapters on the effects of bilingualism on cognitive functioning" (p. 147).

The most telling criticisms of the immersion language programme advocates compared the educational circumstances of the immersion classes with those of immigrant or low socio-economic status minority group children whose parents spoke a heritage language at home. These children, it was argued, were immersed in the second language at school and spoke their mother tongue at home as did the "immersion" children, but were commonly found to be deficient in linguistic ability in both languages, and, in general, were not seen to be achieving academically as they progressed through the education system[2]. Surely the practice of immersion language education could not in and of itself explain both sets of results? In answer to these criticisms, Lambert suggested that the discrepancy might be due to the fact that, for these children, the first language was being gradually supplanted by a more socially prestigious second language (Cummins, 1984a, p. 106). It was proposed that in such a social circumstance, the decline of the first language could occur more rapidly than the acquisition of the second, an outcome which could lead to a general degradation of linguistic ability and poor academic achievement. In forwarding this explanation, Lambert was allowing that linguistic considerations could not tell the whole story; at a minimum, the

[2] It should be remembered, however, that certain sub-groups such as the Asian-Canadian and -American communities were an exception to this rule.

social status of the languages concerned had to be taken into account as an intervening variable in how well they were learned.

The protests of sceptics were thus ignored or explained away as research results supporting the notion that bilingualism could constitute a positive force in children's academic development and could, moreover, bring about subtle cognitive advantages such as greater cognitive flexibility began to accumulate (see Cummins, 1984a, pp. 105-106; Hamers, 1991, pp. 130-131). Eventually, one advocate of the advantages of bilingualism "went so far as to claim that parents who didn't enrol their children in these new cognitively enriching programs which produced 'cognitive flexibility' should be liable to be charged with child neglect!" (Carey, 1993a, p. 5). It was in an effort to reconcile these opposing bodies of evidence, then, that Cummins first proposed the Threshold Hypothesis (Cummins, 1991):

> Initially (Cummins, 1976), I was concerned with the apparent contradiction between the findings of empirical studies that appeared to suggest cognitive deficits resulting from bilingualism and more recent studies that associated bilingualism with cognitive advantages. Based on the patterns of L1 [first language] and L2 [second language] development manifested by bilinguals in each of these two types of studies, I suggested that there may be threshold levels of linguistic competence that bilingual children must attain both

> in order to avoid cognitive disadvantages and allow the potentially beneficial aspects of becoming bilingual to influence their cognitive functioning. The possibility of lower and higher thresholds associated with cognitive disadvantages and advantages was later linked to Lambert's ... distinction between subtractive and additive types of bilingual acquisition contexts. (p. 76)

The Threshold Hypothesis, therefore, essentially accepted a causal link between bilingualism and cognitive functioning and attempted to make sense of the conflicting data on the basis of that assumption, though it was very evident from his own research that Cummins was acutely aware of the interference of sociological variables in that relationship. In his text Bilingualism and Special Education, Cummins (1984a) states, "it is a fallacy to regard psycho-educational factors as in any sense an explanation of the [minority] students' academic problems" (p. 92), but later goes on to elaborate that, because psycho-educational "child process" variables, which "mediate" the social and educational variables in academic achievement, "are also those that are of most direct concern to educators concerned with assessment and pedagogy of bilingual students," it is necessary to explore their nature "to better understand their interactions with social and educational variables and also to specify more precisely the theoretical generalizations that are relevant to educational policy-making in this area" (ibid., p. 127). He further clarifies his rationale for pursuing this particular psycho-educational path as follows:

> This analysis suggests that the levels of proficiency bilingual children attain in their two languages may be an important intervening variable mediating the effects of bilingualism on children's cognitive and academic development. Specifically, it has been hypothesised that there may be threshold levels of linguistic proficiency bilingual children must attain in order to avoid cognitive deficits and allow the potentially beneficial aspects of becoming bilingual to influence cognitive growth. The threshold hypothesis assumes that those aspects of bilingualism that might positively influence cognitive growth are unlikely to come into effect until children have attained a certain minimum or threshold level of proficiency in the second language. Similarly, if bilingual children attain only a very low level of proficiency in one or both of their languages, their long-term interaction with their academic environment through these languages, both in terms of input and output, is likely to be impoverished.
>
> The form of the threshold hypothesis that seems to be most consistent with the available data is that there are two thresholds ... The attainment of a lower threshold level of bilingual proficiency would be sufficient to avoid any negative effects; but the attainment of a second, higher level of bilingual proficiency might be necessary to lead to

> accelerated cognitive growth. (ibid., p. 107)

To this hypothesis was added a few years later the Interdependence Hypothesis. This addition incorporated the growing number of studies on bilingual education in both Canada and the United States which supported Lambert's contention that education through a minority language entailed no adverse effects for the development of the majority language. Expressed formally, the Interdependence Hypothesis states:

> To the extent that instruction in Lx is effective in promoting proficiency in Lx, transfer of this proficiency to Ly will occur provided there is adequate exposure to Ly (either in school or environment) and adequate motivation to learn Ly. (Cummins, 1987, p. 199)

The implication of the Interdependence Hypothesis is that there is a common underlying language proficiency upon which both the first and second languages can draw. This hypothesis, then, along with the Threshold Hypothesis became the foundations for the subsequent theoretical edifices constructed by Cummins in the following years.

1.2 An Examination of the Validity of the Classification of the Threshold Hypothesis as an Hypothesis

In the Threshold Hypothesis, it is proposed that there is a direct, causal relationship between the degree of bilingualism and the cognitive ability of the learner, the direction of the causality being from the degree of bilingualism to cognitive ability. Below a certain minimum threshold level of bilingualism, the learner is expected to experience negative cognitive effects when compared to the unilingual norm, while above that threshold, a period of development will be entered into during which neither positive nor negative cognitive benefits will be accrued. As the level of bilingualism increases, a second upper threshold of proficiency will be reached beyond which cognitive benefits as compared to the unilingual norm will begin to manifest themselves in the learner.

In his delineation of the hypothesis, however, Cummins gives no definition of either of the two main variables, cognitive ability and degree of bilingualism, nor does he define the features of a threshold level of bilingualism. From some of his statements (Cummins, 1984a, p. 105), cognitive capability can be construed as a conglomerate of "academic and cognitive skills," "cognitive flexibility," "divergent or 'creative' thinking abilities," and "general intellectual development." It seems unlikely that such a diverse and loosely related group of abilities could be represented accurately by a single global construct labelled "cognitive ability," especially when the measures of cognitive variables used in the studies the results of which prompted Cummins to propose his hypothesis were equally disparate, comprising various subtests out of various intelligence test "kits." In a study purporting to support the Threshold hypothesis, for example, Cummins (1977) himself used the following measures: the Kuhlmann Finch Intelligence Test as indicators of verbal ability and general reasoning, the Utility Test from the French, Ekstrom and Price (1963) Kit of Reference Tests for Cognitive Factors scored for fluency, flexibility and originality, and an adaptation of Wallach and Kogan's (1965) Patterns test as a measure of non-verbal intelligence. No rationale was put forward as to why these particular subtests were considered more appropriate than others. That Cummins appears to be having difficulty deciding on which measures to use to reflect cognitive ability is hardly surprising, however, given the very broad range of phenomena grouped under the definition of cognition (see Introduction).

As far as degree of bilingualism is concerned, no indication is given by Cummins in his elaboration of the hypothesis as to how this should be measured, whether it would be dependent on verbal or written fluency, reading ability, conversational skills, or perhaps a medley of all of these of the type seen in standardized language tests. Even if the latter suggestion were used, there is reason to doubt that such an assessment offers an accurate measure of language proficiency (Edelsky et al., 1983). The measures Cummins used in his 1977 study, which were adapted from those pioneered by Peal and Lambert, were directed at assessing the degree of balance shown by the student but could ostensibly be employed to assess the degree of bilingualism. The reliability of these tests as accurate measures of balance, however, can also be called into question since they are dependent on a) a word association test which appears to compare only stored lexical knowledge in both languages, b) a subjective self-rating test, and c) a teacher rating of balance. A composite balance score was formed by combining scores on the three measures of balance. Although it can be seen that an attempt is being made at obtaining a comprehensive and reliable measure of degree of bilingualism, the nature of the ingredients that make up the composite balance score leave much to be desired in terms of conceptual clarity and reproducibility. These considerations combined with the nebulous notion of a "threshold" level of bilingualism, which is both defined and proven by the onset of certain levels of cognitive effects and is thus a tautology (see section 1.3.2), would render the Threshold Hypothesis extremely difficult if not impossible to test empirically.

In this penury of definition of its constructs, then, the Threshold Hypothesis can be seen to be lacking the conceptual clarity required for it to be termed an "hypothesis" taken in its strong sense, a provisional supposition from which to draw conclusions that are in accordance with known facts, and which serve as a starting-point for further investigation by which it may be disproved and an enduring theory arrived at. Cummins himself appears to have reached this conclusion in recent times (Cummins, 1991, pp. 83 and 85):

> Of the three theoretical constructs, the threshold hypothesis is clearly the most speculative and also the most difficult to test in any precise manner. Initially, it was intended to serve a *heuristic* [italics added] function in highlighting the possible relationships between patterns of bilingual development and cognitive consequences.... In view of the difficulties in conceptualizing and operationalizing the nature of "language proficiency" itself in any precise way ..., it is probably unrealistic to expect significant progress in elucidating the precise nature of the hypothesized thresholds. Thus, the positing of specific thresholds must be regarded, at this point, as primarily a *heuristic* [italics added] to orient research to explore the relationship between aspects of bilingual students' emerging proficiency in two languages and the possible cognitive implications of different patterns. Since the nature of the hypothesized thresholds are not specified, neither version of the threshold hypothesis [Cummins' or Diaz' (1985)] can be falsified at this level.

In referring to the "heuristic function" of the Threshold Hypothesis, which he does in several of his publications including those in which the hypothesis is first proposed, and in concluding that it is "primarily a heuristic," Cummins is approaching a more reasonable classification of his "hypothesis." Given that an heuristic may be defined as an argument intended to stimulate interest as a means of furthering investigation, to classify the Threshold Hypothesis as an heuristic would seem appropriate, since it has indeed encouraged a great deal of discussion and research interest within the second language research community, a fact to which the following section will attest.

To suggest that the classification of the Threshold Hypothesis should be downgraded to that of an heuristic is not so serious a matter as might be supposed, however, since the particular classification given to a theoretical proposition is dependent on the level of inquiry at which falsifying information is sought. What might qualify as a testable hypothesis at the level of general psychological inquiry, for example, might be considered hopelessly imprecise at the neuropsychological. Thus, theoretical classification is strongly related to the degree of generalization sought. It could be said that the degree of generalization sought by the Threshold Hypothesis is too great, and that its lack of testability is a direct consequence of its overly general conceptualization. That is, it is too susceptible to deconstructive critiques. In consideration of this difficulty (which dogs all theoretical propositions), therefore, the term "hypothesis" will continue to be applied to the Threshold Hypothesis on the understanding that, rather than a "test hypothesis," it is to be conceived as more of a "working hypothesis," one which serves "to guide and organize the investigation, providing us something to go on with" (Kaplan, 1964, p.88), since this term is closer in meaning to that of an heuristic.

1.3 Review of the Critiques of Cummins' Threshold Hypothesis

1.3.1 Introduction

As Cummins himself has noted, the Threshold Hypothesis and the theoretical constructs built thereupon have been "both influential and controversial" (Cummins, 1991, p. 75). The natural consequence of this has been critical assessment of the work from both advocates and opponents. Very few of the critics have treated only one aspect of the overall theory, such as the Threshold Hypothesis, as a separate entity, preferring instead to see all of Cummins' work as a related whole, undoubtedly a valid perspective. Troike (1984, p. 44), for example, notes that Cummins along with Toukomaa and Skutnabb-Kangas

> have suggested that *lower* and *higher thresholds* of proficiency should be recognized, with the lower threshold being minimally adequate for non-academic purposes, and the higher threshold being a precondition for academic success. The lower threshold would presumably have some similarity to what Cummins ... has since called 'basic interpersonal communication skills' (BICS).

The upper threshold, then, appears closely related to a postulated "'cognitive/academic language proficiency' (CALP), differing from other aspects of language proficiency, which forms the basis for academic achievement" (ibid., p.44). As a result of critiques of these constructs, Cummins (1984b) subsequently proposed that

> "language proficiency" can be conceptualised along two "continuums". The first is a continuum relating to the range of contextual support available for expressing or receiving meaning. The extremes of this continuum are described in terms of "context-embedded" versus "context-reduced" communication.... The vertical continuum is intended to address the developmental aspects of communicative proficiency in terms of the degree of active cognitive involvement in the task or activity. (pp. 12-13)

These continua are viewed as intersecting and are represented graphically as a grid.

It is important to recognize the interrelatedness of Cummins' theoretical constructs in a review of the critiques such as this one which will focus more narrowly on only the Threshold Hypothesis, because, as mentioned above, many of the critiques touching on the Threshold Hypothesis deal with the entirety of Cummins' proposals. It is precisely this interrelatedness of concepts that allows much of the criticism aimed at later constructs to be extrapolated back onto the original proposals which gave them birth. Throughout this review, therefore, criticism directed at the larger issues underpinning the later constructs will be considered as relevant to the Threshold Hypothesis. It should also be noted that, due to the enormous amount of research that the work has spawned, only the most important critiques, those referred to by Cummins himself and those judged by the author to be of significance, have been surveyed here. The intent, then, is to give comprehensive insight into the types of criticisms invoked by Cummins' theoretical stance on bilingualism and academic achievement, but not to conduct an all-inclusive synthesis of critiques of his theoretical constructs relating bilingualism to academic achievement.

In general, the criticisms fall into three basic categories: conceptual, philosophical, and analytical concerns. Though it is difficult at times to tease these apart, an attempt to do so will be made for the sake of clarity.

1.3.2 Conceptual concerns

Isolation of a psycholinguistic factor of language proficiency from sociolinguistic concerns

When outlining the rationale for his "new" theoretical framework for second language learning, Cummins (1979) noted that, in several papers

> the marked difference between the outcomes of immersion

> programs and "submersion"[1] programs for the minority child has usually been attributed to socio-cultural and attitudinal factors such as socioeconomic status (SES), community support for the school program, relative prestige of L1 and L2, teacher expectations, etc. Unlike earlier attempts to explain the poor academic achievement of many minority language children, little importance has generally been attributed to specifically linguistic explanatory factors. Bowen goes so far as to argue that linguistic factors are unimportant and that "the choice of language of instruction in our schools is linguistically irrelevant"...
> In contrast to this position, I shall propose a theoretical framework which assigns a central role to the *interaction* between socio-cultural, linguistic and school program factors in explaining the academic and cognitive development of bilingual children. (pp. 222-223)

In spite of his avowed intent to incorporate sociological factors, Cummins has been taken to task by several authors (Canale, 1984; Commins & Miramomtes, 1989; Edelsky et al., 1983; Genesee, 1984; MacNab, 1979; Martin-Jones & Romaine, 1986; Troike, 1984; Wald, 1984) for giving primacy in his theoretical constructs to the role

---

[1] The term "submersion" is used to indicate a language immersion program in which the first language of the students is neither used nor promoted in the classroom, a situation which is thought to lead to negative linguistic and cognitive effects.

of a putative[2] linguistic factor. The perspective of these authors is succinctly put forward by Genesee (1984, p.21) who protests:

> Cummins stresses that the type of language proficiency that is associated with school-related language use, is "socially grounded and could only develop within a matrix of human interaction". However, in subsequent discussions of the relationship between language proficiency and academic achievement, these social foundations are virtually ignored.

Genesee goes on to point out that if stress is placed on cognitive and linguistic dimensions of language proficiency as an important explanatory factor in accounting for disparate academic achievement outcomes, it will lead to cognitive and linguistic treatments of the problem. What disturbs these authors about the assumption of an intervening linguistic variable is first, a lack of adequate support for such a concept in the data (Edelsky et al., 1983; MacNab, 1979), and second, the direction of the causal chain that is thereby invoked (Diaz, 1985; MacNab, 1979; Martin-Jones & Romaine, 1986), both issues which will be dealt with in more depth at a later stage of the discussion. Martin-Jones and Romaine (1986, p. 31) see such an assumption as constituting "a deficit theory in the making," a view which is "most clearly expressed in Cummins'

---

[2] The adjective "putative" is used here to indicate that, in the minds of many researchers, especially those interested in sociological explanations of the facts, that linguistic link has not been proven.

threshold hypothesis."

Commentary

The basic charge made here of the ignoring of any explanatory power of sociological variables in Cummins' theoretical constructs, and in particular, in the Threshold Hypothesis does seem to be justified. The heuristic describes only two covarying dimensions, level of bilingualism and cognitive ability, neither of which includes, nor is intended to include, any measure of sociological interference. Cummins (1984a) characterizes them as "psycho-educational 'child process' variables that mediate the social and educational determinants of bilingual students' achievement" (p. 127), and appears to be attempting quite purposely to winnow out such variables from the sociological chaff. As MacNab (1979, pp. 241-244) and Reynolds (1991, pp. 148-162) have shown, considerable effort has been made by some researchers, especially those purporting to show positive correlations between the two dimensions, in the design of their experiments just so as to isolate the independent variable, degree of bilingualism, by careful matching of subjects on other relevant variables including sociological factors. With respect to the dependent variable, cognitive ability, similar care has been taken to ensure that the tests used are not biased against second language learners. That such efforts are ill-fated in the views of these authors will be discussed later.

What Cummins appears to be seeking is the isolation of a "brain processing" variable from other mediating variables, an effort perhaps to find an answer to the question: "What is going on *in there* as a result of what is going on *out here*?" It will be argued that, though a valid and laudable quest, his approach to the problem is lacking the resolving power to settle such a question.

Definition of constructs

(1) Cognitive issues

Critiques of Cummins' approach to cognition

Cummins has been faulted by several authors on a lack of clear-mindedness with respect to what might best be termed "cognitive issues." Edelsky et al. (1983, p. 7) complain that, "in his claims for evidence of cognitive advantage or disadvantage," he makes use of tests which "fail to illuminate the cognitive processes they claim to assess, focusing as they do on answers rather than processes," adding

> Almost none [of the tests] investigates cognitive functioning from a broad theoretical framework of cognition. Instead, investigations of static surface performances which are highly susceptible to variations in context substitute for investigations of such cognitive processes as memory, hypothesis generation, abstracting from and re-coding input,

etc. ...

There are therefore good reasons to suspect the data offered in support of these hypotheses.

MacNab (1979) also showed concern over what exactly was being measured as indicators of the dependent variable, "cognitive skills." He determined that in the post Peal and Lambert era, different instruments were being used, ones which included a greater degree of cognitive flexibility measurements, and concluded that "the measures can probably be categorized into divergent thinking measures and two kinds of convergent thinking measures (where the problem is to find the right answer) - verbal and non-verbal" (p. 237).

While agreeing in general with the spirit of the critique of Edelsky et al., Martin-Jones and Romaine (1986) focus primarily on the tightly drawn relationship which Cummins' theories presuppose between language and cognitive functioning, citing his statement that "'semilingualism'[3] `cannot be used as a strictly linguistic concept at all' ..., but relates specifically to cognitive development" (p. 28). They comment:

Cummins appears to be equating semantic development with

[3] It should be noted that the term "semilingualism" has been widely criticized (see, for example, Martin-Jones and Romaine, 1986) and is now outdated. It has been replaced by the less conceptually burdened term "subtractive bilingualism" proffered by Lambert.

> cognitive development in these definitions. The relationship between language and thought processes is by no means as straightforward as he suggests. Cognitive categories are not the unequivocal pace-setters for the acquisition of linguistic categories hoped for by some researchers who have done work on language acquisition. (p. 29)

Canale (1984) concurs noting

> one must not confuse linguistic demands with other cognitive demands made by a given task: To do so may contribute to incorrect conclusions about the language proficiency and even general cognitive proficiency of individuals (e.g. in the form of labels such as "linguistic deficit" and "cognitive deficit." (p. 37)

Just such a mislabelling of students with limited English and Spanish proficiency was discovered in the study by Commins and Miramontes (1989) which was critical of the constraining of the parameters of students' abilities by theoretical constructs such as Cummins' BICS/CALP dichotomy. In an ethnographic study, they found that teachers

> had underestimated [the students'] linguistic and academic abilities in both English and Spanish based on their classroom performance. Seemingly, teachers mistake a lack of vocabulary

> and a lack of verbal clarity for the absence of underlying thinking abilities.... Although students may not have demonstrated these skills [synthesis, analysis, and evaluation] in this classroom's normal activities, they did show evidence of thinking at complex levels when they were observed over a variety of contexts. (ibid., pp. 465-466)

What was lacking in the school context, according to these authors, was "the essential element of relevance that characterized the situations in which the students displayed their greatest competency" (p. 467). The worries of Canale and Commins and Miramomtes are echoed by Wald (1984) who, in explaining the work of Labov and Shuy, talks of "the falseness of associating nonstandard, often stigmatized, forms of language with cognitive deficiencies in their speakers" (p. 57).

Finally, MacNab (1979, p. 232) "takes issue with the proposition of Cummins (1978) that children who are exposed to a bilingual education but who do not achieve balance may be cognitively harmed by the experience," a concern shared by Diaz (1985), and one which relates directly to the Threshold Hypothesis. In an experiment designed to test this premise, Diaz found that, not only was it possible to attribute the observed cognitive differences to group differences in socioeconomic variables, but also that "the positive effects of bilingualism are probably related to the initial efforts required to understand and produce a second language rather than to increasingly higher levels of bilingual proficiency" (p. 1387). MacNab, on the other hand, argued that a more parsimonious explanation of the data would be that brighter children learn second languages faster and better and that, in this respect, the learning of a second language is no different from the learning of any other subject. The results of the study conducted by Diaz, however, who also investigated the causal direction of the relationship between "degree of bilingualism" and "cognitive ability," lend support to the premise of the Threshold Hypothesis that it is level of bilingualism which drives cognitive variance, and not the reverse.

Commentary

It seems clear that at the heart of all these concerns lies one burning question: How does Cummins define cognition? The answer, however, is very difficult to pin down, since, while he has written a great deal on issues related to the question, he has not addressed the question as posed in any of the papers reviewed. Nevertheless, in light of the critiques levelled at him by the authors above and the interest of this author in the topic of cognition, it seems only fair to examine Cummins' own statements on cognitive matters to try to synthesize, through extrapolation of concepts, a comprehensive and well-focused picture of his conceptualization of this important construct.

Cummins' conceptualization of cognition

On examination of several of Cummins' papers, particularly those dealing with the elaboration of the Threshold Hypothesis and his ideas on the nature of language proficiency, it was noted that the information on cognition fell into three general categories, answers to What develops?, How does it function?, and How can it be measured?, all questions that psychologists have pursued assiduously for many years. These three categories being somewhat independent, they will initially be dealt with separately, and a synthesis of the underlying concepts will be presented at the end of the section.

What develops?

(i) Cognition itself

It is evident from his dissection of the components of language proficiency into BICS and CALP, "surface" and "deeper" levels (Cummins, 1984b, p. 138), that Cummins conceives of the various language processes necessary to the proficient use of language as hierarchical in nature, with the "deeper" levels being more cognitively demanding than the "surface" levels (e.g. pronunciation is less cognitively demanding than grammar, which is less cognitively demanding than functional meaning). It is equally evident by his juxtaposition of various cognitive processes in parallel to the language processes, that these are conceived of as separate from but interactive with language processes, and are similarly hierarchical in nature, increasing in cognitive demands with increasing abstraction (e.g. comprehension is seen as less cognitively demanding than analysis, which is less demanding than evaluation). These conclusions are ratified by statements made by Cummins when discussing the nature of the thresholds in the Threshold Hypothesis (Cummins, 1979):

> Thus, in the early grades the lower threshold may involve only a relatively low level of listening comprehension and expressive skills, but - as the curriculum content becomes more symbolic and requires more abstract formal operational thought processes - the children's "surface" L2 competence must be translated into deeper levels of "cognitive competence" in the language. (p. 231)

Fluency in a language, then, is no guarantee that the speaker is capable of succeeding on tests which require "complex cognitive operations to be carried out" (ibid., p. 231).

In keeping with these notions of the separateness yet relatedness of cognitive processes and language processes, Cummins has taken pains to distinguish "basic cognitive abilities" from cognitive abilities related to language skills and academic achievement:

> The fact that, in comparison to middle-class children, low SES minority language children may be more dependent on the school to provide the prerequisites for the acquisition of literacy skills does not imply that these children's basic cognitive

> abilities are in any sense deficient nor that their command of the linguistic system of their L1 is necessarily inadequate. (ibid., p. 240)

It is notable, though, that he does perceive of language, and in particular, bilingualism, as capable of having a profound effect on cognitive and academic abilities over the long term:

> However, in bilingual learning situations where the child fails to overcome difficulties in coping with two languages the research evidence suggests that his bilingual learning experiences might have a negative effect on his cognitive functioning, at least in so far as this functioning involves language. Continued difficulties with language over a prolonged period of time are likely to mean that a bilingual child's interaction with an increasingly symbolic environment will not optimally promote his cognitive and academic progress. (Cummins, 1976, p. 23)

The ambivalence of Cummins' position on the influence of language on cognition will be returned to at a later stage in the discussion.

Implicit in the last citation and explicitly stated in the Threshold Hypothesis is the notion that cognitive development can be accelerated or retarded by experience including language experience. By describing level of bilingualism as an "intervening variable" in accounting for cognitive growth, one which is itself dependent on "more fundamental social, attitudinal, educational and cognitive (e.g. language learning aptitude) factors" (Cummins, 1976, p. 23), Cummins appears to be attempting to diminish the importance of language factors in the face of sociocultural factors. Although it would seem that this is a fitting refutation of the criticisms of his giving primacy to the influence of language on cognitive growth, this impression too turns out to be contentious as will be seen later in the discussion.

The rough sketch of Cummins' view of cognition which is emerging to this point is that of an isolated core of abilities which have their own cohesive integrity ("basic cognitive abilities"), but a core which has extensions into specific areas of knowledge such as language (cognitive competence in a language). It would also appear from some of the above statements that the cognitive abilities developed in these specific areas of knowledge can exert a reciprocal effect on core or "basic" abilities. Given this notion of cognition as separate yet integrated, it is interesting to examine just what Cummins perceives cognition to be separate from. In the citation above (Cummins, 1976, p. 23), "cognitive (e.g. language learning aptitude) factors" are clearly separate from social, attitudinal, and educational factors, and appear to be equated with individual difference factors or "child input" (Cummins, 1979, p. 243) factors. These child input factors are seen as having "motivational aspects [which] are also likely to interact both with linguistic input variables and Educational Treatments" (p. 243) in determining the academic and cognitive outcomes of children learning a second language. It is evident, then, that according to this perspective, in the determination of academic and cognitive outcomes, there are two types of influences; ones that are particular to the child or, more generally, the learner and are internal, and external ones stemming from the environment which interact with these internal variables in such a manner as to accelerate or retard their growth. However, it is noteworthy that the former, which are most closely identified with cognitive factors, are considered to include attitudinal aspects such as motivation towards learning which, in the statement above, was considered separate from cognition. Attitudes, it would appear, have influence from both without and within, not an unreasonable position, since the attitudes expressed in our environment are inevitably incorporated into our thinking in one way or another.

In replying to criticisms that he placed too little emphasis on social variables as determinants of school success, Cummins adamantly contends that this is not the case: "In short, the causal primacy of sociopolitical factors is not in question; however, cognitive and linguistic factors are also of obvious relevance to consider as intervening variables in interaction with educational treatment" (Cummins, 1984c, p. 72). The problem is that, as Cummins (1979) himself notes, "[t]he 'threshold' and 'developmental interdependence' hypotheses attempt to integrate data which suggests that linguistic factors are *important* [italics added] in understanding the dynamics of the bilingual child's interaction with his educational environment" (p. 236). Just why and to what extent linguistic factors are important is precisely what these critics are questioning, an issue which will be taken up in more detail below.

One problem that raises its head here is that of how related Cummins considers cognition and academic achievement to be. Although on many occasions in his writings he separates the two concepts, often, as in his elaboration of the BICS/CALP dichotomy, he appears to lump the two together as closely tied concepts. For example, in the very naming of the construct cognitive/academic language proficiency, where the slash appears to indicate inter-relatedness. Given this blurring of the boundaries between the two concepts, then, it is interesting to note that the Threshold Hypothesis, an hypothesis aimed at explaining differential results in academic achievement for second language learners, proposes a relationship between level of bilingualism and cognitive abilities, and *not* bilingualism and academic achievement. A reasonable interpretation of this is would be that it is through the mediating effects of cognition that language interacts with the educational environment. But the directional arrow of the effect could ostensibly go in two directions: either cognition has a mediating effect on language ability as was suggested by MacNab (1979), or language ability alters cognitive processes, or perhaps both. Whatever the case may be, it is very apparent that Cummins does regard the two as interdependent. To what degree he perceives them to be so, and in which direction he perceives the causal arrow to point will be examined below.

(ii) Language as an influence on cognition

In the paper in which he discusses in greatest detail the theoretical scaffolding for the Threshold Hypothesis (Cummins, 1976), Cummins outlines the two major contrasting positions on the inter-relationship of language and thought, those of Piaget and Vygotsky. While the Piagetian position considers that the development of cognitive operations is essentially independent of language, Vygotsky's view holds not only that the development of logical thought is dependent on the internalization of speech, but that "the totality of an individual's personality is closely integrated with linguistic experience" (p. 28). It is immediately evident that the latter perspective offers no difficulties for the concepts of the Threshold Hypothesis, and, according to Cummins, even the Piagetian point of view can accommodate it. He argues that, though in the view of "Genevan psychologists" "linguistic experience... is capable of accelerating cognitive growth only to a limited extent" (p. 34) due to the fact that operational thinking derives from action not language, "language" in the context of their experiments usually refers to "specific short-term training procedures" (p. 34), not to language as it is normally employed in social situations. He sees the emphasis placed by these researchers on the generation of feedback derived from the child's own actions on the environment for the promotion of operational growth as holding the key to the effectiveness of bilingualism in channelling cognitive development. Bilingual training, in his view, constitutes a special case of action on the environment, and even of language action. Cummins contends that, during the language training period, the bilingual is

> constantly *generating* data, through his speech actions on the environment, which provide a qualitatively different form of feedback from that provided by a unilingual's speech activity. It is the feedback from these bilingual speech actions which can (according to the objectification hypothesis) accelerate cognitive growth. (p. 35)

This position, he believes, is "in no way incompatible with Piaget's ... statement that 'language does not seem to be the motor of operational evolution, but rather an instrument in the service of intelligence itself'" (p. 35). As an instrument in the service of intelligence, bilingualism "might lead to a faster separation of sound and meaning, thereby directing the bilingual child's attention both to the essential or conceptual attributes of objects...and to the characteristics of his two languages" (p. 32). The importance accorded to the directing of the child's attention is of interest as it is a point which will be taken up in great detail in a later chapter. For the moment, it is sufficient to wonder why this directing of attention to conceptual and linguistic characteristics of the two languages would not be equally effective in promoting intellectual growth at all levels of bilingual training.

Commentary

Although Cummins has gone to great pains to try to reconcile Piaget's position on cognitive development with his own, his arguments are not thoroughly convincing. First of all, the spirit of Piaget's statement above would seem to be that intelligence (cognition) operates on language rather than the other way around. It is clear from Cummins' comments, however, that, though he may agree with the first part of the statement, he believes there is a reciprocity of action from language onto cognition which is strong enough to cause cognitive retardation or acceleration. This is a far cry from the Genevan psychologists' position that it is "capable of accelerating cognitive growth only to a limited extent" (Cummins, 1976, p. 32), even taking into consideration the possibility that the definition of language employed is particular to the test situation. Secondly, if a bilingual's speech actions are generating data or feedback which bring about cognitive change, this would appear to constitute a motoric action: a motor, after all, is the instrument that generates the power to effect change. Furthermore, any action which does effect change, whether directly or indirectly, must be classified as causal to some degree. Thirdly, in spite of the fact that he does not openly espouse one theory of cognitive development over the other, Cummins does not disassociate himself from Vygotsky's position on the relationship between language and thought, and even cites Vygotsky's arguments on the possibility of bilingualism acting as a powerful tool in the development of cognition:

> In terms of Vygotsky's theory of cognitive development, it is inconceivable that bilingualism should not affect the child's development (either positively or negatively) since not only intellectual development but also "character development and emotional development all reflect the direct influence of speech". (ibid., p. 34)

Now, if one considers the importance of speech to the "more fundamental social, attitudinal [and] educational" variables which Cummins considers to be causal in reaching the lower threshold of bilingualism, it becomes apparent that there is nothing that is not influenced by language in his mind.

What emerges on analysis of Cummins' statements on the relationship between language and thought (or cognition) is that he considers language to be of the utmost importance in determining both cognitive and academic outcomes. Thus the critics who claim that he places too much emphasis on language in explaining academic success appear to be on solid ground. However, as noted above, language is indeed all-pervasive in society and cannot be effectively isolated from any of the so-called causal variables affecting cognitive growth as it is defined and measured. Hence it will undoubtedly have a significant effect on what students learn. This is not to say that Cummins is correct in his approach; but the problem appears to be far more complex than was suspected. For the moment, suffice it to say that a more productive avenue of inquiry might be to question the usefulness of attempting to isolate a linguistic variable given the omnipresence of language in the environment, and even to question whether cognitive processes can be accurately discerned at this level of investigation, a point which will be developed further later in the discussion.

How does cognition function?

In elucidating the manner in which cognitive functioning can be altered by language, Cummins describes language, especially in the bilingual experience, as generating data (feedback) which has the potential for accelerating cognitive growth. Data or feedback is, of course, information. Hence, it would appear that it is the additional information processing load that the brain has to accommodate which accelerates cognitive growth. That this is indeed the explanation that Cummins intended is supported by his later elaboration of the nature of language proficiency as being best described by postulating that degree of difficulty in language tasks varies according to the positions of the features of a language task on two intersecting continua: one specifying the context-embeddedness of the language, and the other the intensity of the cognitive demands placed on the brain (Cummins, 1984a, p. 138; Cummins, 1984b, p. 12). The conceptual basis for the vertical continuum specifying cognitive demands has evidently been taken from information processing theory which delineates two basic types of information processing, automatic and controlled. *Automatic* processes have been learned by repetition of a task and are completed very quickly without the necessity of a high degree of attention, while *controlled* processes are not learned responses and therefore do require focused attention. Because of this, it is difficult to perform more than one controlled process simultaneously without interference. The latter are also tightly capacity-limited and require more time for their activation (see McLaughlin, Rossman, and Macleod (1983) for more detailed information on this topic). A critical factor in determining the cognitive demands of a task is the degree of attention required. As McLaughlin and his group put it:

> [T]he predetermination of relations between elements reduces the amount of cognitive effort required to handle isolated bits of information. By dealing with related units of information rather than isolated bits, more efficient processing becomes possible.

In recent years researchers have investigated the effects of practice, rehearsal, and familiarity on information processing. A particularly important variable appears to be the degree of attention involved. The more attention required, the more resources are consumed and the slower the processing. Greater practice, rehearsal or familiarity with the material allows information to be handled more routinely without as much cost in attentional monitoring as when there is less practice, rehearsal, or familiarity with the material. (ibid., p. 138)

In his conceptualization of this vertical continuum of language proficiency, Cummins (1984b) has manifestly borrowed from information processing theory:

The vertical continuum is intended to address the developmental aspects of communicative proficiency in terms of the degree of active cognitive involvement in the task or activity. Cognitive involvement can be conceptualized in terms of the amount of information that must be processed simultaneously or in close succession by the individual in order to carry out the activity ... Thus, the upper parts of the vertical continuum consist of communicative tasks and activities in which the linguistic tools have become largely automatized (mastered) and thus require little active cognitive involvement for appropriate performance. (p. 13)

It is interesting to note at this point that "cognitive demands" have essentially been equated to "attentional demands". How these concepts explain learning to do a task (of any kind including language tasks) is quite simple: if a task is unfamiliar to the learner, it requires high attentional monitoring and, because of this, is costly in terms of cognitive resources, and moreover, is accomplished more slowly than a task at which the learner is practised. As the task becomes more familiar, less and less attentional monitoring is necessary, therefore it can be accomplished more quickly and the capacity-limited cognitive resources are freed up for other tasks.[4] Thus, what was once cognitively demanding in terms of attentional capacity becomes less so with practice, hence its position on the vertical continuum changes according to the degree to which the task has become well-learned or automatized. Now, since this is true for all tasks, any language task whether speaking or writing or reading is more cognitively demanding if it is not practised, and, as Cummins points out, "skills that have become automatized for native speakers of a language may very well be highly cognitively-demanding for learners of that language as an L2" (ibid., p. 15). He believes, however, that there are some language tasks that have intrinsic characteristics which make them more cognitively demanding than others, and that tasks which are both cognitively-demanding and context-reduced are generally more related to

---

[4] There is evidence now that the attentional system is divided into two distinct anatomic systems one of which subserves the slower capacity-limited processing, and the other the fast essentially unlimited capacity processing (Raichle et al., 1994). This will be discussed in more detail in chapter 2.

academic achievement. Context as defined for the horizontal continuum refers to the "range of *meaningful* [italics added] paralinguistic and situational cues" (ibid., p. 12) available to support the language task.

Commentary

The problem for this schema is that the context embedded/reduced continuum is not free of cognitive demands. If we re-examine McLaughlin's statement above, it can be seen that "the predetermination of relations between elements" of the information will reduce cognitive demands. In other words, if some of the information in a given situation (or information processing task) is already related in the learner's brain, that is, if some of the information is already *meaningful* to the learner, and this can only come about with familiarity with the input, the cognitive demands of the task are reduced. Thus, to have simply a wide range of paralinguistic and situational cues which are not meaningful would not necessarily be any more helpful to a second language learner than to have a narrow range. It is the meaningfulness of the available cues which is of primary importance precisely because it reduces the cognitive load, and this comes about as a result of the reduced attentional load[5]. To exemplify this point, take the situation of a second language learner who has been schooled to a high academic level in his native tongue. For such a person, the

[5] The role of *meaningfulness* in capturing attention will be elaborated upon in chapter 2.

familiarity with the academic environment and the language tasks required within it would render the learning of language related to academic achievement in the second language less cognitively demanding than for someone without academic training. Similarly, unfamiliarity with the customs of a new culture and hence with the meaningfulness of paralinguistic and situational cues would render the learning of language tied to a "context-embedded" situation more cognitively demanding. Note that these concepts support the tenets of the Linguistic Interdependence Hypothesis, but contradict the notion that some tasks are intrinsically more cognitively demanding than others. To repeat for the sake of clarity, then, when something is meaningful, it has already been learned, and what has already been learned demands less of the capacity-limited attention system thereby reducing cognitive demands. Hence the necessity for meaningfulness in the context in order for it to be helpful renders the axes as conceptualized by Cummins overlapping in information content in that both include a cognitive dimension, and consequently they are not orthogonal.

Cummins appears to have sensed that the axes cannot be cleanly disassociated from other factors such as the "personality, learning style" and "L1 and/or L2 proficiency" (Cummins, 1984b, pp. 14, 15) of the particular language users, and he also does appear to have recognized that they cannot be cleanly disassociated from one another. For instance, he answers critics who take him to task on the "vagueness and ambiguity in the notion of 'contextual support'" (Cummins, 1984c, p.72) by distinguishing between internal and external context, by which he means respectively context particular to the language user, and context which is "more or less objectively specifiable along the context-embedded continuum" (ibid., p. 72). These internal contextual factors he perceives as being an influence on the location of any particular task on both the vertical and the horizontal continua, therefore he is in essence conceding that the continua do overlap. However, he conceptualizes these internal factors as "degree of familiarity and acceptance of the task/activity" without clarifying that these are dependent on past learning and must therefore be considered cognitive to some extent. Thus, he sees no need to alter the framework as delineated believing "that these [sociolinguistic] factors can be incorporated in more detail than has been the case up to now without relinquishing the parsimony of two basic dimensions" (ibid., p. 73).

It will be argued, however, that the essential flaw in Cummins' framework is embodied in his effort to isolate a "cognitive" variable as an independent dimension of language proficiency. The information emerging from neuroscientific research is elucidating the degree to which learning is integrated into every aspect of brain processing and, since learning is undoubtedly intimately involved in "cognitive functioning" as it has been defined, cognition itself must determine to some extent even the most fundamental of brain processing abilities including the perception of ostensibly objective information. In other words, it will be argued that attempting to designate some brain processes as "cognitive" and others as "not cognitive" is a futile and even misguided exercise.

How can cognition be measured?

There is no doubt that Cummins believes that psychological tests purporting to measure certain aspects of cognition do just that. Not only does he accept the results of well-conceived research involving cognitive measures as being valid in the development of his theoretical concepts, but he uses these tests in his own research without questioning their efficacy. Several comments (Cummins, 1977) attest to his faith in these measures:

> Scores on Raven's Progressive Matrices (a nonverbal measure of intelligence) did not significantly predict French speaking ability. (p. 4)
>
> *Cognitive measures*: Two subtests of the Kuhlmann Finch Intelligence Test designed to measure verbal ability (subtest 1) and general reasoning (subtest 4) were administered. The Utility Test from the French, Ekstrom and Price (1963) Kit of Reference Tests for Cognitive Factors was administered as a measure of verbal divergence. (p. 5)

The matter-of-fact statement of what these tests measure is indicative of an uncritical stance towards the validity of the tests and even what counts as a cognitive characteristic, though the latter is hardly surprising given the wide range of phenomena that are included under cognition. This is not to say that Cummins is unfailingly uncritical towards testing in general (see for example Cummins, 1984, chaps. 2 and 3), but in his critiques he appears to question the validity of using certain tests in particular situations, rather than questioning the validity of the test itself. For instance, while at one time deploring the "lack of appreciation of the limitations of IQ tests such as the WISC-R" which reflect the "dominant middle-class culture experiences" and on which students from ESL backgrounds perform poorly (Cummins, 1984a, p. 46), at a later time he appears to accept the same test as a valid measure of overall IQ:

> Thus, instruction through a minority language will usually result in transfer to the majority language but the opposite is seldom the case because of the lower status of the minority language.
>
> It is important to note that what is transferred is primarily conceptual knowledge rather than specific linguistic elements. It is this conceptual knowledge together with its linguistic manifestations that is being tapped by most verbal academic measures. For example, the best predictor of overall IQ on the Wechsler Intelligence Scales for Children is the vocabulary subtest. (Cummins, 1991, p. 77)

An examination of the vocabulary subtest in question, however, reveals that the "conceptual knowledge" that it is tapping into can be more accurately described as general knowledge which is highly experience-dependent (not to mention language-dependent) and is therefore not a valid measure of intelligence or cognitive ability if this is conceived of as independent of lived experience or "basic"[6]. Interestingly enough, this is precisely the point which Cummins made previously.

This issue of what constitutes good evidence for conclusions on cognitive ability or language ability will be treated from different perspectives under the section on analytical concerns.

Summary of Cummins' view on cognition

It is clear on analysis of Cummins' statements on cognition that, in his opinion, cognition consists of an inner core of (basic) abilities which can be applied to the learning of particular subject areas such as language, and that the result of this application will be changes in the cognitive abilities themselves,

---

[6] Sattler (1982) notes: "The Vocabulary subtest, a test of word knowledge, may involve a variety of cognitive functions or features - including learning ability, fund of information, richness of ideas, memory, concept formation, and language development - *that may be closely related to the child's experiences and educational environment* [italics added]" (p. 175). In spite of this, he concludes that the subtest "likely reflects their [the children's] ability to learn and to accumulate information," "provides an excellent estimate of intellectual capacity," and "is valuable in deriving an index of the examinee's general mental capacity" (ibid.). Hence Cummins is merely echoing the accepted opinion of psychological assessment experts.

changes which will reflect the particular learning tasks accomplished. The cognitive abilities themselves are hierarchical in nature, with tasks which require the processing of unfamiliar information being more cognitively demanding than those which process familiar information. Although Cummins did not clarify this point with respect to the role of cognition, it is evident from the application of information processing theory that tasks in which there is a large range of meaningful task-related (e.g. paralinguistic) and situational cues will be less cognitively demanding than those supported by only a small range of cues. Cognitively demanding tasks are in general equated with tasks requiring high attention. In terms of its ability to effect change on cognition, language is considered to be extremely effective, and bilingualism even more so due to the putative ability of bilingual speech actions to generate a qualitatively different kind of feedback data for cognition to act upon. Cognitive factors are seen as separate from social, political, attitudinal and educational factors but as capable of being influenced by them through the mediating effect of language. Since it can also be influenced by internal factors, termed "child input" factors, cognition may be said to operate by processing data from the environment, both internal and external. As far as measurement of cognitive abilities is concerned, Cummins shows a general acceptance of the validity of familiar tests of cognitive ability, and also appears to consider that cognitive growth is mirrored by academic achievement.

This portrait of cognition will be compared in a later chapter with what can be gleaned about the underlying processes of cognition, such as learning and attention, from a consideration of connectionism and neuroscientific discovery on the inner workings of the brain.

(2) <u>Thresholds, linguistic competence and degree of bilingualism</u>

The concept of a *threshold level of linguistic competence* has always been ill-defined as Cummins himself readily admits. In the original presentation of the idea he inquires into its potential nature in these words:

> What are the characteristics of this threshold level of bilingual competence? In the first place the threshold cannot be defined in absolute terms; rather, it is likely to vary both with the amount of time that is spent through L2 and with the type of cognitive operations that must be expressed through L2.... The threshold is also likely to vary according to the type of cognitive operations appropriate for a child's stage of cognitive development. (Cummins, 1976, p. 24)

According to this description, then, the linguistic threshold depends on both internal (stage of cognitive development) and external situational factors relating to cognition. But the stage of cognitive development of a learner is itself dependent to a large extent on past experience as well as individual ability and must therefore be distinct for each individual at any age. In addition, the cognitive operations demanded of the learner will change from learning situation to learning situation. By Cummins' own defining musings, therefore, the threshold level of bilingual competence must be considered intrinsically unspecifiable on any precise basis, a conclusion which Cummins himself appears to have reached in recent times (see section 1.2). Furthermore, given the relationship drawn between level of bilingualism and cognitive ability by the Threshold Hypothesis, a decision as to whether or not a child has reached the threshold level of *linguistic* competence in the second language would have to be made on a reading of whether or not the child could perform adequately the *cognitive* tasks that were being demanded of him or her. Hence, the definition of the threshold is the capability of performing certain cognitive tasks, which may or may not be language related, considered appropriate to the child's stage of cognitive development and educational circumstances[7]. And when one considers that what explains the child's ability to perform these tasks is that she has attained the threshold level of bilingual competence, it can be seen that the argument is circular and the construct is a tautology. Edelsky et al. (1983) have levelled the same type of criticism at Cummins' CALP construct. Moreover, the concept of a threshold level of bilingual competence provides no insight whatsoever into what has happened in the brain of the child to

---

[7] These are points which will be expanded upon in chapter 3, but see also Hawson (in press-a).

promote the required learning and is therefore of little value in directing educational efforts which, given Cummins' interest in academic achievement, is presumably the ultimate aim of his hypothesis.

Another difficulty arising out of the notion of a threshold level of bilingual competence, one that has been raised by Martin-Jones and Romaine and was alluded to by Cummins himself (1991, p. 85), is the conceptualization of *language proficiency* (or linguistic competence) which it promotes. According to Martin-Jones and Romaine (1986), the very idea of "levels of linguistic competence" attracts quantitative descriptors of the same type as those born of a "container metaphor" of linguistic competence:

> The literature on "semilingualism" abounds with terms such as "full competence," "threshold level," "additive" and "subtractive bilingualism".... Thus, it seems that linguistic competence has so far been conceptualized in terms of an implicit "container" metaphor: a container which can be either "full" or "partially full." Terms such as "semilingualism" are, in our view, misleading because they implicitly foster the belief that there is such a thing as an ideal, fully competent monolingual or bilingual speaker who has a full or complete version of a language. (p. 32)

That Cummins' theories do tend to attract "container metaphor" terminology can be seen in a study by Collier (1987) which was designed to follow up on Cummins' research (as well as that of other researchers) on the significance of the BICS/CALP dichotomy of language proficiency. Collier not only uses the term "full proficiency" but even describes its nature:

> Second language is acquired to varying degrees of proficiency depending on the context in which the acquirer needs to use it. Immigrants of school age who must acquire a second language in the context of schooling need to develop *full* [italics added] proficiency in all language domains (including the structures and semantics of phonetics, phonology, inflectional morphology, syntax, vocabulary, discourse, pragmatics, and paralinguistics) and all language skills (listening, speaking, reading, writing, and metalinguistic knowledge of the language) for use in all content areas (language arts, mathematics, science and social studies). (p. 618)

While descriptions of this type can be useful in demonstrating just how complex a problem the school age immigrant is faced with, the word "full" is redundant, since the proficiency needed is precisely that demanded by the tasks themselves and may be neither full nor even adequate to the demands of other tasks, as for example, the understanding of street slang or technologically specific terminology. The point being made by critics like Martin-Jones and

Romaine is that no one is fully competent in a language in that precise knowledge of all highly specific codes within a target language is virtually impossible for a single individual. They consider, moreover, that the terminology related to the container metaphor

> reflects the ideological bias of a linguistic theory which has been primarily concerned with the idealized competence of monolingual speakers in the speech communities of western Europe and the United States; communities which, on the whole, have a high degree of stability, autonomy, and historicity, and possess highly codified standard languages. (Martin-Jones and Romaine, 1986, p.33)

They argue that in many socio-cultural settings "mixed codes" defined by the specific circumstances of the setting emerge, though these are often stigmatized by the majority language groups. Given the demographic instability of some of the old established nations in Europe today, it will be interesting to observe how the notion of full proficiency in a language will evolve. Thus, according to these critics, to try to define an appropriate threshold level of linguistic competence outside of the specific context in which the language is to be used is misguided, and may promote a negative assessment of the real abilities, cognitive and linguistic, of certain socioeconomic groups. Notice that this argument is essentially a return to the concerns of critics worried about the problems of isolating a psycholinguistic factor of language proficiency and how cognition was being conceptualized.

Related to the issues surrounding the concept of threshold levels of bilingualism and the container metaphor is that of how to conceptualize the *degree of bilingualism* of a learner. MacNab (1979) notes that many of the studies in which bilingualism was associated with positive cognitive and academic results used "balanced bilinguals" and "unilinguals" whose degree of bilingualism or unilingualism was determined according to the criteria established by Peal and Lambert in 1962. Cummins adapted some of these criteria in several studies. MacNab questions, not only the validity of the criteria, but also some of the basic assumptions of their originators. He believed that the selection procedures biased the sample selection "towards brighter bilinguals and duller unilinguals". As he notes:

> There is certainly reason to think that the 'balance' criterion has altered the sample in some fundamental way on the basis of numbers alone.... where the earlier [negative] studies tended to take whole schools or classes, the language balance criteria used by Peal and Lambert defined over half of the otherwise eligible ten-year-olds as neither bilingual nor unilingual. (MacNab, 1979, pp. 238-239)

Whereas Anisfeld (formerly Peal) and Lambert (1969) argue that "the balance measure would in no way penalize the child who was slow in two languages: he had only to be equally slow" (p. 125), MacNab (1979) concludes that "the assumption that a slower child will tend to learn two languages equally, rather than emphasizing one to learn well, is doubtful. It certainly should be explored rather than assumed" (p. 241). It seems strange that Peal and Lambert should object to criticism on this point when they themselves commented on the possibility of the sample being unintentionally biased in their own study (see Reynolds, 1991, pp. 157-158), and stranger still that researchers like Cummins would have continued to use the criteria developed knowing of the potential for bias in the sample. This is all the more true when consideration is given to the fact that the requirement of "balanced bilingualism" severely limits the generalizability of the results (Reynolds, 1991, p. 155). Perhaps at the time, as Carey (1993a) has indicated, in Canada at least, the political wind encouraging positive results was so strong that it swept concerns about applicability aside.

Martin-Jones and Romaine (1986) object to the notion of balanced bilingualism on different grounds; they note that, as defined by Peal and Lambert, the balanced bilingual can be described as "one who has equal, though not necessarily full knowledge of two languages" (p. 32). Thus even people who have a low level of competence in both languages can potentially be balanced bilinguals. Martin-Jones and Romaine contend, however, that in practical terms, "the notion of balanced bilingualism has...functioned as an implicit synonym for 'good' or 'complete' bilingualism and has been used as a yardstick against which other kinds of bilingualism have been measured and stigmatized as 'inadequate' and/or 'underdeveloped'" (p. 33). It is the value ascribed to balanced bilingualism, therefore, that they object to rather than the term itself, as well as the reliance of the determining of such states on assessment procedures at a time when there is no general agreement about "the 'normal' course of development among monolingual, let alone bilingual children" (p. 34). The basis for the objection of Martin-Jones and Romaine, then, appears to lie in the association of balanced bilinguals with positive cognitive and academic results, and that of "unbalanced" bilinguals with negative cognitive and academic results, associations which, as MacNab and Reynolds have pointed out, rest on treacherous ground. In his profound and detailed analysis of the research on the cognitive consequences of bilingualism, Reynolds (1991) concludes:

> Due to the prevalence of design problems and other logical difficulties I have described, the central question we should be asking ourselves is not whether the thesis that bilinguality affects cognitive performance is true or whether the reverse is true, *but whether there is a relationship at all* [italics added]. (p. 159)

Given the degree of uncertainty in all areas of this field of study, this is a question which demands serious inquiry.

1.3.3 Philosophical concerns

The Threshold Hypothesis as a deficit hypothesis[8]

> [T]he assumption of an intervening variable such as the child's so-called "conceptual linguistic knowledge" (or lack of it) in explaining the outcome of educational programmes for minority children constitutes a deficit theory in the making. This deficit view is, as we see it, most clearly expressed in Cummins' "threshold hypothesis." (Martin-Jones & Romaine, 1986, p. 31)

The opinion of Martin-Jones and Romaine expressed above is shared by several others, most notably Edelsky et al. (1983), Canale (1984), Wald (1984), and Commins and Miramontes (1989). What these authors are protesting is the designation of minority students (or any student) as being deficient cognitively if not able to achieve adequately on tests administered in an academic environment. They perceive a designation founded on such a basis as reflecting uncritical acceptance of the societal status quo, within academic circles and without, with respect to whose knowledge is valued. This is a position with which Edelsky et al. (1983) take great

---

[8] Although it is evident that the term "theory" should not be applied to the Threshold Hypothesis, much of the discussion around deficit charges against the latter has employed the word "theory". For the sake of accuracy and simplicity, therefore, this discussion has been recorded as it occurred in the literature.

umbrage:

> Underlying both Cummins' and Swain's attempts to account for differential school success among second language learners is the premise that the predominant current definition of success in school and particularly success in literacy is right, acceptable, reasonable, etc.; that what is, should be. They do not question the usual definition of reading as the ability to perform well on a reading achievement test, the definition of writing as the ability to do work-sheet-type exercises on mechanics, vocabulary, synonyms, analogies, etc.; definitions that equate literacy with performance in discrete, otherwise purposeless tasks intended as practice for some other time when the "real event" occurs. This unquestioning acceptance of current school definitions and current school curricula is the flaw that we believe leads to all the rest of the errors. (p. 4)

What is suggested by these authors is that given "real event" tasks, these students might perform much better than they do on tests. This view is supported by the work of Commins and Miramontes (1989), a study of "four Hispanic bilingual students whose language dominance was not clearly defined and whose academic achievement was perceived by teachers to be limited by their language abilities in both English and Spanish" (p. 448). Observing these students over a variety of contexts both inside and outside the classroom, they concluded that the teachers had underestimated the linguistic and academic abilities of the students based on their classroom performance, and that what was lacking in the school context for them was "the essential element of relevance" (p. 467). Relevance, except perhaps as it is reflected in the desire to perform well to please others and satisfy the demands of the academic system, is precisely what is lacking in formal test situations. Unless the student "buys into" the system, there is little if any motivation to do well on tests. Thus, ability to do well on tests of intelligence or "cognitive" measures has been termed "test-wiseness" by Edelsky et al. (1983), and defined as "an ability that incorporates a desire to do well on artificial (as opposed to real-world) tests" (p. 4). In a similar vein, Canale speaks of the importance of "socialisation and acculturation" to performance on a given language test, and comments: "In this light, test results may be better viewed as *created* rather than simply observed or found" (Canale, 1984, p. 38).

Underpinning all of these protestations, then, is the notion that what is being labelled a deficiency in ability is simply an inability, or even a lack of desire, to perform as expected on standardized tests. This "deficiency" is determined according to the standards of the status quo, and may or may not reflect some underlying cognitive deficit. In other words, what is being implied is that terms such as "cognitive deficit" and "cognitive disadvantage" are being defined by the demands of the society as well as those of the task, demands which incorporate the necessity for a positive attitude towards the successful accomplishment of the task. It is implied by the authors above that this is at best an uncertain footing for the founding of theories or hypotheses which relate cognitive benefits to the academic performance of learners.

For his part, in an article written jointly with Merrill Swain, Cummins strongly defends his theoretical constructs against "deficit theory" allegations (Cummins and Swain, 1983):

> Our contention is that the theoretical distinctions we have made do not constitute a deficit position; on the contrary, they have led directly to an increased awareness on the part of many North American educators of the ways in which current school programs and personnel are actively creating educational deficits in minority students. (p. 23)

In an effort to understand how his theoretical constructs could have been read in this light, he criticises Edelsky et al. for not having defined what in their eyes constitutes a deficit theory and comments:

> In the absence of any analysis of this issue, we must conclude that they base their deficit allegation on the fact that we view attributes of minority children (and *all* children) as

relevant to their educational performance, within the context of an overall causal model involving sociopolitical, educational, and other factors. Virtually all other theories in this area of which we are aware (including those endorsed by Edelsky et al.) invoke similar causal models, either implicitly or explicitly ... (ibid., p. 26)

Commentary

Though Cummins' refutation of the "deficit theory" charges laid before him is lengthy, in depth, and, on many counts, highly persuasive, there are some dangling issues which still breed cause for concern. First of all, it is interesting to note that, in the first citation above, he has interpreted the "deficit" charge to be that of an "educational deficit" rather than a "cognitive deficit" which Edelsky et al. suggest his Threshold Hypothesis implies (Edelsky et al., 1983, p. 2), or even a "conceptual linguistic deficit" as referred to by Martin-Jones and Romaine. This confusion of education, language, and cognition is yet another example of the compounding of the concepts of academic achievement, linguistic ability, and cognitive ability discussed previously, and one which only serves to cloak the issues. Secondly, in the Threshold Hypothesis, Cummins talks of students suffering "cognitive disadvantages" and benefitting from "cognitive advantages" at opposite ends of the degree of bilingualism scale. In his discussion of the charges, however, he avoids the terms "advantages" and "disadvantages," preferring instead to refer to learner "attributes" and "individual differences." In doing so, he appears to be acknowledging that the latter are more neutral than the former which carry implicitly the notion of a judgement, or at least a reference base against which performances can be compared. The very use of the terms "advantage" and "disadvantage" in the elaboration of the theoretical constructs, therefore, would appear to support the contentions of Cummins' critics with respect to a deficit position charge. Evidently, what he perceives as being of value are indeed the academic criteria set by the status quo, since it is against such criteria that the performances of the students are judged. Thirdly, the modifier "cognitive" in "cognitive disadvantage or advantage" is not neutral in value as it is used in education. Since cognition is considered to be a facet of brain functioning which nurtures thinking and its underlying processes (as opposed to emotion, volition, and motoric functioning) it is especially revered by educators. Hence, to say that someone is cognitively disadvantaged is to pronounce that he or she is in a state of double jeopardy educationally, and conversely, to be cognitively advantaged is to be doubly blessed. That this view of cognition is in need of revision, will be broached in chapters 2 and 4.

In conclusion, then, it appears that, in spite of protestations to the contrary, there is some justification for the "deficit theory in the making" charges levelled at Cummins' Threshold Hypothesis.

That it was never his intent to create a deficit hypothesis is very evident, however, and it is equally evident that practical applications of Cummins' theoretical constructs have had beneficial effects for many minority group learners.

Research tradition and choice of data

It is notable that the studies upon which the Threshold Hypothesis and related theoretical constructs are built are in the tradition of scientific inquiry in that they are empirical in nature with the dependent variable being a performance score on a standardised test of some sort, academic or psychological. As has been noted above, this dependence on tests has drawn fire from critics who view the test situation as specific to school culture. On this view, ability to perform well on tests is seen as a learned response, which necessitates both a learning period and motivation towards learning to do well on tests. In short, good performance on tests is due to some unknown extent on the degree of acculturation of the student. In a more general vein, Edelsky et al. (1983) are critical of the dependence on data representing the philosophical perspective of a single research tradition: "Research in this tradition attempts to separate confounding variables and to isolate some particular notion for analysis. It thus risks misrepresenting human existence, transforming its interactive, constitutive nature ... into something simplified, static, and non-interactive" (p. 5).

This complaint is echoed by Martin-Jones and Romaine (1986, p. 29) who note that "[l]anguage skills cannot be neatly compartmentalized in the way that Cummins suggests," and that the assumption "that it is possible to separate analytically different aspects of language competence without reference to the context of usage" is highly questionable. These authors, along with others (Commins and Miramomtes, 1989) are in favour of a more naturalistic approach to research in second language learning, so that the interactive nature of language learning can be documented.

Paradoxically, standardized test results have also been criticised by these authors for failing to be analytical enough in highlighting individual differences, or at least small group differences within larger group designations. As Edelsky et al. (1983) argue, it is averaged data which "rarely analyses the ranges or clusterings within ranges that give rise to different means among different groups, making it appear that an entire group is different from another entire group" (p. 6). For this reason also, then, naturalistic inquiry is viewed as the more informative research option by these authors.

In response to these critiques, Cummins and Swain (1983, p. 28) have noted that "authentic reading measurement procedures" such as miscue analysis and cloze tests do relate strongly to standardized reading test scores, and that many of the tests they have developed require students to read "real life" texts like "printed stories, newspaper articles, poetry, recipes, magazine clippings" and so on.

Cummins also points out that the results of these tests and standardized reading tests are consistent. Though Cummins does not endorse the widespread use of standardized tests for either minority or majority students, he does believe that they can provide useful information under certain circumstances, and considers universal condemnation of them "simplistic and unsupported" (ibid., p. 29). "A more constructive strategy", he contends, "would be to document for practitioners the ways in which their [Edelsky et al.'s] current assessment practices are prejudicial for some groups of students, i.e. the conditions under which tests are nonsense" (ibid.).

Commentary

In spite of the validity of Cummins' arguments on the particular points he has chosen to address, there are two issues which he appears to have overlooked. First, a test situation is a test situation whether the test be a standardized reading test or a cloze test; to perform well on any test, a student must be knowledgeable of the general meaningfulness of a test situation and its larger ramifications. This is, in part, what is meant by "acculturation" to the school system or, as Edelsky et al. term it, "test-wiseness". Thus changing the nature of an item on a test would not necessarily lead to a changed result. Second, even documenting the ways in which tests are prejudicial to certain groups, or subgroups, of students would not explain why individual students do well or badly while their confreres do not. Hence, although standardized tests may provide useful information with a "broad strokes" validity within the constraints of the test parameters (dictated by the actual test items), they cannot probe the precise nature of how an individual learns a language. However, given the number of variables interacting in such a situation, any attempt at elucidating this process would appear to be doomed to failure. Nevertheless, ethnographic studies of the type carried out by Commins and Miramontes (1989) can serve as checks on information obtained through standardized test results and prevent researchers from making the kind of sweeping conclusions which have too often in the past led to faulty theorizing.

The major point being made by critics of standardized testing is that it is provides "top-down" information which is shaped to a high degree by the philosophical biases of the researchers, and that "bottom-up" considerations must be taken into account to counteract these biases, especially where human beings are the subjects in question. This is an issue which will be returned to in a later chapter.

1.3.4 Analytical concerns

Methodological inadequacies of the studies used

In a coherent critique of the methodological problems with both the studies pointing to cognitive disadvantages arising from bilingualism, and those purporting to show cognitive advantages, MacNab (1979) arrives at the conclusion that there is no solid evidence to support the claims of these studies. A more recent analysis of the data by Reynolds (1991) agrees. Both of these authors fault researchers on the inadequacy of the research designs employed, noting that many of both the pre- and post-Peal and Lambert study era were *ex post facto* designs which are "replete with threats to internal and external validity" (Reynolds, 1991, p. 148). Reynolds discusses a variety of the research designs used up until the time of publication (1991) and found the validity of all of them, both pre-experimental and quasi-experimental, to be compromised in some way. MacNab (1979) notes that associational studies do not warrant a directional arrow of causality to be drawn, and Reynolds' (1991) analysis of issues related to design and the independent variables, which includes more recent studies, ends on this sombre tone:

> The research prior to 1962 has repeatedly been criticized as methodologically inadequate and I fear we are notomuch [sic] better off today; even sophisticated techniques like regression analysis, factor analysis, path analysis, etc. do not compensate for inherently trouble prone designs. In addition, this body of research is plagued by procedural shortcomings, problems with the definition of bilinguality, and the dangers of selection artifacts and effects occurring with the manipulations. Due to the prevalence of design problems and other logical difficulties I have described, the

central question we should be asking ourselves is not whether the thesis that bilinguality affects cognitive performance is true or whether the reverse is true, but whether there is a relationship at all.... we find the situation is not much better when dependent variables are examined. (p. 159)

Commentary

Given this indictment of the research upon which the Threshold Hypothesis was founded, it would appear that this hypothesis has emerged from a pool of tainted research, especially since it was in part delineated taking into consideration the least reliable early studies. In accepting the validity of the bilingual research data on the cognitive consequences of bilingualism, Cummins accepted implicitly the validity of a relationship between cognitive ability and level of bilingualism, which both of the above authors have questioned, and, though critical of some of the results reported, sought merely to make sense out of the pattern they presented. If, as Reynolds suggests, however, there is reason to question the existence of such a relationship, the rationale for his argumentation collapses. It should be noted, nevertheless, that, using a "within group" design, Diaz (1985) investigated the directionality of the relationship between degree of bilingualism and cognitive ability and came out in agreement with Cummins: that is, his results supported Cummins' assumption that bilingualism is causal in its effect on cognitive ability. Hakuta (1987) corroborated the results obtained by Diaz using a similar longitudinal design to assess causality. These authors, then, agree with Cummins with respect to both the existence of a relationship between these two variables, and the direction of the causal arrow. It will be argued in a later chapter, however, that both these authors are defining cognition in very specific ways, and that their results can be better understood through a connectionist vision of second language learning which gives a clearer view of the particular systems in the brain that may be involved in any changes in information processing brought about by the immersion second language learning setting. Also, as Reynolds (1991, p. 152) points out, even the more sophisticated designs of these experiments have inbuilt hazards to validity. In chapter 3, therefore, the question of the nature of the relationship between level of bilingualism and cognitive functioning will be probed more deeply to clarify just how cognition is being conceptualized by the bilingual research community.

1.4 Summary of Chapter One

Research with respect to the cognitive consequences of bilingualism has yielded very conflicting results. It has been proposed by some that prevailing political attitudes in the countries in which the research was carried out coloured the outcome of the research itself. In political climates of assimilation, most studies reported a negative correlation between bilingualism and cognitive consequences. The Peal and Lambert study in 1962 in which a positive correlation between bilingualism and cognitive ability was found marked the beginning of a new era during which, increasingly, positive correlations were being found. Proponents of the two opposing positions were in constant conflict and extreme statements were made by both sides. Eventually, social factors were recognized as also being influential in determining cognitive and linguistic outcomes.

Cummins attempted to make sense out of the conflicting data by proposing in the Threshold Hypothesis that there may be a threshold level of bilingual ability that must be attained by a learner in order to avoid negative cognitive consequences, and another higher threshold beyond which positive cognitive effects associated with bilingualism would begin to manifest themselves. To this hypothesis, Cummins later added the Interdependence Hypothesis which postulates that there is a common underlying linguistic proficiency upon which all languages, first or second can draw, and that because of this, learning a second language need not detract from linguistic ability in the first language. It was determined, that, due to a chronic lack of definition of terms, the Threshold Hypothesis is more appropriately classified as an heuristic than as a test hypothesis.

The Threshold Hypothesis has been a source of controversy since its introduction. Criticisms of it have been launched on conceptual, philosophical, and analytical grounds, and a critical assessment of these critiques was undertaken in this chapter. Conceptual concerns can be gathered under two main headings: the difficulty in isolating a psycholinguistic factor of language proficiency from sociolinguistic concerns, and problems involving definition of constructs. Of major import among the latter are what has been termed "cognitive issues." Several authors have been highly critical of Cummins' conceptualization of cognition, and, because of this, an in depth examination of Cummins' own statements relating to cognition and cognitive issues was undertaken. It was discovered that Cummins views cognition as consisting of a core of "basic" cognitive abilities that are distinct from cognitive abilities accrued through learning in the various subject areas, but are capable of being affected by such learning. Essentially a proponent of Vygotsky's theoretical position on language and cognition, Cummins sees language as a particularly powerful influence on cognitive ability, and second language learning as having a unique capacity to affect cognitive development as a result of its being able to generate a qualitatively different kind of feedback for the learner. Criticisms of the definitions, or lack thereof, of the concepts linguistic thresholds, linguistic competence, and degree of bilingualism were also probed and found to be justified.

Philosophical concerns clustered around claims that the Threshold Hypothesis is a deficit hypothesis which casts second language learners at the lower end of the degree of bilingualism scale in a negative light, and complaints that there is a bias in Cummins' choice of seminal data for the hypothesis which favours the scientific research tradition over naturalistic inquiry. The use of the results of standardized tests and averaged data was particularly criticized. Support was found for both these charges.

Lastly, the analytical concerns of critics addressed the methodological inadequacies of the studies upon which Cummins constructed the Threshold Hypothesis. Past and recent reviews of the methodologies employed in research into the cognitive consequences of bilingualism highlighted problems of internal and external validity in even the most carefully crafted designs. As a result, whether there is a relationship at all between bilingualism and cognitive ability has been called into question by Reynolds.

Chapter 2. A CRITIQUE OF CONNECTIONISM

2.1 Historical Perspective

Understanding the mind may not be as intricate as our vanity hoped or our intellect feared.

Rodolfo Llinas, 1986

The study of the nature of conscious intelligence has long preoccupied humankind, and opinions on this topic have always been contentious. As far back as the fifth and fourth centuries B.C., Greek philosophers debated the issues and took their stands: Plato postulated that the world could be divided into two, the world of sensibilia (the physical world) and the world of intelligibilia (the non-physical world), and that Knowledge of Reality was to be found only in the non-physical world, the proper sphere of operations for the rational soul which was itself non-physical; Democritus, on the other hand, claimed that the world was made up of fundamental particles called atoms which linked together in different patterns to form the various physical substances, and that since humans were merely highly complex combinations of the fundamental particles, human minds were fundamentally material, their remarkable capabilities being a product of the complexity of their organization. Thus Democritus expressed what has come to be known as the *materialist* view of the mind, while Plato adhered to a *dualist* perspective, and, though volumes upon volumes of texts have been written since their day, contemporary upheavals in philosophical thinking still follow roughly the same fault lines.

The problem of the nature of conscious intelligence or the mind has been termed the *ontological* problem or the "mind-body" problem, which, according to Paul Churchland, attempts to deal with questions of the type:

> What is the real nature of mental states and processes? In what medium do they take place, and how are they related to the physical world? Will my consciousness survive the disintegration of my physical body? Or will it disappear forever as my brain ceases to function? Is it possible that a purely physical system such as a computer could be constructed so as to enjoy real conscious intelligence? Where do minds come from? What are they? (Churchland, 1984, p. 7).

Ruminations on these questions have led not only to the basic dualist and materialist[1] divisions in philosophical thinking, but to splintering of opinions within these two camps. Dualism, for example, has split into "substance dualism", which considers the mind to consist of a fundamentally different substance than the body (the opinion of Plato and later Descartes), and "property dualism", in which, though the substance of the mind is conceded to

---

[1] Sometimes referred to as "monist," as materialism is sometimes referred to as "monism."

be the brain, the mind itself is thought to have special non-physical properties possessed by no other physical object, as well as into several different flavours within these two subdivisions (see Churchland, 1984, chap. 2). Materialism has also spawned a number of different lines of thinking. "Philosophical behaviourism", for example, claims that any statement about a mental state can be rephrased *without loss of meaning* into a statement about the behaviour that would result in a particular observable circumstance. "Reductive materialism", commonly referred to as the "identity theory" or the "identity problem", on the other hand, postulates that mental states are identical to physical states of the brain or central nervous system and foresees that, as man's knowledge of the workings of the brain grows, everyday theories about mental functioning will be reduced to theories about brain function in much the same way as classical mechanics (Newtonian mechanics) was succeeded by Einstein's theory of relativity. As the popularity of behaviourism waned, it was supplanted by "functionalism", the central premise of which is "the thesis that mental states are defined in terms of their abstract causal roles within the wider information-processing system" (P. S. Churchland, 1986, p. 351), that is, a given mental state is considered to have "abstract causal relations to environmental input, to other internal states, and to output" (ibid.). The heir to a battered and bruised identity theory was "eliminative materialism" which side-stepped the pitfall that tripped its precursor, the demand for a neat one-to-one match-up between what has been coined "folk psychology", our common sense understanding of the world, and theoretical neuroscience, by agreeing with critics that no such match-up would be forthcoming. The rationale put forward by advocates of eliminative materialism, however, is not that of detractors of the identity theory, that a "great variety of quite different physical systems ... could instantiate the required functional organization" (P. M. Churchland, 1984, p.43). Successful reduction between the neuroscientific understanding of the mind/brain and folk psychology will not occur, according to eliminative materialists, because the concepts inherent in our folk psychological understanding of the world including our mental states are radically misguided. So the debate has raged and *still* rages on. (See Searle, 1992, for a recent update on the lines of argumentation.)

At first glance, it is hard to understand why there is so much dissension among philosophers on this question of how mind and body are linked. On closer inspection, however, it becomes evident that the ontological problem is not discrete: the issues surrounding it overlap onto territories covered by semantics, epistemology, and methodology. The *semantical* problem speaks to queries as to how our common-sense vocabulary of psychological functioning gets its meaning. In essence it insists that if we are to attempt to find a working relationship between the mind and the body, we must first have a clear idea of how it is that the mind forms the concepts linked to the words in question. In other terms, the semantical problem asks how it is that words come to be meaningful, and how meaning becomes wordful. The *epistemological* problem, on the other hand, directs its attention to the question of how human beings can acquire knowledge of the inner workings of conscious, intelligent minds, their own and those of others. Thus, the first part of this problem deals with concerns about *self-consciousness*: How is it that a person comes to have "an immediate and privileged knowledge of [his or her] own sensations, emotions, beliefs, desires, and so forth? How is this possible? And just how trustworthy is that knowledge?" (P. M. Churchland, 1984, p. 67). The second part of the epistemological problem examines how we know that something other than ourselves (an alien being, a computer, or even a comatose human) is or is not a conscious being, possessing thoughts and feelings: What are the standards we use or should use? How do we judge? This is the problem of *other minds*. Crucially important to the solution of all of the above problems is the *methodological* problem which addresses the issue of what are the best intellectual methods to employ in a science of mind, a topic upon which there is little agreement, though all parties do agree on the necessity of pursuing investigations into how the mind works. Approaches to date have varied from *idealism*, in which the fundamental stuff of existence is considered to be mind, not matter, to *phenomenology* whose advocates are of the opinion that the methods that guide science cannot lead to a true understanding of the nature of mind since "the world-of-our-experience is in large measure a constructed world" (P. M. Churchland, 1984, p.85), to *methodological behaviourism* which attempts to reconstruct a science of psychology according to the traditions of the highly successful hypothetico-deductive methods of the physical sciences such as chemistry and physics, on to the *cognitive/computational* approach which unites the research programmes of cognitive psychology and artificial intelligence, both of which postulate the occurrence of information-bearing internal states in intelligent creatures which interact collectively to form the mind (this is the so-called information processing or computational approach), and finally to the "bottom-up" perspective of *methodological materialism* whose basic premise is that cognitive activities are simply activities of the nervous system, and that, consequently, the best manner in which to gain access to the nature of cognitive activity is to investigate the functioning of the nervous system itself, that is to say, to study the behaviour of its microscopic[2] elements, their connections and interactions, and their development over time as well as their control over behaviour.

Evidently, then, there is good reason for the lack of concerted opinion on the nature of the relationship between the mind and the

---

[2] With today's technological expertise, the degree of "micro" investigation that can be attained is astonishing. Techniques have been developed which have elucidated the functioning of single macro-molecules such as ion channels, as well as ions themselves within the cells, all of which has added an incredible richness and depth to earlier discoveries of nervous system function using less powerful investigatory tools. Hence, it could be said that within the designation of "microscopic" as a level of investigation, there is now a multiplicity of levels of investigation. In other words, one person's micro is another person's macro.

body! Inclusion of these associated problems into the sphere of the ontological problem brings about a quantum increase in the complexity of the issues. However, just as the ontological problem cannot be cleanly separated from the semantical, epistemological, or methodological problems, the divisions within each of these problems are not categorically discrete. This is particularly clear in the descriptions of the various methodological approaches advocated: Is it not more proper, for example, to consider some of today's artificial intelligence models in which a great deal of the structural detail is based on the fine-grained structure of the brain itself to be as much products of methodological materialism as of the information processing approach?

The arbitrariness of many of the lines drawn between the methodological approaches outlined gave birth to a more synthetic approach, which came into prominence in the late 1970s and early 1980s, called *connectionism*, but also entitled "parallel distributed processing"[3] (see Rumelhart & McClelland, 1986, Vol. I, chap. 2). Clark (1989) says that connectionism is

> an attempt to provide slightly more biologically realistic models of mind. Such models, though hardly accurate biologically, are at least *inspired* by the structure of the brain. Moreover, they are tailored ... to evolutionarily basic

[3] However, in recent times these terms have come to represent different sub-divisions within the discipline. This is discussed at the end of this section.

problem-solving needs, like perceptual pattern completion.... [and although] [b]oth our deeper explanatory understanding of cognition and on some occasions our actual processing strategies may well *demand* the use of higher level, symbolic, neoclassical descriptions.... connectionist models offer insights into the way nature may provide for certain properties that seem to be quite essential to what we consider intelligent thought. (p. 83).

Since the theoretical models of mind developed by connectionism are instantiated in computer models, connectionism reaches into at least a triad of disciplines, neuroscience, artificial intelligence, and cognitive psychology, and exists at the interface of all three (and beyond).

The relationships within this triad can be described in simple terms as follows: cognitive psychology deepens our understanding of the behavioural phenomena to be explained; computational neuroscience studies how the neurons, the basic units of the brain, when given a sensory input pattern, integrate the information and compute an output which initiates specific behaviour; artificial intelligence uses information obtained from neuroscientific studies to build "biologically valid" models of brain systems (e.g. the visual system) to see if behavioural tendencies observed in humans (and other animals) can be reproduced. When they are reproduced successfully, it can be concluded that the model must mimic to some degree the connectivity of that same system in the brain itself and hence knowledge of the modus operandum of that system has been gained[4]. (Hawson, in press-b).

This dependency on empirical research in the sciences, as well as the conceptualization of the mind as the product of elemental biological processes, then, places connectionism very firmly in the materialist philosophical camp, echoing distinctly the views of Democritus.

The premises of connectionism have, of course, been subject to critical review, not only from dualists, but also from materialists of the non-connectionist persuasion, as well as by some who, while persuaded of the value of connectionism, feel that claims as to its explanatory power may have been overblown. (For example, see Smolensky, 1988; Clark, 1989.) One familiar general criticism of connectionist thinking is that it is too "reductionist": many critics consider that it is unlikely, and for some preposterous, that mental states will be reducible to brain states as knowledge

---

[4] The degree of mimicry that is achieved between connectionist models and actual brain systems, however, is an issue that stirs up considerable debate. It is contended by experimenters that the macroscopic behaviour of a system may be reproduced through judicious adjustments to model parameters (commonly called "tweaking") without mimicking to any significant degree the underlying processes in the brain. Churchland and Sejnowski (1992, pp.130-137), on the other hand, argue that a model net that is highly constrained by biological data can achieve a high degree of overlap with real neural nets in terms of "parameter space". This problem will be discussed in more detail later in this chapter.

of brain functioning increases. It has been pointed out by Patricia Churchland (1986, p.278), however, that the terms "reductionism" and "reductionist" are not used consistently in the literature, and are often employed as epithets of abuse or insult. To be labelled reductionist, however, is not necessarily a negative thing in Churchland's opinion, and so as to construct an accurate framing of the sense in which connectionism is reductionist, she has treated the subject of intertheoretic reduction in detail in her text Neurophilosophy (Churchland, 1986, chap. 7). Her point of view is expressed succinctly below:

> In the sense of "reduction" that is relevant here, reduction is first and foremost a relation between theories. Most simply, one theory, the *reduced theory* T_R, stands in a certain relation ... to another more basic theory T_B. Statements that a phenomenon P_R reduces to another phenomenon P_B are derivative upon the more basic claim that the *theory* that characterizes the first reduces to the *theory* that characterizes the second. (p.278)

Thus statements claiming that a certain phenomenon has been reduced to another phenomenon are in fact statements that the theory underlying the first phenomenon has been reduced to the theory underlying the second phenomenon:

> For example, the claim that light has been reduced to

electromagnetic radiation means (a) that the theory of optics has been reduced to the theory of electromagnetic radiation and (b) that the theory of optics is reduced in such a way that it is appropriate to identify light with electromagnetic radiation. (p.279)

What is important in this elucidation of the issues is the emphasis that it places on the requirement that, for a successful intertheoretic reduction to take place, *both* theories must be "correct"[5], at least as far as current knowledge goes: should either one be even partly erroneous, whether due to false logic, an incomplete information base, or other such inadequacies, completely successful reduction cannot be expected. Indeed, reductions between theories typically involve some revision to the reduced theory thereby creating an analogue of the original. In cases of extreme wrongmindedness, complete elimination of the reduced theory may be mandated.

How much of, or how closely, the old theory is mirrored in the analogue varies from case to case and is a function of how much correction the old theory requires. Most of the theory of optics could be preserved, only some of classical thermodynamics, and almost nothing of the caloric theory of

[5] It is recognized that the correctness of any theory is never absolute as it is a function of the data chosen to form the basis of the theory, and this may be deficient in many ways, out-of-date, ill-chosen, insufficient, imprecise and so on.

heat and temperature. (p.283)

Since one of the major benefits of intertheoretic reduction is explanatory unification, considerations of this sort are crucial in responding to critics of connectionism. Implicit in many of the criticisms is the assumption that any new theory of mental states must be able to explain the conceptualizations of mental states which have emerged from different theoretical perspectives. If, however, the theory to be reduced requires substantial correction, no smooth reduction will, nor should be, forthcoming, and the goal of explanatory unification will not be attained. In such a case, the requirement of cross-theoretical identification of states and suchlike constitutes an unacceptable constraint on the reducing theory. A far more satisfactory situation arises if theories at different levels of inquiry co-evolve by allowing the discoveries in one area to inform and initiate modification in the other, the ultimate goal being to construct a grand unifying theory of the topic under investigation. Happily, co-evolution of theories has already occurred in, for example, the fields of transmission and molecular genetics, and there is evidence of mutual influencing in research on memory carried out in the fields of neurobiology, neurology, and psychology (see P. S. Churchland, 1986, p. 285). For such co-evolution to develop, however, there must be an openness on the part of theorists (as well as experimentalists, philosophers, linguists, etc.) to the possibility of error within the premises of their espoused theory, hence, it is argued by connectionists, for researchers into the nature of mind to insist that our intuitive understanding of mental states and processes cannot be in error is to beg the very question being examined. (See also Searle, 1992, pp.144-149 for a discussion of this point.) Furthermore, inherent in the triadic relationship of the disciplines which form the root theories of connectionism, the main goal of which is to arrive at a "unified science of the mind/brain" (P. S. Churchland, 1986, see subtitle of text), is the potential for the emergence of such a unifying theory through co-evolution of theoretical progress in the fields of neuroscience, artificial intelligence, and cognitive psychology. Nevertheless, connectionists are often to be found in the eliminative materialist philosophical camp (for example Stich, 1988; P. M. Churchland, 1989, pp. 125-127) which, it may be remembered, favours the elimination of erroneous theories rather than their reduction. For these philosophers, folk psychology, "our commonsense conceptual framework for mental phenomena" (Churchland, P. M., 1989, p. 2), is merely a rough and ready, subjective, and culturally specific theory of mental processes, which may be proven by advancing neuroscientific discovery to be utterly false. It is on this stance, then, that connectionists are often taken to task (see for example Searle, 1992, chap. 1).

One of the most prominent critics of the eliminative materialist stance of connectionism is John Searle. He has argued (Searle, 1992) that elimination of folk psychology through intertheoretic reduction is implausible because

> [f]olk psychology includes such claims as that people sometimes drink because they are thirsty and eat because they are hungry; that they have desires and beliefs, that some of these beliefs are true, or at least false; that some beliefs are better supported than others; that people sometimes do something because they want to do it; that they sometimes have pains; and that these are often unpleasant. And so - more or less indefinitely - on.... If the theory goes, the theoretical entities go with it: to demonstrate the falsehood of folk psychology would be to remove any justification for accepting the existence of the folk psychological entities. (p. 6).

As all of the claims detailed above are surface phenomena of conscious intelligence in humans, Searle is also maintaining in his line of argumentation that consciousness will not be reducible to brain states. In his opinion, however, the irreducibility of consciousness is inconsequential, or at least "has no deep consequences" (p. 118), because it "is a trivial consequence of the pragmatics of our definitional practices" (p.122). Taking the case of the "reduction" of the phenomenon of heat to the kinetic energy of molecular motion, Searle argues that in such intertheoretic reductions the surface or subjective appearances of the phenomenon are explained as being caused by the behaviour of some underlying microstructure. When this explanation is accepted as plausible,

> [w]e then *redefine* heat ... in terms of the underlying causes of both the subjective experiences and other surface phenomena. And in the redefinition we eliminate any reference to the subjective appearances and other surface effects of the underlying causes. "Real" heat is now defined in terms of the kinetic energy of the molecular movements, and the subjective feel of heat that we get when we touch a hot object is now treated as just a subjective appearance caused by heat, as an effect of heat. It is no longer part of real heat. (p. 119).

As Searle explains, however, this does not eliminate the actual experience of heat as a valid phenomenon. In the case of consciousness, the formula outlined above of carving off the subjective experience of the phenomenon from its "reality" does not work and cannot work according to Searle for the simple reason that where consciousness is concerned the reality *is* the "appearance":

> Consciousness fails to be reducible, not because of some mysterious feature, but simply because by definition it falls outside the pattern of reduction that we have chosen to use for pragmatic reasons. Pretheoretically, consciousness, like solidity, is a surface feature of certain physical systems. But unlike solidity, consciousness cannot be redefined in terms of an underlying microstructure, and the surface features then treated as mere effects of real consciousness, without losing the point of having the concept of

consciousness in the first place. (p. 123)

Should these standard patterns of reduction change as a result of a major intellectual revolution, however, Searle can envisage the possibility of a presently inconceivable reduction of consciousness coming about. Finally and unexpectedly, however, in musing on various approaches to representing the processing of information by the brain, Searle arrives at a defense of connectionism:

> Among their other merits, at least some connectionist models show how a system might convert a meaningful input into a meaningful output without any rules, principles, inferences, or other sorts of meaningful phenomena in between. This is not to say that existing connectionist models are correct - perhaps they are all wrong. But it is to say that they are not all obviously false or incoherent in the way that the traditional cognitivist models that violate the connection principle are. (p. 246)

The subject of the difference between "traditional cognitivist models" and connectionist models highlighted by Searle is an important one which has been explored in fine detail by Paul Smolensky.

In a landmark paper which stirred up enormous controversy (see Smolensky, 1988, pp. 23-74), Smolensky, who is himself a researcher and advocate of the connectionist approach to modelling cognition, examined connectionism with a view to attempting to "articulate the goals of the approach, the fundamental hypotheses it is testing, and the relations presumed to link it with the other theoretical frameworks of cognitive science" (Smolensky, 1988, p. 2). He labels the formulation of the connectionist approach which he advocates the PTC approach: the Proper Treatment of Connectionism. Though a believer in connectionism, his beliefs about connectionist models are not uncritically weighted towards the positive. In fact, in the introductory passages of his famous paper he lays out an almost equal number of positive and negative assessments of connectionist modelling in an even-handed and contradictory fashion. For instance, Smolensky (1988, pp. 2-3) reports that, among other things, he believes

> (1) a. It is far from clear whether connectionist models have adequate computational power to perform high-level cognitive tasks....
>
> (1) c. It is far from clear that connectionist models can contribute to the study of human competence....

while at the same time believing

> (1) g. It is very likely that connectionist models will turn out to offer contributions to the modelling of human cognitive

performance on higher level tasks that are at least as significant as those offered by traditional symbolic models.

(1) i. It is likely that connectionist models will offer the most significant progress of the past several millennia on the mind/body problem.

One of Smolensky's major interests in carrying out this detailed analytical investigation of connectionism was to establish the "level of analysis" at which connectionist models of cognition operate. In his opinion, this is neither at the neural level, as has been proposed by many, nor at the "symbolic" level adopted by traditional cognitive models of cognition, but at a so-called "subsymbolic" level somewhere between the two. Defining the traditional approach to cognitive modelling as the *symbolic paradigm* (details of which will be discussed below), he goes on to explain that in the PTC approach to connectionism the correct perspective to adopt is that of the *subsymbolic paradigm* which "is intended to suggest cognitive descriptions built up of entities that correspond to *constituents* of the symbols used in the symbolic paradigm; these fine-grained constituents could be called *subsymbols*, and they are the activities of individual processing units in connectionist networks" (p. 3). As Smolensky points out, the models emanating from both the symbolic and subsymbolic paradigms can be analyzed at various levels, but he suggests that the preferred level of analysis for models of the symbolic paradigm is the *conceptual* level, while that of the subsymbolic is the *subconceptual*. Since the level of analysis adopted "greatly constrains the formalism for analysis" (p. 3), and since "[p]robably the most striking feature of the connectionist approach is the change in formalism relative to the symbolic approach" (pp. 3-4), Smolensky goes to great lengths to elaborate how working at the subsymbolic level alters formulations of cognitive performance. He notes that:

> Since the birth of cognitive science, *language* has provided the dominant theoretical model. Formal cognitive models have taken their structure from the syntax of formal languages, and their content from the semantics of natural language. The mind has been taken to be a machine for formal symbol manipulation, and the symbols manipulated have assumed essentially the same semantics as words of English. (p. 4)

The structure of language, then, is the basic structure underpinning the symbolic paradigm. The basic structure for connectionist models, however, is a *dynamical system* in which "the state of the system is a numerical vector evolving in time according to differential evolution equations" (p. 6). At any moment in such a system, the state vector (which determines the behavioural output) is represented as a pattern of activity scattered over a wide distribution of units, each of which computes its value from inputs received from other units before passing on information with respect to its activity to other units in the system. Thus, the meaning of a unit's activity value in a connectionist model cannot be likened to an element of a natural language, such as a word, as in a traditional cognitive model: because many units are active in the representation of any concept, the units themselves do not have conceptual semantics, but subconceptual. In other words, the activity value of a unit has no recognizable meaning to us, a point which refers back to Searle's argument above which casts doubt on the notion of the existence of "meaningful phenomena", or "universal grammars" of various kinds, which intercede causally between meaningful input and meaningful output. Higher level or "hard" behaviours such as "rule following" in a language are better described, according to this view, as emergent properties of an inherently "soft" system in which no such rules exist. In connectionist subsymbolic models, then, to try to understand the causal factors for high, conceptual-level behaviour through subconceptual level descriptions demands a "dimensional shift" which is not necessary in traditional models involving the symbolic paradigm, one which moves the semantics of the description away from language, and into the realm of mathematics.

As was mentioned above, though the level at which connectionist models function is lower than the symbolic level, it is still higher than the neural level despite the fact that much of what has been learned about brain processing has been incorporated into model designs. This lack of "neural realism" has been another source of criticism of connectionist models. There is no doubt, however, that the functioning of connectionist models is much closer to the neural level than that of symbolic models, and some connectionist models which are designed specifically to test theories of brain functioning of more fundamental abilities such as visual functioning come very close to being "neural models". According to Smolensky, neural models of cognitive behaviour are not yet feasible, not for want of data on the brain, but for want of the right kind of data:

> Our information about the nervous system tends to describe its structure, not its dynamic behavior. Subsymbolic systems are dynamical systems with certain kinds of differential equations governing their dynamics. If we knew which dynamical variables in the neural system for some cognitive task were the critical ones for performing that task, and what the "equations of motion" were for those variables, we could use that information to build neurally faithful cognitive models. But generally what we know are endless static properties of how the hardware is arranged. (p. 10)

Thus, only a model of the brain which provides a description at the neural level can legitimately be called a neural model. Nevertheless, connectionist models can provide approximate descriptions of the functioning of the brain at a higher than neural level, that is, at a subconceptual level, and though approximate, the description of the microstrucure of cognition that they offer is more precise by far than that of symbolic models. Smolensky clarifies the relationship between symbolic and subsymbolic models of cognition as follows:

> Subsymbolic models accurately describe the microstructure of cognition, whereas symbolic models provide an approximate description of the macrostructure. An important job of subsymbolic theory is to delineate the situations and the respects in which the symbolic approximation is valid, and to explain why (p. 12).

In other words, the job of subsymbolic theory is to try to reduce symbolic theory by explaining at a more fundamental level how symbolic phenomena arise, and in the process to correct any apparent errors. Churchland and Sejnowski (1992), however, respond to "lack of realism" criticisms of connectionist models in a different manner: they question not just the practicality but also the desirability of realism in a model by noting that in the scientific tradition no model is 100% realistic, and moreover, need not be 100% realistic in order to be a good and useful model. Speaking of models of brain function in general, they comment:

> Realistic one-upmanship needs to be put in perspective. First, models that are excessively rich may mask the very principles the models were built to reveal. In the most extreme case, if

> a model is exactly as realistic as the human brain, then the construction and the analysis may be expensive in computational and human time and hence impractical. ... modelling may be a stultifying undertaking if one slavishly adheres to the bald rule that the more constraints put in a model, the better it is. Every level needs models that are simplified with respect to levels below it, but modelling can proceed very well in parallel, at many levels of organization at the same time, where sensible decisions are made about what detail to include and what to ignore. There is no decision procedure for deciding what to include, though extensive knowledge of the nervous system together with patience and imagination are advantageous. The best directive we could come up with is the woefully vague rule of thumb: make the model simple enough to reveal what is important, but rich enough to include whatever is relevant to the measurements needed. (p. 137)

A crucial insight into the debate surrounding the modelling of cognition may have been provided by Clark (1989) who maintains that there are two types of cognitive science: *descriptive cognitive science* which "attempts to give a formal theory or model of the structure of the abstract domain of thoughts, using the computer program as a tool or medium" (p.153), and *causal cognitive science* which "attempts to give an account of the inner computational causes of the intelligent behaviours that form the basis for the ascription of thoughts" (p. 154). Connectionist or, more generally, subconceptual models are undoubtedly causal, while traditional or symbolic models are more descriptive. The importance of the distinction made by Clark lies in the light it sheds on the futility of comparing the performances of models of cognitive behaviour which have arisen out of the different divisions in cognitive science: since they have entirely different goals they are hardly in competition with one another. In spite of the fact that both Clark and Smolensky are convinced of the value of connectionist models of behaviour, they do not believe that higher level symbolic models of cognition have nothing to offer to the study of cognition. Though "[a]ccording to the subsymbolic paradigm, the unity underlying cognition is to be found not at the conceptual level, but rather at the subconceptual level, where relatively few principles in a single formal framework lead to a rich variety of global behaviours" (Smolensky, 1988, p. 22), Clark foresees the possibility that "we may view the constructs of symbolic AI [Artificial Intelligence], not as mere approximations to the connectionist cognitive truth, but as a means of highlighting a higher level of unity between otherwise disparate groups of cognitive systems" (Clark, 1989, p. 200).

And so the debate as to the nature of conscious intelligence and how best to investigate it rages on and (to borrow Searle's phrase) more or less indefinitely on[6]. In the meantime, as the amount of data on brain processing has increased, the sophistication and power of connectionist models of brain function has evolved accordingly: researchers have constructed highly realistic working models of brain systems such as the visual system, which they lesion to simulate frequently occurring insult to the brain and then observe how the model recovers its function in order to gain insight into the recovery of the brain itself (Kosslyn et al., 1990a, 1990b); models which have been created to study the mechanics of spatial representation in the visual system have discovered that the response properties of units in the model mimic the response properties of neurons in the visual system of the brain (Anderson & Zipser, 1990); language acquisition is being modelled (Plunkett & Marchman, 1991), and robots with senses such as vision, touch and kinesthesia are being devised to study how perceptual categorization is built up in the brain (Edelman & Reeke, 1990)[7]. But, it may be protested, how much of this is truly relevant to the study of consciousness? What insight into consciousness is gained by modelling the visual system? How is visual perception related to the central problems being investigated? The short answer is that nobody really knows. No one has yet discovered a "seat of consciousness" in the brain, nor by

---

[6] For a recent up-date on the issues, particularly the role of science in the study of consciousness, see Horgan, 1994.

[7] Of major concern to critics of this type of modelling is the fact that they are more in the business of mimicry than prediction of outcome, as is more normal in the scientific tradition.

some estimates are they likely to do so, since much evidence now points to the fact that awareness, which according to Crick (1994) is synonymous with consciousness, is represented in the brain in a widely distributed fashion (see for example, P. S. Churchland, 1986, chap. 5; Crick & Koch, 1992). In the opinion of Crick and Koch (1992), however, to be able to explain *visual* awareness or consciousness, a well-studied phenomenon compared to other aspects of consciousness, must be to have solved at least a piece of the puzzle. And what has been learned about visual awareness so far highlights the complexity of even this reduced undertaking. Crick and Koch (1992) note:

> Although the main function of the visual system is to perceive objects and events in the world around us, the information available to our eyes is not sufficient by itself to provide the brain with its unique interpretation of the visual world. The brain must use past experience (either its own or that of our distant ancestors, which is embedded in our genes) to help interpret the information coming into our eyes. (pp. 153-154)

Hence, if it is conceded that "learning and memory are at the dead center of cognition, if anything is" (P. S. Churchland, 1986, p. 151), it must be recognized that the dependence of visual awareness on experience or learning means that the study of consciousness is not separable from the study of cognition. Furthermore, it has become evident that, for forms of awareness other than the most fleeting, attention is necessary, and, of course, attention has long been known to interact with the process of encoding memories (see Crick, 1994, chap. 5). Thus, the study of consciousness mandates a knowledge of the brain's attentional mechanisms, as well as a knowledge of how memories are encoded, presently subjects of intensive research in information processing and the neurosciences. It is for this reason, that Crick criticizes the conservative attitudes of scientists who carry out research in these areas but who refuse to tackle the problem of consciousness, preferring to consider it a philosophical matter. In Crick's opinion,

> [i]t will no longer do to study in detail some particular aspect of vision and at the same time to ignore the overall question: What happens in the brain to make us see. A layman would regard such an attitude as excessively narrow-minded and he would be right. As I have tried to show, the problem of visual awareness is approachable now, both experimentally and theoretically. What is more, by actively confronting the problem, one begins to think in new ways; to ask for information (e. g., on dynamic parameters or short-term memory) that had previously appeared to be irrelevant or of little interest. I hope that before long every laboratory working on the visual system of man and the other vertebrates will have a large sign posted on its walls, reading: *CONSCIOUSNESS NOW*. (Crick, 1994, p. 253).

In advocating the pursuit of consciousness, from a "bottom-up" scientific perspective, Crick is echoing the call of connectionists for the inclusion of scientific theorizing in the solving of complex problems such as consciousness and cognition which, in the past, have been seen as outside the realm of scientific inquiry at its most fundamental levels. That such calls are somewhat overenthusiastic at times is understandable given the resistance of the disciplines which have long considered these topics to belong exclusively in their territory.

Examination of the problems surrounding the study of consciousness from both top down and bottom up perspectives, therefore, have shown, if nothing else, that consciousness is an exceedingly complex phenomenon, and this inherent complexity of consciousness has been well captured by one particular "definition" cited by Crick (1994) at the beginning of his text:

> consciousness. The having of perceptions, thoughts, and feelings; awareness. The term is impossible to define except in terms that are unintelligible without a grasp of what consciousness means. Many fall into the trap of equating consciousness with self-consciousness - to be conscious it is only necessary to be aware of the external world. Consciousness is a fascinating but elusive phenomenon; it is impossible to specify what it is, what it does, or why it evolved. Nothing worth reading has ever been written on it.

Stuart Sutherland,
The International Dictionary of Psychology

In the face of the elusiveness, complexity and mystery of the true nature of consciousness, it would appear reasonable to stalk it cautiously but determinedly along several fronts, as was suggested by Clark and Smolensky, including those currently being explored by connectionism. Connectionism is undoubtedly interested in the causal aspects of cognitive behaviour and, since cognition is intimately intertwined with consciousness, it should be able to fill in some pieces of the overall picture.

Just as Crick narrowed the focus of his search for an understanding of consciousness to visual consciousness, in attempting to understand the fine structure of cognition, another complex phenomenon, it would seem prudent to target only one aspect of cognition in order to increase the chances of reaching some satisfactory conclusions as to its nature and its workings. For educators, the nature of the learning processes of the brain are of the utmost importance, and since this dissertation is exploring the relationship between second language learning and academic achievement, of central concern here will be how well connectionism describes the learning process in human brains, particularly, but not exclusively, as this is related to second language learning.

As has been noted above, "connectionism" has been considered synonymous in the past with "parallel distributed processing". It has also been confused with "neural net modeling", another information processing framework. Recently, Patricia Churchland and Terrence Sejnowski have injected some order into this muddle by clarifying how the meanings of these terms have evolved with time:

> Coined more recently, "connectionism" usually refers to modeling with networks that bear only superficial similarities to real neural networks, while "neural net modeling" can cover a broad range of projects. Ironically perhaps, "neural net modeling" is usually identified with computer modeling of highly artificial nonneuronal networks, often with mainly technological significance such as medical diagnoses in emergency wards. "PDP" ("parallel distributed processing") is generally the preferred label of cognitive psychologists and some computer scientists who seek to model rather high-level activity such as face recognition and language learning rather than low-level activity such as visual motion detection or defensive bending in the leech. (Churchland & Sejnowski, 1992, p. 6)

These authors prefer to call the approach they take to modelling information processing in the brain *computational neuroscience*, an approach which "draws on both neurobiological data and computational ideas to investigate how neural networks can produce complex effects such as stereo vision, learning, and auditory location of sound-emitting objects" (ibid., p.6). This take on information processing in the brain is very much in keeping with that detailed by Paul Churchland (1989), who has described connectionism as "a recently developed theoretical and experimental approach to the phenomena of human cognition that is at once (a) naturalistic, (b) reductionistic, and (c) capable of explaining both the radical plasticity of human consciousness, and its intricate dependence on the extended cultural surround" (P. M. Churchland, p. 130), and it is in this sense and no other that the terms "connectionism" and "connectionist" should be interpreted throughout this dissertation.

The naturalism and reductionism of connectionism come from setting down its roots in neuroscientific research on plasticity or learning[8] in the brain, human and animal (though it is recognized that, as argued above, connectionist models of brain function operate at a level above the neural level in general); connectionism's ability to explain the radical plasticity or learning capability of the human brain, on the other hand, is dependent on the astuteness of its proponents in abstracting accurately from the microscopic detail of neuroscientific discovery

[8] Due to the fact that, as Patricia Churchland (1986) has noted, "[t]he general category of learning has fragmented into a variety of kinds of process" (p. 151), the term learning is now commonly replaced by the less theoretically burdened term plasticity in many neuroscientific papers. In this thesis, however, "plasticity" as it refers to brain function will be taken to be synonymous with "learning."

basic structural principles which can be incorporated into the designs of the models, as well as basic principles of learning which can be described by computational mathematical equations. To what extent connectionists have been successful in depicting the learning processes in the brain will be examined below.

2.2 An Examination of Connectionism with the Object of Classifying It Appropriately

As was discussed at the end of the previous section, the definition of connectionism which will be adhered to throughout this dissertation is that of Paul Churchland (1989, p. 130; also see glossary). Churchland has termed connectionism a "new approach [to the phenomena of cognition] ... which resides at the interface of computational neuroscience, cognitive psychology, and artificial intelligence" (p. 130). Although computational neuroscience is mentioned as only one of the disciplines feeding into connectionism, Churchland's description of connectionism is seen as having much in common with the definition of "computational neuroscience" adopted by Patricia Smith Churchland and Terrence Sejnowski (1992) which they term "an evolving approach that aims to discover the properties characterizing and the principles governing neurons and networks of neurons" (p. 6). The understanding of connectionism outlined in these descriptions should not be confused, therefore, with notions which perceive connectionism as "modeling with networks that bear only superficial similarities to real neural networks" (ibid.). In the context of this dissertation, the critical feature of connectionism as detailed by Churchland's definition is its ability to describe the plasticity (learning capability) of the human brain.

The definition chosen refers to connectionism as an "approach."

This is consistent with the terminology used by Churchland and Sejnowski, "an evolving approach", in referring to computational neuroscience, as well as with that of Clark (1989, pp. 118-119) when discussing "PDP [parallel distributed processing] approaches" to modelling human cognition[1]. If we take an "approach" as meaning "the method used or steps taken in setting about a task" (see glossary), connectionism could then be considered to be "a method taken to set about studying the phenomena of human cognition, including learning." In other words, connectionism is perhaps best described as a methodology*, "a set or system of methods, principles, and rules used in a given discipline." This classification is in keeping with Paul Churchland's (1984) terminology in detailing the intricacies of the "methodological problem." He includes the "cognitive/computational approach" (arising from cognitive psychology and artificial intelligence) and "methodological materialism" (based on examination of the fine-grained structure of nervous system function) as different perspectives in dealing with this problem. Given that connectionism amalgamates these perspectives, it seems entirely reasonable to suggest that it also falls under the jurisdiction of the methodological problem. Hence, it would appear that it is appropriate to classify connectionism as a methodology for the study of human cognition.

---

[1] It may be remembered that Clark equates parallel distributed processing with connectionism (Clark, 1989, p. 83).

The premises upon which connectionism is constructed are instantiated in models, whether theoretical or empirical, of particular phenomena of human cognition, for example recognition and understanding, or learning, where each model is "a simplified representation of a system or phenomenon with any hypotheses required to describe the system or explain the phenomenon" (see glossary). As was argued in the previous section, connectionist models are simplified representations of the nature and principles governing neurons and neural networks within particular neural systems, and the working hypothesis* underpinning these models is

> that emergent properties are high-level effects that depend on lower-level phenomena in some systematic way. Turning the hypothesis around to its negative version, it is highly improbable that emergent properties are properties that cannot be explained by low-level properties ..., or that they are in some sense irreducible, causally *sui generis*, or as philosophers are wont to say, "nomologically autonomous", meaning roughly, "not part of the rest of science." (Churchland & Sejnowski, 1992, p. 2)

Thus, it would seem appropriate to describe connectionism as a *methodology* which gives rise to a particular class of *models* of cognitive function, all based on a single working hypothesis about how the brain works.

It will be in its capacity to effectively model how the brain learns that connectionism will be examined here.

2.3 An Examination of the Validity of Connectionist Models of Learning

2.3.1 Biological constraints on connectionist learning procedures

The main goal of connectionist researchers was, and still is, to build more biologically plausible models of learning in brains, human and animal. In order to do so, it is necessary to take into consideration essential elements of the brain's structure and function, which neuroscientific studies have elucidated to an astonishing degree of precision due to ever improving technological advances in investigatory techniques, and to incorporate these elements into the design of the models. The wealth of information available on the brain, however, constitutes a problem for connectionist modellers, since it would be impossible, and perhaps even counterproductive as Churchland and Sejnowski (1992, p. 137) have commented, to include every known detail. The questions of which details to incorporate into the model, and which to either ignore or to represent in abstract form, are the first hurdles that modellers have to clear. The decisions made on these matters, of course, depend on what level of organization of brain function the model is aiming to elucidate[1]. As Churchland and Sejnowski have noted "[a] model that captures the salient features of learning in

---

[1] Churchland and Sejnowski (1992, p. 11) have identified seven levels of organization in the brain, ranging in spatial scale from the molecular, to the synaptic, neuronal, network, map, systems, and finally to the CNS (central nervous system).

networks will have a different face from a model that describes the NMDA channel[2]" (Churchland & Sejnowski, 1992, p. 11). These authors also point out, nevertheless, that a model aiming at one level must take into account information available at higher and lower levels in order to maintain consistency with the overall picture of brain function, especially as it relates to the phenomenon being modelled. Most connectionist models are models of neural networks, and so details of the fundamental elements of neural nets, neurons, are usually taken as the starting point in the designs. However, as connectionist researchers readily admit, "[b]ecause our knowledge of neurons is incomplete and our computing power is limited, our models are necessarily gross idealizations of real networks of neurons" (Hinton, 1992, p.145). That is to say, much but not all of what is known about neurons is incorporated into connectionist models. To appreciate to what extent this is so, it is necessary to detail the structural and functional aspects of neurons which are most salient to connectionist designs.

Structure and function of neurons

The typical *neuron* in the human brain is an excitable cell which consists of a *cell body* out of which have grown a *dendritic arbour* and an *axon*: the dendritic arbour (or simply dendrites) is a host

[2] The NMDA "channel" is more properly termed the NMDA receptor. It is a ligand-gated ion channel which is implicated in certain learning processes and is studied at the molecular level.

of fine structures which collect incoming electrical signals[3] from other cells, and the axon is usually a long thin strand, branching profusely at its tip, which carries the signal from the cell body to other neurons. At the end of each axonal branch is a structure called a *synaptic terminal* or *bouton* which makes contact, though not usually physically, with the dendrites of other neurons. A *synapse* is a functional contact between neurons, and consists of a *presynaptic* terminal, a narrow gap, and a *postsynaptic* site capable of receiving the signals from the presynaptic terminal, signals which are generally passed on by chemicals called *neurotransmitters*[4]. The postsynaptic site is most commonly a *dendritic spine*, a small outgrowth on a dendrite, but can also be a small portion of the cell body. When it becomes sufficiently excited by incoming signals, and this is determined by the specific threshold level of activation of the cell, the neuron "fires", that is, it sends a self-sustaining, regenerative "spike" of electrical activity down its axon and through its synapses to as many neurons as it has contacted. A generally accepted estimate of the number of synaptic connections a neuron makes is 1000, however, in some cases it may be as many as 10,000 (see Churchland & Sejnowski, 1992, p. 51.) Hence, the information that the neuron communicates to others

---

[3] The "electrical signals" collected by dendrites, and passed on by other neurons, are chemical in nature and are caused by the inflow and outflux through the cell membrane of charged particles called ions.

[4] Electrical synapses are exceptions to this general picture. They do make physical contact between the cells, allowing a bidirectional flow of ions.

about its *state* or *activity* (encoded in its rate of firing) is passed on in a *highly distributed* fashion, some of it locally, that is, to neurons in its close vicinity, but most of it to neurons in other areas of the brain. Forward projections to an area are usually equal in number to backward or recurrent projections to the neuron from that area, so the "conversation" between connected neurons is two-way. The message of a particular neuron can have either an excitatory effect on another, or an inhibitory effect. When the effect is excitatory, the probability that the postsynaptic neuron will fire is increased; when inhibitory, that probability is decreased. Which effect it will have is dependent on which particular receptor is reached by the neurotransmitter the neuron employs as its messenger. In most cases, the effect of any one neuron on another, however strongly connected, is small, and it typically takes a convergence of inputs from many neurons in both proximate and distant parts of the brain to make a neuron fire. This arrangement constitutes a safeguard of sorts for the function of the neuron, because it means that the loss of some connections will not result in a great reduction (or increase) in its rate of firing or activity.

Learning occurs in the brain when there is a change in the ability of one neuron to communicate effectively with another, that is, when there is a change in the strength of synaptic connectivity between neurons. The strength of a synaptic connection is often referred to as the *synaptic weight*. At the synaptic level, learning can be manifested in different ways depending on whether it takes place on a conscious plane (explicit learning) or an unconscious plane (implicit learning) (see Kandel & Hawkins, 1992), and whether it takes place during development (see Shatz, 1992), though it is doubtful that these processes can be cleanly separated "because learning and development almost certainly share mechanisms and principles" (Churchland & Sejnowski, 1992, p.293). No matter what the specifics of the change in efficacy, learning causes the activity of the connected neurons to become either more, or less[5], closely correlated. As mentioned above, in the case of an excitatory message, a strengthening of synaptic contact between two neurons will result in a higher probability that, when the first neuron fires, the second will too. This will result in an increase in the correlation of firing between the two neurons. A weakening of synaptic connections will result in a correspondingly weakened correlation in the firing. The reverse is true, of course, for an inhibitory connection. However, in one important theoretical model of synaptic change, it has been suggested that inhibitory synapses do not undergo modification (Bear et al., 1987). According to Crick (1994, p. 209), this strengthening in synaptic connections can be considered a kind of "binding" (though certainly not the only kind) between neurons, and the stronger the synaptic connections, the stronger the bond. It is through this kind of binding of activity,

---

[5] For example, there are incidences of anti-Hebbian learning where correlated activity at the pre- and post-synaptic sites results in a decrease in synaptic strength. (See Churchland & Sejnowski, 1992, pp. 252-253)

then, that networks, and even systems, of neurons are formed, and once formed, these networks have their own governing principles.

Organization at the neural network and map levels

Undoubtedly, an understanding of the general characteristics of neuronal function and structure is critical to understanding brain organization at the cellular level. Nevertheless, other levels of cerebral organization are equally important to comprehending how the brain processes information. For example, within many sensory and motor systems, a major principle of organization is the *topographic map*, the key feature of which is that adjacent neurons have adjacent receptive fields. That is to say, adjacent neurons process similar representations, a situation that appears to be brought about by competitive interaction between neurons, and which may have evolved as a strategy for saving and sharing neuronal "wiring." In the primary visual cortex, to give but one well-known example, neurons are arranged topographically such that neighbouring neurons collectively form a map of the retina. However, all areas of the retina are not accorded equal representation within this map: the fovea, for instance, covers a relatively large area of primary visual cortex. Thus the map of the retina reproduced in this area of the cortex is distorted with respect to the original. Likewise, the hands occupy a relatively large area of somatosensory cortex in humans and primates. Both the fovea and the hands have a higher than normal density of receptive neurons which have small receptive fields, and since

> the amount of cortex activated by a stimulus in a given receptive field is roughly constant, whatever the peripheral area.... [t]his means that when a body part is innervated by many neurons, each of which has a small receptive field, that body part will appear magnified on the cortical map relative to an area in which there are fewer neurons, each of which has a large receptive field. (Churchland, P. S., 1986, p. 129)

It is because of this high density of receptive neurons that both the fovea and the hands are extremely important informational sources for their respective sensory modalities.

Thus, the general rules which seem to govern the allocation of territory in topographic mapping are (a) neighbouring sensory sites maintain their spatial proximity in the cortex, and (b) the greater the amount of information a sensory area provides, the greater the amount of territory it is allocated. Once formed, these topographic maps are relatively stable, but retain a degree of plasticity throughout life, such that the responses of the cells in a particular territory can alter to accommodate changed sensory circumstances, an alteration in the intensity of sensory stimuli, or perhaps nerve damage, the net effect of which is a territorial shift in the cortex (see Merzenich et al., 1990). According to Churchland and Sejnowski (1992), that the mechanisms by which this rearrangement occur are still a matter for conjecture does not diminish the importance of the implications of this proclivity of neurons to exhibit plastic behaviour:

> It is not yet known how much of this rearrangement is due to plasticity in cerebral cortex, or perhaps in subcortical areas that project to cortical maps.... Nonetheless, this evidence, and further evidence for synaptic plasticity ... make it difficult to think of the machinery in the adult brain as "hardwired", or static. Rather, the brain has a remarkable ability to adapt to changes in the environment, at many different structural levels and over a wide range of time scales. (p. 34)

This adaptability of neurons to changed circumstances is further evidence that networks of neurons are not only able to learn, but to relearn. This is a property that it is important to reproduce in any biologically valid model of learning.

As well as topographic organization, many areas of cerebral cortex exhibit *laminar organization*. Several non-cortical anatomical structures have also been shown to display striation (laminar organization), the superior colliculus, the cerebellum, the hippocampus, and the lateral geniculate nucleus being a few of the most significant. The laminae or layers are caused by the segregation of different cell types and of cell processes, though no one layer is completely homogeneous in cell type. A given lamina, however, shows high conformity with respect to where it receives input from, and where it projects its output to. For example, in the cerebral cortex, there are typically six laminae, and layer 4, the middle layer, receives specific sensory input from the thalamus, while output to subcortical motor areas comes from layer 5, and layer 6 mainly projects back to the thalamus.

To complement this horizontal organization, there is vertical organization in the cerebral cortex in the form of *columns* of cells traversing layers which reflect anatomical and physiological commonalities: cells in a particular column show rich interconnectivity and similar response properties. Neurons in the same column of the visual cortex, for instance, might all show a preference for a stimulus with a particular orientation. However, boundaries in response properties between columns are rarely sharp (ocular dominance columns in layer 4 of primary visual cortex being one notable exception), hence the term "vertical column" may be somewhat misleading.

A word on brain "design"

These different organizational modes may have a common genesis, according to Churchland and Sejnowski (1992):

> Topographic mapping, columnar organization, and laminae are special cases of a more general principle: the exploitation of geometric properties in information processing design. Spatial proximity may be an efficient way for biological systems to

> assemble in one place information needed to solve a problem. To consider a simple case, suppose it is necessary to compare differences between stimuli at neighbouring locations, where comparison requires signals be brought together. Then topographic organization may achieve this efficiently while minimizing the total length of the connections. (p. 37)

One feature of brain organization which appears to emerge from consideration of its substructure, then, is its pragmatism. According to the above analysis of its structure, the brain apparently seeks out economical solutions to practical problems, as would be expected of good engineering design. But can the brain be accurately portrayed as an example of good design, or is it more of a "kludge", as Clark (1989, p. 69) suggests? A kludge in Clark's definition is

> a term used in engineering and computer science to describe something that, from a pure (i.e., ahistorical), design-oriented viewpoint, looks messy and inefficient. But it gets the job done. And it may even count as an elegant solution once all the constraints (e.g., the available skills and resources) are taken into account. (p. 69)

Most brain researchers, including Churchland and Sejnowski, agree that indeed the brain is more of a kludge than a well-designed organ, and that its kludge-like character is a result of its having evolved from simpler structures which themselves evolved in accordance with constraints, both external (environmental) and internal (materials and structures available to build upon), as stable units capable of carrying out tasks necessary for the survival of the host organism[6] (For an overview of brain evolution, see P. M. Churchland, 1984, chap. 7). As Clark (1969) notes,

> [c]omplex biological systems, then, have evolved subject to the constraints of gradualistic holism.... Gradualism requires that each structural step in the evolutionary process involve only a small adjustment to the previous state. Jacob compares evolution to a tinkerer who must use whatever is immediately at his disposal to achieve some goal. This case contrasts with that of an engineer who, within certain limits, decides on the appropriate materials and design, gathers them, and then carries out the project... The point, then, is that what the tinkerer produces is heavily dependent on his historical situation in a way in which the engineer's product is not. (p. 70)

Evidence of the "tinkering" progress of human evolution can be found in the eye, the light sensitive photoreceptors of which face away from the light, so that light must first pass through overlaying layers of nerve cells and their processes before it

---

[6]According to Clark (1992, p. 70), this may be true down to the level of such fundamental structures as mitchondria and chloroplasts which now are integral components of the eukaryotic cell.

reaches the photoreceptors, not what one might call an optimal engineered solution to the problem of constructing a viewing device.

Parallelism in information processing in the brain

It is on consideration of this gradual construction of complex functions from autonomously developed sub-functions that the massive parallelism of information processing in the brain becomes comprehensible. Information gathering systems such as the visual system have come together in a piecemeal fashion from already functioning sub-systems, the ability to analyse form, to detect motion, to process colour, all of which may at one time have operated simultaneously and separately, but which have co-evolved in humans to form one complex whole, so that today, each sub-function still maintains its integrity of processing while admitting integration of information from the others to influence its output (see Zeki, 1992). Nor is sensory processing alone among complex human abilities to depend on "lower level" competencies: the comprehension and production of language, it would appear, is dependent on specialized sequencing regions in the left hemisphere, which may be equally important to the production of ballistic movement, that is, throwing and such, and indeed may have evolved to serve this purpose (see Calvin, 1994).

One of the main benefits of parallel processing of information to an organism is speed in production of output. Though the human brain may be considered slow at some high level tasks such as playing bridge or mental computation when compared to the speed of a conventional digital computer, it accomplishes extremely complex tasks like visual recognition astonishingly quickly, especially given the time constraints imposed by neuronal transmission and cerebral architecture. Parallel processing of information gives a rational account of how this comes about. But it is important to remember that parallel processing of information was not designed to speed up processing, although the fact that it did was almost certainly crucial to its selection by evolution. Speed of processing is an emergent property of parallel processing of information, a result rather than a goal of cerebral architecture, in much the same way as "spandrels [in Saint Mark's Cathedral in Venice] are a necessary structural by-product caused by mounting a dome on a number of rounded arches" (Clark, 1989, p. 77). That such emergent speed be exploited by the organism which has acquired it is only natural. This point is important to keep in mind when considering how complex cognitive capabilities came into being; they also may very well be emergent properties of the architectural features which underpin them, and if so, an understanding of the fine-structure of these features becomes of primary concern to an understanding of the "higher level" ability. (See Clark, 1989, chap. 4 for a detailed argument along these lines.)

Additional considerations

There are some higher order biological constraints that it is important to keep in mind when considering how to produce biologically valid models of brain function: for example, the ability to cope with conflicting and inadequate data and still be able to function, that is, to make a decision or to remember an event given time. Clark notes that "[p]erformance should gradually become less satisfactory as available data decreases; it should not suddenly cease. Systems, in short, must be robust enough to survive in an informationally hostile environment" (Clark, 1989, p. 62). Brain systems should also be insensitive to the loss of a few neurons or a blow to the head. Healthy older people, after all, do not forget *everything* about an event, and an injury which causes amnesia is not permanent in most cases. Even when the brain has suffered a more severe loss of neurons as a result of a stroke or surgery, recovery of function can and does occur, sometimes to a surprising degree. How does this come about? All of these phenomena common to the human experience, and more, should be taken into consideration in modelling human information processing. Moreover, an overriding constraint to biological validity in modelling human performance must be that models should perform these feats in time frames comparable to those of human performance. This is a tall order for any model, but it is one which connectionist models should and do attempt to fill.

2.3.2 Attention as a biological constraint on learning

Attention and learning

It has long been recognized by educators that "paying attention" has a beneficial effect on learning, in that information to which attention is paid is remembered more clearly and for a longer period of time. This intuition on the part of educators has recently been corroborated more formally by attention researchers for whom the concept of "attention" is at one and the same time much more complex and more precise, yet more difficult to define than most educators would imagine, a problem which will be elaborated upon below. Attention has been shown to be effective in promoting neuronal plasticity at the synaptic level along with other central brain states such as motivation, behavioural significance of the stimulus, and level of arousal (Singer, 1990; Churchland & Sejnowski, 1992, chap. 5). Attention is thought to promote higher activity in the targeted neurons relative to the "unattended" surround, and in this way to effect positive synaptic change in general, that is, an increase in synaptic strength between active neurons[7]. (For a theoretical explanation of the effect of high neuronal activity on synaptic plasticity see Bear et al., 1987.) As well as having an excitatory effect on selected neurons, however, attention has been shown to have inhibitory effects on the activity of neurons: for example, when an ineffective (or irrelevant) stimulus is attended to, the activity of the cell is decreased (Moran & Desimone, 1985). This is a case

[7] An exception to this scenario may be found in anti-Hebbian learning in which activity in pre- and post-synaptic sites leads to a decrease in synaptic strength (see Churchland & Sejnowski, pp. 250-254).

of an inhibition of the cell to which attention is allocated. The operation of attention appears to routinely involve both excitation of cells receiving attention, and inhibition of cells outside this attentional focus. Rizzolatti (1983) has noted with respect to "passive sensorial attention" (more on which below):

> A necessary condition for assigning an attentional function to sensory stimuli, is that, besides their excitatory properties, they are also able to "withdraw" attention from other stimuli present in the environment. In other words, that a stimulus, beside[s] an excitation, could produce a concomitant inhibition of other simultaneously performed activities. If one examines the effects of the presentation of "attentional stimuli" to animals, there is little doubt that inhibition of preexisting behavior does indeed occur.... Thus, passive attentional processes utilize inhibitory processes which, on one hand, decrease the degree of clearness of the non-relevant stimuli and, on the other hand, permit a smooth execution of the motor responses toward the relevant stimuli. (p. 263)

For cells being inhibited, then, the outcome of the allocation of attention may be a decrease in synaptic strength due to the reduction of activity. A recent study employing PET (Positron Emission Tomography) scans has demonstrated that this inhibiting action of attention can also occur at the level of entire cortical areas. Haxby et al. (1994) discovered that "selective attention to one sensory modality ... is associated with decreased activity in cortical areas dedicated to processing input from other sensory modalities" (p. 6336). This cross-modal reduction in activity due to selective attention is of particular importance to this thesis and will be discussed in greater depth at a later time.

In short, the ability of attention to affect the activity of neurons renders it an important force in learning, and it is for this reason that its characteristics must be considered in any attempt to model how human beings learn. But what exactly is attention? This is a difficult question to answer. Indeed, no universally accepted characterization of either the structure or function of attention exists. When asked by the author if he would define attention at the West Coast Attention conference in Oregon in 1993, the well-known attention researcher Mike Posner declined to answer, presumably because the phenomena grouped under attention are extremely diverse, and therefore attention can be studied from various perspectives. To date, no body of evidence is conclusive enough to endorse one particular viewpoint on attention, even Posner's own. Nevertheless, much data on attention has been gathered and it does seem worthwhile at this stage to attempt to glean from the disparate positions on its nature and function some features which will be useful to educators. Since the purpose of this investigation is to clarify how best to conceptualize attention, the definition of attention which will be used throughout this dissertation will only be formalized at the end of this section.

At the outset, it should be noted that much of the early research on attention was directed at auditory system functioning, while a higher proportion of more recent efforts focuses on visual system functioning. Though there may indeed be significant differences between attention in audition and vision, as Van der Heijden (1992, pp. 59-62) has pointed out, due to fundamental differences in the nature of light and sound, evidence on the nature and functioning of attention will be drawn from studies involving either (or both) sensory systems. Justification for this approach can be found in the fact that general theories of attentional function have done likewise (ibid., chap. 2), that there is evidence that certain aspects of attention generalize across sensory modalities (Schneider, Dumais, & Shiffrin, 1984, p. 4), and furthermore that there is evidence that

> the attention system of the brain is anatomically separate from the data processing systems that perform operations on specific inputs even when attention is oriented elsewhere. In this sense, the attention system is like other sensory and motor systems. It interacts with other parts of the brain, but maintains its own identity. (Posner & Petersen, 1990, p. 26)

Thus, certain aspects of attention may be thought of as separate from but in interaction with the various sensory and other information processing systems.[8] Also, Van der Heijden (1992, chap. 8) argues in his model of attention that visual attention is a product of stimulation of the cells within the visual system. There seems little reason to suppose that this would not similarly be the case within the auditory and other sensory systems.

The controversy surrounding the nature and function of attention

As mentioned above, one of the most surprising outcomes of an investigation of the literature on attention is the finding that there is no generally accepted definition of this phenomenon. Indeed many papers and texts on attention begin by noting this fact (see for example Enns, 1990; Van der Heijden, 1992). Presumably, the inability of researchers to agree on an adequate definition is due to the controversy surrounding its nature, that is, its structure and function. The situation is not entirely hopeless, however. As Posner and Petersen (1990) note, although "our knowledge of the anatomy of attention is incomplete.... we can now begin to identify some principles of organization that allow attention to function as a unified system for the control of mental processing" (p. 26). This reference to attention as "*a unified system*," rather than a single phenomenon is important, since it is clear from research data that attention is not unitary, but can be

[8] It is interesting to note that this view of attention mirrors that of Cummins' vision of cognition delineated in chapter one. This is particularly so when the degree to which attention was considered to be active in "cognitively demanding" information processing is taken into account. The question that arises here, then, is whether Cummins is equating cognitive functioning with brain functioning that demands high attention.

subdivided into "intensive phenomena, such as arousal, alertness, or attentiveness" and "selective phenomena" (Rizzolatti, 1983, p. 261). From cognitive accounts of attention, Posner and Petersen (1990) have identified three major functions of attention: "(a) orienting to sensory events; (b) detecting signals for focal (conscious) processing, and (c) maintaining a vigilant or alert state" (p. 26). Though few attention researchers would disagree with these conclusions, for many the most important aspect of the attention system is its ability to select items from either the external or internal environment for "focal (conscious) processing."[9] It is *selective attention*, then, which is most implicated in consciousness or awareness (and therefore in what is generally considered to be "cognitive" processing), whereas "[a]rousal [or alertness] is a general condition affecting all of one's behaviour" (Crick, 1994, p. 59).

Even when dealing with this reduced notion "selective attention", however, agreement on a definition is hard to come by. Many authors refer back to the classic definition of attention by the great psychologist William James. Rizzolatti (1983) begins his study of the neurophysiology of attention in just this way:

> The definition of selective attention that I will use is that of

---

[9] Enns (1990) provides a functional taxonomy of visual attention in which the "highest level construct in the taxonomy is *selectivity*" (p. 140). Lower level constructs that he highlights are: integration, filtering, search, and priming.

> William James (James, 1950). In "The Principles of Psychology" (p. 403) he writes: 'It (attention) is the taking possession by the mind, in clear and vivid form, of one out of what seems several simultaneous possible objects or trains of thought. Focalization, concentration of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others'. (p. 261)

The reason why this definition has endured where others have been cast aside may be that it captures the essence of the function of selective attention, variously termed "focal", "focused", or "central" attention, the capacity to concentrate consciousness on one area of experience (that is, a single stimulus or a group of stimuli), while at the same time withdrawing it from other areas. As Schneider et al. (1984) note,

> [t]his concept presupposes that there is some bottleneck, or capacity limitation, in the processing system and that subjects have the ability to give preference to certain stimuli so that they pass through this bottleneck easily and at the expense of other stimuli. (p. 3)

The existence of the putative "capacity limitation" within the selective attention mechanism is inferred from *divided attention* research showing that "subjects exhibit reduced performance when they try to accomplish an increased number of tasks or to attend simultaneously to an increased number of stimuli" (Schneider et al., 1984, p.3). Whether the hypothesized bottleneck is one which screens out information before or after it is processed by the brain has long been a subject of intense debate in attention research circles. (See Moray, 1969; Broadbent, 1982; Kahneman and Treisman, 1984; Van der Heijden, 1992 for some interesting reviews of the topic.) There does appear to be agreement, nevertheless, on the fact that the main reason why one stimulus is given preference over another is its *relevance* or *importance* in the context of the situation at hand. What defines the characteristics of the importance of the stimulus appears to depend on whether the allocation of attention is stimulus or "world" driven (*passive attention*), or actively directed by the person, or other organism, (*active or voluntary attention*).

In the case of passive attention in which attention is "captured" by a stimulus, Broadbent has noted: "The selection is not completely random, and the probability of a particular class of events being selected is increased by certain properties of the events and by certain states of the organism" (cited in Van der Heijden, 1992, p. 44).

Properties of events which might affect the probability of a stimulus being selected are physical intensity, recency of the perception of information from a specific class of events, the frequency level of incident sounds. Sensitivity to these properties tends to be genetically pre-programmed, although it could be argued that this is only so if the surrounding environment is relatively normal for the species being observed. In the opinion of Deutsch and Deutsch (1963), however, the property of a stimulus which is crucial to determining the probability of its capturing attention given an adequate level of arousal is its "weighting of importance", which these authors consider to be determined by "past learning". They cite in support of this view studies which report "that during sleep a subject tends to respond selectively to his own name" (p. 82). Schneider et al. (1984) come to a similar conclusion, though they conceptualize the situation somewhat differently:

> We argue that the consistent training over trials leads the targets [stimuli] ... to attract attention automatically.... our suggested basis for the effects seen in these studies is a training of attention itself. Thus attention may be thought of as a trainable response in its own right. (pp. 4-5)

An interesting hypothesis on how a stimulus might acquire a high (or low) importance weighting has recently been proposed by Antonio Damasio (1994). He suggests that, for each perceived stimulus or given representation in the brain, there is an associated *somatic marker*, a feeling[10] about body (soma) state which "marks" the image.

[10] Damasio uses the terms "feeling" and "emotion" in particular ways. Essentially an "emotion" is a bodily response to a perceived stimulus, and can be pre-wired (primary emotions) or learned

These somatic markers are usually built up over time in accordance with the emotional response of an organism to a stimulus or situation and are thus learned and thoroughly individualized, but some may come about as a result of genetically pre-programmed emotional responses, in which case learning may still be operative but in a species-wide manner akin to Jungian archetypes. According to Damasio (1994), somatic markers act as "a biasing device" (p. 174) during the process of rational thinking to speed up decision making by effectively providing "an automated detection of the scenario components which are more likely to be *relevant* [italics added]" (p. 175). Somatic markers, then, are evaluations of the importance (or relevance) of the information being processed to the situation at hand, evaluations which, it should be emphasized, emerge out of emotional response and experience. What is of interest for this discussion is that Damasio proposes that these somatic markers which reflect value are the driving force behind basic attention and working memory[11]:

(secondary emotions); a "feeling", on the other hand, is the conscious experience of the changes in body state brought on by an emotion: "If an emotion is a collection of changes in body state connected to particular mental images that have activated a specific brain system, *the essence of feeling an emotion is the experience of such changes in juxtaposition to the mental images that initiated the cycle*. In other words, a feeling depends on the juxtaposition of an image of the body proper to an image of something else, such as the visual image of a face or the auditory image of a melody" (Damasio, 1994, p. 145).

[11] Damasio (1992, p. 197) defines basic attention and working memory as follows: *basic attention* permits "the maintenance of a mental image in consciousness to the relative exclusion of others. In neural terms, this probably depends on enhancement of the neural activity pattern that sustains a given image, while other neural activity around it is depressed;" *basic working memory* "holds

> There is, of course, an important question to be asked at this point: what drives basic attention and working memory? The answer can only be *basic value*, the collection of basic preferences inherent in biological regulation....
>
> In the full somatic-marker hypothesis, I propose that a somatic state, negative or positive, caused by the appearance of a given representation, operates not only as a *marker for the value of what is represented, but also as a booster for continued working memory and attention*. The proceedings are "energized" by signs that the process is actually being evaluated, positively or negatively, in terms of the individual's preferences and goals. The allocation and maintenance of attention and working memory do not happen by miracle. They are first motivated by preferences inherent in the organism, and then by preferences and goals acquired on the basis of the inherent ones. (Damasio, 1994, pp. 197-198)

The implications of this hypothesis for learning will be explored in a later chapter.

To summarize, then, the ability of a stimulus to capture attention is dependent on its "importance," also referred to as its "meaning" (Broadbent, 1982, p. 257) or its "relevance" (Schneider et al., 1984), which may be genetically pre-programmed, but which may

---

separate images for a relatively 'extended' period of hundreds to thousands of milliseconds." This definition of attention is in keeping with the one developed in this chapter.

equally be built up over time as a function of its value or learned significance for the organism[12], and which may be determined by emotional state. It can be concluded, therefore, that a dichotomous separation of selective attention into a world driven component and a voluntary or actively directed component cannot be cleanly accomplished since, even when actively focused on a location or object, attention can be captured by "important" stimuli. This "weighting of importance" feature of a stimulus is often referred to as its *saliency*. The saliency of a stimulus, at least in visual attention, is considered to be a function of the conspicuity of its location on a hypothesized "saliency map" which

> gives then a "biased" view of the visual environment, emphasizing interesting or conspicuous locations in the visual field.... Saliency at a given location is determined by how different this location is from its surround in colour, orientation, motion, depth, etc. It is possible, however, that the relative weight of the different properties contributing to this representation can be modulated by the activity of some higher centers, as for instance during prolonged practice with a particular set of targets and distracters. (Koch and Ullman, 1985, p. 221)

---

[12] It is interesting to note that even in artificial intelligence models of cognitive behaviour such as the robot, Darwin III, incorporation of a set of values (conceptualized as "innate goals") was seen to be crucial to adequate information processing performance (Edelman & Reeke, 1990, p. 235).

This dependence of saliency on the difference of a particular location from its surround explains why the *novelty* factor of a stimulus, especially in terms of abruptness of onset, has great power to capture attention in most situations[13]. In Koch and Ullman's artificial intelligence model which is based closely on physiological evidence, the registering of saliency corresponds to a high activity level in the neuron-like elements in the map: "The higher their activity, for instance their firing frequency, the higher the saliency of the corresponding location in the visual field" (ibid., p. 221). And since, according to the argument detailed above, the higher the activity of a unit, the higher is the likelihood of long term synaptic change leading to learning, it would appear that what is learned (even when direct instructions are given to a subject, but especially when they are not) is determined to some extent by both genetic predispositions and past learning.

Dichotomous concepts in attention research

The study of attention abounds with dichotomous concepts as to its nature and function, and, as in the case of passive and active attention, these dichotomous notions rarely stand up to empirical scrutiny. In essence, they are heuristics: they serve the purpose of furthering discussion. Because many of them relate to information processing issues, before elaborating on some of the

---

[13] See Van der Heijden, 1992, pp. 264-265 for a list of the kinds of abrupt visual events that can capture attention even when explicit directions for focussing attention have been given.

most important of them, it seems worth highlighting a point made by Van der Heijden (1992) when discussing visual attention:

> In very general terms, structured visual stimulation results in structured activation in the visual information processing system. The study of this stimulation-caused activation is essential for visual information processing psychology. *All "further processes", including selective attention, are based upon, make use of, or interact with stimulation-caused activation.* Visual information processing as a whole can be regarded as the interaction between, or integration of, stimulation-caused activation and, what can be called, internally generated activation. (p. 209)

The importance of Van der Heijden's statement lies in the emphasis placed on the interactiveness of the information processing system. Without stimulation of some kind, there is no processing whatsoever, therefore attention is pointless; without attention, on the other hand, there is no selection of certain stimuli over others and hence confusion results. In effect, then, without attention, the information processing system itself is useless to the organism. These are truths that are self-evident, but ones which researchers appear to lose sight of at times in their conceptualizations of the nature and function of attention.

To return to the examination of dichotomous notions in attention research, the concepts of passive and active attention find a rough correspondence in *parallel* and *serial processing* of information. Speaking of the visual search paradigm, Cheal and Lyon (1992) have commented that "[s]everal analyses of visual search have suggested that some stimulus arrays can be searched in parallel, whereas others require that an attentional 'spotlight'[14] be directed serially to different display locations until the target is found." (p. 113) In keeping with these suggestions, parallel processing is thought to be "pre-attentive," that is, not requiring active or conscious attention, is fast, and initiated by stimuli referred to as "features"[15] of a particular sensory system, while serial processing is thought to require active conscious attention, is slower, and is initiated by conjunctions of features. However, empirical support for this type of division in processing is not clear cut: using a variety of targets representing several different kinds of defining features, Cheal and Lyon (1992) found that "[s]ome targets showed strongly parallel or strongly serial search, but there was evidence for several intermediate search classes" (p. 113), and later (Cheal & Lyon, 1994a), discovered in an examination of the allocation of attention in three different

---

[14] The attentional "spotlight" metaphor will be elaborated upon below.

[15] Treisman (1988) describes features as "an elementary alphabet of visual building blocks or primitives ....[which] may be 'hard-wired' into the structure of the visual system, either innately or through early or prolonged experience" (p. 230). Cheal and Lyon (1992) complain, however, that "it has been difficult to arrive at an independent definition of a feature.... Sometimes, features are just those qualities in a character that allow parallel search" (p. 118).

research paradigms that no simple dichotomy between responses to "feature" targets and those defined by a conjunction of features was evident. Thus, as with passive and active attention, a hard and fast separation between parallel and serial processing of information seems untenable. A reasonable resolution of this matter may be found in theories which suggest that serial processing is in fact guided by parallel processing (e.g. Guided Search Theory), and others in which a clear separation of the two is not deemed necessary (e.g. Similarity Theory). (See Cheal and Lyon, 1992, for a recent discussion of these and other theories of visual search.)

Another influential partitioning of information processing, one which incorporates parallel and serial processing as well as other dichotomous concepts, is to be found in the notions of *automatic* and *controlled* processing. Schneider, Dumais, and Shiffrin (1990) define these terms as follows:

> *Automatic processing* is a fast, parallel, fairly effortless process that is not limited by short-term memory (STM) capacity, is not under direct subject control, and is responsible for the performance of well-developed skilled behaviors. Automatic processing typically develops when subjects process stimuli consistently over many trials.... *Control processing* [more usually termed "controlled processing"] is characterized as a slow, generally serial, effortful, capacity-limited, subject-regulated processing mode that must be used to deal with novel

or inconsistent information. Control processing is expected when the subject's response to the stimulus varies from trial to trial. (pp.1-2)

It is immediately obvious that the concepts of automatic and controlled processing not only incorporate parallel and serial processing of information, but also the notion of limited capacity (or "bottleneck") mentioned with respect to selective attention. What is interesting in these definitions, however, is the suggestion that the capacity limitation is determined by "short-term memory" rather than being conceived as within some "attention system." That attention is involved in these two kinds of processing, though it is not mentioned directly, follows naturally from the inclusion of parallel and serial processing as defining characteristics of automatic and controlled processing.

A recent positron emission tomography (PET) study carried out by Raichle and the group at Washington University lends some support to the reality of the separateness of these two types of information processing, in that they have been able to image two separate anatomic systems in the human brain which appear to be involved predominantly in either novel/naive language tasks or practised tasks (Raichle et al., 1994). An interesting feature of these two systems is that, according to this study, when one system is active, the other is suppressed. This finding speaks to the involvement of attention with its ability to screen out stimulation from irrelevant or unimportant information, and to choose a specific information "channel" to attend to (see Broadbent, 1982; also Haxby et al., 1994). Another fascinating discovery in this study is that it takes as little as 15 minutes practice for a task which was initially processed by the novel/naive route to pass to the practised route. This appears to confirm the observations of psychological studies that "[c]onsistent practice develops automatic component processes that exhibit fast, accurate, parallel processing" (Schneider, Dumais, & Shiffrin, 1984, p. 2). As these authors and Raichle and his group point out, however, even a relatively simple task in everyday life would involve a mixture of both automatic and controlled processing systems "possibly organized in a systematic network or hierarchy, with many of the automatic processes operating in parallel" (ibid.), hence reference to a task being carried out by automatic or controlled processes in general implies merely that these are the major component processes in the task. This is an important point, since it implies that no task is truly "pre-attentive": selective attention must be involved to some degree for successful completion of any task.

One notable feature of automatic processing, however, is that it is "not under direct subject control." Experiments designed to examine the difference in a subject's ability to control targets used both consistent mapping processing trials [CM trials, that is, trials in which "the subject makes the same overt or covert response each time the stimulus occurs"], in other words, trials which should show practice effects, and variable mapping trials [VM trials, that is, trials in which "the responses change across trials"] (Schneider, Dumais, & Shiffrin, 1984, p. 2). The results of these experiments showed that "subjects could not focus attention sufficiently to ignore stimuli that were previously CM targets" (ibid., p.10), demonstrating that CM processing (or the processing of practised tasks) is not under direct subject control even when the targets "are known to be irrelevant, when they occur in consistently invalid display locations, or when subjects are instructed to ignore them" (ibid.). Given the involvement of attention in information processing, the findings on the difficulty of blocking automatic processes speak to the difficulty of exerting active control over attention allocation in situations in which "subjects are asked to perform tasks incompatible with previously learned automatic processes," although further studies indicate that some control is possible (p.11). It seems reasonable to suppose that, once learned, stimuli interact with selective attention via the passive attention route.

As was mentioned above, incorporated into the concepts of automatic and controlled processing are the notions of a *limited capacity* channel as well as an *unlimited capacity* channel for information processing in the brain. There seems little doubt that two such channels exist. There is controversy, nevertheless, and has been for decades, about where in the information processing system the limited channel which selectively filters out irrelevant information exists (see Moray, 1969; Broadbent, 1982; Van der Heijden, 1992; for historical overviews of the arguments plus some recent updates). Since shades of the original arguments still colour today's theorizing, it is instructive to summarize the major perspectives of the past. In the opinion of Broadbent, for example, the filtering occurred early in the system and affected which information was allowed to be processed. According to Deutsch and Deutsch, however, the ability of unattended stimuli which had contextual relevance or a high weighting of importance to interfere with attended stimuli pointed to a late selection filter such that all stimuli impinging on the sensory systems was processed to the limits of its representation in the brain, that is, up to and including meaning and learned significance. A third influential theory which combined features of the two mentioned above was that of Treisman who proposed, as in Broadbent's model, an early selective filter, but one which merely attenuated unwanted messages rather than blocking them completely, thereby allowing them to pass further into the system, where they would enter a module of "dictionary units." "One important property of these units is that they have variable thresholds: importance or relevance, conveyed by explicit instructions or context, is reflected in lower thresholds and therefore a higher probability of a unit's being triggered, even by a weakened input" (Van der Heijden, 1992, p. 49). These dictionary units, then, incorporated the idea of a late selection mechanism seen in Deutsch and Deutsch's theory. As Van der Heijden suggests (1992, chap. 2), the issue of selection has been unnecessarily dichotomized in much of this discussion into a debate on whether the selection process is one of "input selection" or "response selection," though Treisman's model had shown that a successful marriage of the two could be achieved. In his view, this is unfortunate since most tasks not only involve some aspects of both, but also involve a third factor related to intention to act (or respond), "an advance preparation whether to react and, if so, how, with what class of responses, to react" (p.50). Forwarding arguments to support his contention, Van der Heijden concludes:

> This analysis of selection in vision again suggests the conclusion that in (nearly) all (experimental) tasks more attentional factors or selective processes have to be distinguished: one process concerned with the "where"' (i.e., controlling the source of stimuli to be responded to), a second concerned with the "what" (i.e., controlling what aspect of the stimulus is of importance) and a third concerned with "what to do" (i.e., controlling whether, and if so, how, to react). These processes are very close to the processes already suggested by Von Helmholtz, James, and Gibson, respectively: attention, expectation, and intention. (p. 55)

This tripartite nature of the attention system will be taken up again in an elaboration of Van der Heijden's information processing model. The concept of "attention" employed in the model is much narrower than that used in other models, representing as it does only one selective process of three, all of which work in close cooperation to enhance selectivity in brain processing.

Metaphors for attention

Because of the difficulties experienced in attempting to characterize the nature and function of attention and its role in information processing, attention researchers have resorted to the time honoured practice of employing metaphors to give conceptual form to their empirical findings. Studies of attention allocation in the visual system have been particularly fruitful in spawning such visualizing aids. One abiding image, initially proposed by Broadbent, is that of a *"filter"* selectively allowing only some information to pass further into the information processing system, or as in the case of Treisman's model, giving a selective facilitation of passage for some information while attenuating the passage of other information. A second influential metaphor promoted by Posner and others likens attention to a *"spotlight"* which moves across the sensory system in an analogue manner. When attention is focused on a particular spot, it is said to be "engaged" at that location, hence a relocation of attention would involve "disengagement", followed by a "shift" and then re-engagement, the time necessary for relocation increasing with increasing distance from the original site of engagement. Similar to the spotlight model, but allowing for a variable size of attention gradient centred at the focus of attention is the *"zoom lens"* model. While these metaphors have been widely used over the past decade, doubt has been cast on their usefulness by a growing body of data (see Cheal, Lyon, & Gottlob, 1994b, pp. 700, 701). Introduced more recently is the *"gradient"* model of attention, in which "attention may be distributed in a gradient around the cued location, or it may form an irregular field with multiple peaks of different intensities (heights)" (ibid., p. 701). In yet another model, it is proposed that, when a location cue appears, a *"gate"* is opened which initiates a flow of information in an appropriate channel. This is the so-called *"attention gating"* model. One of the most recent models combines elements of both the gradient and the attention gating models. This is the *"gradient filter"* model, introduced by Cheal, Lyon and Gottlob (1994b), which is based on the concept of semi-permeable filters which control the rate of flow of information from the visual display, and which themselves are regulated by the amount of attention at each location. In this model, attention may be thought of as being spread out across the visual field in an uneven fashion, depending on the probabilities of a target appearing at a particular location. If one location is highly probable, the permeability of the attention filter at that location will be high, rendering the rate of information flow correspondingly high at that spot, and thereby enhancing access of information at that location. With respect to the "limited capacity" issue, the model suggests that if there is an overall capacity limitation, changing the permeability of a filter at one location will affect the permeability of filters at other locations, either positively or negatively in accordance with the direction of the initial change. According to its authors, "[t]he gradient filter model is differentiated from other metaphors in that various filters become more or less permeable and are not dependent on movement (as in the spotlight metaphor) or on the simultaneous opening or closing of gates" (pp. 703-704).

Although none of these metaphors or models has been wholeheartedly endorsed or rejected by the research community, some of them seem less helpful in their heuristic intent than others, in that they appear to conceptualize attention as a force operating on the information processing system proper from somewhere outside its boundaries (see Van der Heijden, 1992, chap. 7 for an elaboration of this point of view). The spotlight and zoom lens metaphors are particularly prone to this type of interpretation. More recent information processing models based on neuroscientific discovery, however, conceptualize attention (defined in a narrow sense) as internal to the system, indeed, arising out of the properties and organization of the cells themselves as well as the modular structure and functioning of the brain (see, for example, Laberge, Carter, Brown, 1992; and Van der Heijden, 1992, chap. 8). Van der Heijden (ibid.) criticises "external" attention models with these words:

> In the recent past the term "attention" has often been used to indicate a magical force, capable of solving quite unsolvable problems in a non-specified way ... to prevent the need of

introducing an empty, magical, concept of attention our strategy will be to account for as much as possible in terms of the fixed structure and functioning of the information processor [the brain]. (p. 209)

Since he and others (see Laberge et al., 1992) have not only described such a model but have successfully constructed connectionist working models of attention based on their theorizing, it is this level of modelling which would appear to be most relevant to this thesis, and which will be detailed below for its heuristic value.

A system internal conceptualization of attention

The model of attention outlined here will be that of Van der Heijden (1992). It is a model of selective attention in vision, a conceptualization which is not only specific to the visual system but is seen as internal to the system, that is, arising out of the organization and structure of the visual system itself. As such, it may not be entirely applicable to attention in the other sensory systems, nor may it explain other aspects of attention which appear to be sensory system external. The latter will be dealt with in more detail later in the discussion. Nevertheless, the notion of attention as acting at the neuronal level to activate or inhibit cells is an important one for the concept of learning as a process and product of synaptic change, hence variation between systems, though potentially significant, will be taken as simply variations on a theme of neuronal excitation and inhibition through attention allocation.

Before elaborating the details of his model, Van der Heijden (1992, chap. 7) establishes very clearly his position on the unlimited/limited capacity of the human information processor debate. In his view, all stimuli impinging on the brain are processed to the limit of their representations, that is to say, beyond mere perception to complete identification, which implies accessing information on the value or learned significance of the stimuli. As he points out, this perspective is in agreement with those of others working in the field (see Van der Heijden, 1992, pp. 236-237). Thus, he conceives the selective action of attention as operating *after* the simultaneous processing of many stimuli to impose temporal order on spatially ordered information so that which information should be acted upon first may be determined. The selective function of attention, then, is a *selection for action*, a selection which solves a "temporal priority problem: a problem of ordering actions in time.... The function of attention is in the time domain. Attentional selection has to impose order or structure in time" (p. 245). According to Van der Heijden, this type of selection is necessary precisely *because* of the unlimited capacity of the information processor. However, selection for action is insufficient for adequate performance of a task.

[A] second form of selection is needed because a selected action

> can generally be directed to only one among a number of simultaneously available objects at a time. .... This form of selection has therefore to solve the problem of which object to act upon at a certain moment in time, i.e., the problem of where the action is now to be directed (p. 245).

This second form of selection is necessary to avoid behavioural chaos and is carried out by selective attention in accordance with the constraints of the expectations and intentions of the subject performing a specific task. As Van der Heijden notes:

> For an adequate theory of selective attention it is essential to know what stimulation contributes ..., what preparation (i.e. intentions and expectations) contributes, and what further contribution is left for the information processing system that deserves the name selective attention. (p. 247)

Thus, selective attention does not operate in a vacuum, but in close cooperation with stimulus-invoked activation, experience or instruction-invoked expectations, and intention to act. It is interesting that, temporally, these factors reflect the present, the past, and the future. In real life situations, therefore, attentional selection can rarely be considered free of historical context.

But how exactly does Van der Heijden conceptualize attention?

Before answering that question, it is necessary to describe briefly the neurophysiological framework of the visual system within which visual attention operates. There is evidence that, in the visual information processing system, "two major pathways are distinguished: the X- or parvocellular pathway, presumably concerned with the 'what', and the Y- or magnocellular pathway, presumably concerned with the 'where'" (ibid., p. 210; see also Zeki, 1992), and that within the parvocellular pathway two broad groups of specialist modules can be distinguished, one in which information is spatially ordered in the form of topographic maps and which can be considered to deal with what it means "to see," and another in which information is not spatially organized and is generally thought to be more abstract in content, reflecting the identity, meaning, and learned significance of the visual information, all of which leads researchers to speculate that the areas within the second group are concerned primarily with "to know." This latter group Van der Heijden terms "domains." Thus, in his basic model, input received at the "map" level (the input module) is automatically and simultaneously sent on to the "domain" level (the identity module) to be identified, as well as to the magnocellular pathway to be located in space (the location module) (see Van der Heijden, 1992, chap. 8). So far in this scenario, attention is not active, and if it is not allocated to the stimulus invoked activation, the latter will fade rapidly. According to this model, it is only when a stimulus is located that attention kicks in, and it does so by means of a feedback loop originating from the location module and ending in the input model. Thus, it is the feedback itself which is regarded as "attention" in this model, and it is by means of the feedback information from the location module that selection occurs.

> The operation of the feedback loop can be regarded as a visual-perception internal kind of selective attention triggered by position information or the onset of position information. So, in this view, it is not attention, coming from nowhere, that is directed at a position in a map of locations. The map of locations is the *source* of the attention. Location information - if fed back to the input module - *is* the attention. Not attention directed at position, but position directed as attention. (Van der Heijden, 1992, p.254)

On this view, then, there is no need for a "magical" concept of attention coming from somewhere external to the sensory system. Attention is a system internal phenomenon according to this model. In a situation in which various stimuli are simultaneously competing for attention, expectation, whether given by verbal instruction, visual cues, or past experience with a task, acts by enhancing the activity of a particular location in the location module, thereby effectively enhancing the speed of processing of the information available at that location, and enabling that information to be identified first, that is, to be selected. The problem with selection based only on expectation, however, is that it selects *too much* for adequate completion of the task. For good task performance, only a small subset of the properties of any object are necessary: for example, in a letter naming task, the shape of the object alone is sufficient; properties such as size, colour, and brightness are irrelevant. Hence, as alluded to above, a second form of selection aimed at screening out irrelevant properties must be carried out by "intention to act".

> This intention to act can be equated with property selection. It can be conceived as an (advance) preparation in the task-relevant modules or an (advance) setting of the perceptuo-motor channels. Then the intention to act picks out the response properties when attentional selection presents all properties of the object to be acted upon.
>
> So, in our conceptualization, attention (in cooperation with expectation) performs object selection ... This object selection might possibly be sufficient for "perceiving" or "seeing" the selected object, as Treisman suggests ... However, for adequate "task performance" or "behaviour", attentional selection is not sufficient. In addition one or another form of property selection is needed. In our proposal a specific intention to act is at the basis of this second form of selection. Both attention (guided by effective[16] expectation) and intention are involved in

---

[16] Van der Heijden (1992) clarifies what he means by "effective" expectation: "The word *effective* is used to indicate that what is meant is not the total of a subject's expectations with regard to stimuli, task, experimenter, room, etc., but only that part of a subject's total expectations that is made effective

the execution of a task. (Van der Heijden, 1992, pp. 270, 271)

That such a conceptualization of attention is reasonable has been demonstrated by artificial intelligence modelers: Van der Heijden (1992, chap. 8) describes the structure and performance of the model SLAM (SeLective Attention Model), a connectionist model which has demonstrated that "only the cooperation and intersection of the efforts of attentional object selection and intentional property selection produce the correct behaviour. The intention to act in a specific way solves the problem created by attention: selecting and emphasizing far too much" (p.273). Other connectionist modelers have shown that physiological effects of attention observed at the cellular level can also be mimicked effectively by models (see, for example, Laberge, Carter, & Brown, 1992; Niebur, Koch, & Rosin, 1993). It would seem, therefore, not only that connectionist models have captured some basic functional principles of the brain, but also that the function of a complex phenomenon like attention does indeed begin at the "unit" level, that of the neurons themselves, but does not end there. In summarizing his approach, Van der Heijden writes:

> [W]e already saw that the bottom-up or world-driven approach only will never really work ... In the course of this study we were more and more confronted with the fact that a top-down or

---

by means of the *instruction* in the performance of the task" (p. 263).

subject-controlled approach only will also never really work.... In general, only a theory that acknowledges the existence and cooperation of attention, expectation and intention and, within attention, both the contributions of involuntary and voluntary attention, will be able to elucidate the real function of attention in visual information processing tasks. Both the subject and the world have to find their proper place. (Van der Heijden, 1992, p. 282)

Thus Van der Heijden's approach attempts, by means of clear and logical analysis of the available data, to cut through the dispute-ridden bottom-up versus top-down approaches in order to acknowledge the validity of both mindsets, but only within an integrated model: not world-driven *versus* subject-controlled, but world driven *and* subject-controlled. Since this approach is not only consistent with the data, but also conforms with intuitive insight based on life experience, it seems one which is eminently suitable to apply to the problems of elucidating the interaction of attention and learning. Hence the description of attention afforded by this model, though specific to the visual system, will serve as a general base for theorizing on the interaction of attention in information processing in the brain (Van der Heijden, 1992, p. 282):

Visual spatial selective attention

- consists of (late) position information joining, and interacting with, (early) identity information,

- and provides temporal order or temporal structure in a spatially structured visual world.
- It performs its job in close cooperation and interaction with expectations and intentions,
- and is controlled by both the subject and the world.

Attention as a selective strategy

The view of selective attention as a sensory system internal phenomenon is indeed a logical and coherent one, and it does explain much of the data on attention. However, it does not explain how it is that attention can be allocated preferentially to one or another sensory system as has been demonstrated by early research such as that of Broadbent (see Broadbent, 1982; Moray, 1969, chap. 6). Whether this type of selection can be triggered by a system internal mechanism or is the product of a system external (or central) mechanism has not yet been decided, and indeed, as has been mentioned above, there is evidence to support both an anatomically separate attentional mechanism which interacts with other parts of the brain while maintaining its own identity (see Posner and Petersen, 1990) and a neural organization which could underlie a system internal mechanism (see Van der Heijden, chap. 7; also Haxby et al., 1994). Much as these findings appear to be at odds with one another, it should be remembered that the purpose of Van der Heijden's model was to try to screen attention out of other so-called "attentional" features such as expectation and intention which must undoubtedly be considered to make up part of a more comprehensive "attentional system".

This ability to select a sensory system to attend to, and other abilities related to it, as, for example, the preferential allocation of attention to one ear or the other, has been termed "filtering."

> [It] was conceived as a strategy that could allow satisfactory performance in cases when interference would otherwise occur. That is, if a large number of complex events are occurring, a person who selects those events possessing a particular physical feature (such as location in a particular point in space) [or "events that share some basic physical property , such as arriving at the eye" (Broadbent, 1982, p. 257)] will be able to cope adequately with those events at the cost of knowing less about the remainder of the things that are happening. Filtering was not however now imagined to block out everything about the unattended parts of the surroundings; some features would break through from other places, perhaps enough to trigger some later processes. (Broadbent, 1982, p. 259)

Filtering, then, is a selective strategy which trades benefits for costs, but only one of several identified by researchers, "categorising" and "pigeon-holing" being two others of importance (ibid., pp. 257 - 261). Since the selection of one sensory channel over another is important to this dissertation, however, only this particular strategy will be dealt with here.

Consideration of filtering as a selective strategy brings two questions into prominence: first, why would one system (sometimes termed a "channel") be chosen over another, and second, how might such a choice be effected? To deal with the "why" some thought must be given to the role of the sensory systems in information gathering. In Van der Heijden's opinion,

> [i]t is ... very easy to defend the general point of view that the perceptual systems are not just for "perceiving" and the particular point of view that the visual system is not just for "seeing". The most important function of the visual system lies in the transformation of information in light into actions (reading, naming, walking, grasping, etc.). Vision, and all other "perceptual systems have evolved in all species of animals solely as a means of guiding and controlling action ...". (Van der Heijden, 1992, p. 250).

Given this view of the functional priorities of the sensory systems, it seems reasonable to assume that a reduction in knowledge of the sort referred to by Broadbent due to the non-selection of a sensory system would only be acceptable if a decision had been made that the system selected carried (or at least was likely to carry) more, or more useful, information for guiding future actions. Ostensibly, such a decision could be made on an immediate in-situation sampling of the data on all sensory channels, on the basis of instructions from an authority figure, on the basis of the probability of gaining useful information from a particular sensory system, that is, on past experience in information gathering, or even on the basis of genetic predisposition towards one sense. Each of these presupposes a strong element of "voluntary control" over the decision making process. The breakthrough of information from the "unattended channel," however, has shown that voluntary control of attention allocation is by no means absolute, and in fact, that it can be superseded by "high priority" information arriving on an unattended information stream (see Broadbent, 1982). Thus, it would appear that, as far a choosing a sensory system to attend to is concerned, attention can be world-driven or subject driven, and perhaps even both. It is evident, then, that the description of attention proposed by Van der Heijden based on consideration of visual system processing only scales up quite nicely to account for the facts at sensory system external levels.

As to how a choice of sensory system might be effected, it appears likely that voluntary control mechanisms would be governed by "higher centres" capable of transforming probability data from past experience (or instruction, etc.) into an "effective expectation" in terms of neuronal excitation within a particular system such that information processing would be carried out there at an enhanced rate. It also seems likely that "intention to act" would have some executive control in the overall decision making process. But given that the predominant characteristic of attentional allocation at the neuronal level is enhanced activation of one cortical area relative to others, how could external stimuli (the world) bring this decision about? The clue to answering this question would appear to lie in the high "importance" factor of the type of information which has been shown by experiment to cause "breakthrough" or attentional "capture." It seems entirely reasonable to suggest that the importance weighting of the incoming information causes it to be processed more speedily and thereby effects the capture of attention through speed of processing alone. The source of the importance weighting may be merely repetition (sometimes called "overlearning"), a process which would lead to the well-established neural pathways which underpin "automatic" processes (see Schneider, Dumais, & Shiffrin, 1984). However, importance weighting may equally be the product of survival factors, as for example, the attention capturing power of the voice of one's own child or someone calling one's own name (see Deutsch and Deutsch, 1963). This aspect of importance weighting is also the result of experience or learning. In most cases, however, importance weighting is decided upon within the context of the situation at hand and can only be determined by an interaction between stimulus activation, expectation, and intention. Once again, therefore, it can be seen that Van der Heijden's restricted model of attention has validity at a more general level.

What is striking about stimulus-induced attention capture is its dependence on experience or learning. To put this another way, it would appear that *familiarity* with the stimulus is of critical importance in attention capture. Consideration of the work of Raichle et al. (1994) may give clues as to how this comes about. It may be remembered that the data from PET scans showed that, when a subject was processing a task under "novel/naive" conditions, a particular anatomical subsystem was called into action (the anterior cingulate system), whereas when the task was "practised" (or familiar?), a different subsystem was activated (the sylvian insular system). It was speculated by the authors that these two subsystems might be the processing routes for "controlled" and "automatic" processing respectively. If so, it seems likely that well practised information is processed more speedily than novel information or less practised information. Thus it would certainly be identified more quickly than new information and may in fact aid in the allocation of attention to specific features of the novel input (i.e., the guided search hypothesis)[17]. According to Van der Heijden's model, however, being identified would also mean activating the learned significance and value of the information. This could then be passed on to a "higher centre" which could transform it into selective activation of a "location module,"

---

[17] As to why the unfamiliar information (novel/naive tasks) are processed more slowly and are more attention demanding, it is likely that the slowness stems from the lack of well-established neuronal connections for the processing of the elements of the task, and that the high allocation of attention provides the high neuronal activation necessary for the changes in synaptic strength which would speed up that processing.

thereby initiating a feedback loop to the input module and calling attention to the familiar and important information. The only flight of fancy to be added to Van der Heijden's basic model to effect this would be a location module on a higher level than that inside a sensory system, one which could pinpoint locations within the range of sensory cortical areas. This would appear to be in keeping with the notion put forward by Posner and Petersen of an anatomically separate attention system which maintains its own integrity while interacting with sensory systems.

The mechanism above is highly speculative, of course. What is not speculative, however, is that, along with the novelty of the stimulus which has long been known to be a strong attention capturing feature, paradoxically[18], what is familiar, especially when it has a high importance weighting for survival, also has the power to capture attention, that is, to bring about a conscious awareness of the particular familiar stimuli through preferential activation of its signal, and it seems not unreasonable to suggest that that preferential activation might spread to the sensory system which carries it.

---

[18] A resolution to this paradox may be found through consideration of the importance weighting of novelty in perception. If the perception of change in the sensory environment has a high importance weighting, either innately or through experience, as it seems likely that it should, then a module which is set up to perceive such change could facilitate the processing of information at that particular location. Thus, the mechanism behind the attention drawing powers of both novelty and experience may be essentially the same. However, they may also be quite different.

A working definition of attention

It was noted at the beginning of this section on attention, that a generally accepted definition of attention is very hard to come by. The reasons why this is so should now be clear. Attention involves a highly complex set of phenomena which range in level of activity from the sub-cellular to the CNS. Nevertheless, it seems important to attempt a working definition of attention in order to have something more manageable than the mountain of facts to conceptualize with. At the risk of being simplistic, then, the working definition of attention which will be adopted for this thesis is the following:

> Attention is a facilitatory mechanism for the selection of information for further processing[19]. It operates in close cooperation with expectation and intention to act, and can be driven by voluntary or stimulus-initiated mechanisms.

It should be remembered, however, that underpinning this definition are *all* the identified complexities of the phenomenon called attention.

2.3.3 Connectionist models of learning

---

[19] In acknowledgement of the conflicting views on how much information processing is carried out before attention clicks in, it should be noted here that the term "further processing" as used in this definition may simply mean the bringing of the information to awareness, rather than the initiating of identification and recognition.

The basic design of connectionist models of learning is that of a *net*, sometimes called a "neural net" or "neural network" due to the fact that the elementary units in the net, as well as their pattern of interconnectivity, are "neuron-like" in their structure and properties. However, as Hinton (1992) states, "[b]ecause our knowledge of neurons is incomplete and our computing power is limited, our [connectionist] models are necessarily gross idealizations of real networks of neurons" (p. 145). Evidently, then, the level of representation sought by these models in the hierarchy of levels of organization in the brain is that of the network. Of course, this necessitates the incorporation of biological constraints gleaned from the cellular and subcellular levels, and, depending on the problem to be solved, constraints from the map or systems level may also be taken into consideration (see Kosslyn and Koenig, 1992, for a recent review of connectionist models of higher functions).

The general approach taken by all of these models of learning is that, in biological systems, the learning process is essentially a process of "error minimization" which is itself a way of optimizing the interactions of the organism with the world surrounding it (see Churchland & Sejnowski, 1992, pp. 130-134, for a discussion of this point). Since the raison d'être of these models is to mimic the learning procedures in the brain in order to gain more knowledge on brain processing in the hope that "model nets may be a valuable source of ideas relevant to real neural nets" (Churchland &

Sejnowski, 1992, p. 134), and that these ideas will lead to predictions regarding the actual nervous system which will be neurobiologically testable, it is of the utmost importance, given the "gross idealization" aspect of the models referred to above, that at least the *critical features* of neurobiological realism are incorporated into the models. Though what constitutes a "critical feature" is a source of ongoing debate, the constructing of a connectionist model of learning involves the making of decisions on a triad of design features, the architecture, dynamics, and parameter adjusting procedure of the net. In any net, the choice within each design feature must be governed by the particular brain process being modelled. The architecture of the net should reflect the anatomy of the system under study, the dynamics should reflect the physiology, and the parameter adjusting procedure should reflect the procedure for changing the synaptic weights. Notice that, in the opinion of connectionist modellers, it is not necessary to replicate the biology of these features exactly: it is sufficient to "reflect" it so long as the features of the net are "relevantly similar" (Churchland & Sejnowski, 1992, p. 132) to the biological reality. Just what constitutes "relevant similarity" for each of these features of the nets will be dealt with in more detail below.

Given that connectionist models are computational ones which purport to model brain processing, it seems important, before detailing the structure and function of these models, to ask in what sense brain systems can be considered to compute? Churchland and Sejnowski (1992) tackle this philosophical problem by first defining what is necessary for a physical system to be considered a "computer":

> [I]n the most general sense, we can consider a physical system as a computational system when its physical states can be seen as representing states of some other systems, where transitions between its states can be explained as operations on the representations. The simplest way to think of this is in terms of a mapping between the system's states and the states of whatever is represented. That is, the physical system is a computational system just in case there is an appropriate (revealing) mapping between the system's physical states and the elements of the function computed. (p. 62)

As these authors point out, however, we do not count as a computer a device or a system for which discovering the underlying function reveals nothing of interest to us. Thus "computer" is not a natural kind because "there is no intrinsic property necessary and sufficient for all computers, just the interest-relative property that someone sees value in interpreting a system's states as representing states of some other system, and the properties of the system support such an interpretation" (pp. 65-66). Nervous systems, therefore, can be considered computers on the premise that they are physical devices (in whose machinations we undoubtedly have great interest) which have evolved so that their states represent the states of the external environment, or the body, or even other parts of the nervous system, and consequently, their physical state transitions can be considered to be performing computations. An interesting example of a brain system which is thought to be computing in just this sense is the cerebellum which translates sensory information into motor execution by matching "the system of relations among objects in the external world, with a multidimensional inner functional geometry, in a manner such that these geometries approach homeomorphism if not isomorphism" (Pellionisz and Llinas, 1985). (For a more complete discussion on this point see Churchland and Sejnowski, 1992, pp. 61-69. This text is also replete with further examples of biological computational systems.)

Typical net architectures

The basic element of any artificial net is the *unit* which is modelled on the neuron. Each unit consists of a "cell-like" body from which "axon-like" *connections* extend out to other units in the net. Each connection has a modifiable *weight* which governs to what degree the signal carried by the connection will affect the unit contacted. Thus the weight acts in essence like a synapse. Unlike neurons, however, units do not have dendritic branching: for simplicity, all connections make contact directly onto the "cell body" of the unit. The unit's job is to convert the incoming pattern of activity into a single outgoing activity (representing the rate of firing of the unit or "cell") which it then passes on to other units. Hence units communicate to each other information on their activity, which is determined in part by the characteristics of the input signal or stimulus, in part by the pattern of weight connections, and in part by the nature of the input-output function assigned to the unit. On receiving input, the unit multiplies each incoming activity by the weight on the connection, sums the weighted input, then pushes this total input through an input-output function which determines the final outgoing signal. The input-output function can be linear, threshold, or sigmoid in character, and though the capabilities of the net will change according to which kind of function is used, each type of unit can be considered an approximation of a neuron in its behaviour[20]. (See Hinton, 1992, for a very clear explanation of the structure and function of units.)

Networks are generally organized into layers of units with similar functions, mimicking the laminar organization of many areas in the brain. There are three typical arrangements of units into networks that are worth examining in the context of learning: associative nets, nets with hidden units, and recurrent nets. In each of these

---

[20] Hinton (1992) describes the difference between the different functions as follows: "For linear units, the output activity is proportional to the total weighted input. For threshold units, the output is set at one of two levels, depending on whether the total input is greater than or less than some threshold value. For sigmoid units, the output varies continuously but not linearly as the input changes. Sigmoid units bear a greater resemblance to real neurons than do linear or threshold units, but all three must be considered rough approximations" (p. 145).

designs, the input units are generally considered to represent sensory input, their levels of activation being determined by some aspect of the environment either directly or by experimental manipulation. Since the activity of each unit can be represented by a number, a set of input values can be referred to as an *input vector*, which is just an ordered set of numbers. The set of weight values impinging on a unit can similarly be described as a *weight vector*. Thus, taking advantage of the rules of linear algebra, an input vector may be converted into an *output vector* by means of a process of matrix multiplication of the weight vectors by the input vector. Describing the input and output patterns as vectors, then, is both economic and practical in that it makes available to computer modelers the considerable resources of algebraic theory. (For more detailed information on the topic of vectors as they apply to connectionist models, see Rumelhart, McLelland, & the PDP Group, 1986, chap. 9.)

Simple associative nets consist of a set of input units that are connected to a set of output units such that each input unit has a connection to each output unit.

> The aim is simply to store a set of associations between input vectors and output vectors by modifying the weights on the connections. The representation of each association is typically distributed over many connections and each connection is involved in storing many associations. This makes the network

> robust against minor physical damage and it also means that weights tend to capture regularities in the set of input-output pairings, so the network tends to generalize these regularities to new input vectors that it has not been trained on. (Hinton, 1989, pp. 189-190)

Because the input units are activated simultaneously, the information processing pathways operate in parallel, rather than in series as occurs in a traditional digital computer, and, as is mentioned above, results in a distributed representation of the output. This is the reason why the term *parallel distributed processing (PDP)* has become virtually synonymous with connectionism. Though associative networks can capture regularities in the input and some, with nonlinear units, can perform vector completion tasks when given degraded input, the regularities they capture tend to be low order such as answering the question "'do feature A and feature B correlate?', which in machine terms means 'are the A-unit and the B-unit on together and off together?'" (Churchland & Sejnowski, 1992, p. 98). "Simple" human feats like object recognition, on the other hand require the sifting out of more complex regularities. Thus, if the model is to come close to mimicking human recognition ability, it should be able to recognize objects when seen from a different angle, when partially occluded, or when deformed in some way. In order to acquire such computing power, "hidden units" must be added to the system described.

Hidden units are simply units which are neither initial input nor final output units, sometimes termed "visible" units. Their function is to sum the input they receive and pass it on as output to the next layer of units. According to Churchland and Sejnowski (1992), "they typically connect to the input units, to each other, and to the output units" (p. 99), though connections within the hidden layer are rarely shown in diagrams of typical networks. These hidden unit layers (which can range in number from one to many) are sandwiched between the initial input layer and the final output layer. In a sense, then, they are *both* input and output units, since the output they produce can be thought of as providing input for the layer they project to (see P. M. Churchland, 1989, p. 162, for an instructive drawing of a simple network with hidden units). Adding a layer of hidden units endows the network with a *global* perspective in spite of its local connectivity pattern. When the number of input units is very high,

> the problem for the hidden units is to discover what combinations of features are ignorable, which features systematically occur together or are otherwise "cohorted," and among *those*, which are the combinations to "care" about and represent. The information for this last task cannot be garnered from inside the net itself, but must be provided from the outside. (Churchland & Sejnowski, 1992, p. 99)

This means that the information necessary for hidden units to determine which features to "care" about must come from an external source of information, which indicates that some nets must be capable of taking instruction as from a teacher. Such instruction can be considered "supervision". In fact, there can be both "supervised" and "unsupervised" forms of learning for these machines. These will be dealt with in more detail below. With respect to concerns about the biological validity of hidden units, Patricia Goldman-Rakic (1990) has discovered that there are areas in the cortex of the macaque monkey lying in-between sensory input and motor output which she believes are candidate structures for hidden units.

The types of nets described so far are *feedforward nets* in that the direction of information flow is from the input units to the output units. Though it is true that information on the measure of error between the real output and the desired output is passed back through the system in the weight adjusting process (more on which below), this cannot be considered "unit activation" and hence is not input in the traditional sense. That is, there is no flow of information as input from the higher to the lower levels in the network. This means that a feedforward system is a reactive system which produces output only as a function of external input and the configuration of weights. *Recurrent nets*, on the other hand, have *internal* input in that the hidden units can receive inputs from the output units, and from one another, as well as from themselves (see diagrams, Churchland & Sejnowski, 1992, pp. 77 and 121), with the result that the stimulation provided by the external input is not isolated from what went on before in the net.

> In recurrent nets the significance of external input varies across three general types of cases: (1) the net may have output only when there is both external and internal input; (2) the net may have output even when there is no external input, but the continuous activity must initially be triggered by external input; or (3) internal input alone will suffice to activate an output, though the net's output can be modified by the addition of external input. (ibid., p. 117)

The effect of internal input via feedback loops is to render the net sensitive to what has happened in the immediate past and thereby to create a type of memory. In some nets, long recurrent connections provide the hidden units with information about the recent history of the output units, while short, self-exciting connections "create a short-term memory, inasmuch as a signal that produces a level of activity in the unit will decay slowly rather than cease abruptly, for the self-excitation maintains the activity" (ibid., p. 122). This type of performance is significant when attempting to model a smooth sequence of behaviour such as coordinated muscle contractions.

Finally, networks in which systems of *specialist mininets* are interposed between the input layer and the output layer have been constructed (see Churchland & Sejnowski, p. 130). These mininets complete with their own array of hidden units, modifiable weights, and connections are trained to represent a particular input pattern more and more accurately within a larger specialized network. Organization of this kind corresponds well with neurobiological studies which point to the existence of specialist substructures. George Ojemann (1991), for example, has noted that there is a good deal of evidence which suggests that

> the cortical area dedicated to language is not unitary, but compartmentalized into separate systems for processing different aspects of language.... Moreover, lesion studies have indicated separate areas for handling different languages.... [and] separate areas are also present for handling different grammatical classes of words. (Ojemann, 1991, p. 2282)

Churchland and Sejnowski (1992) also note that this same researcher "has observed interference with Greek only and not English, or English only and not Greek, where the intervening [cortical] distance may be as small as 10 mm" (p. 130).

Although the net architectures described here in no way constitute an all inclusive list, they give some idea of the different flavours of nets employed by connectionist modellers interested in modelling particular brain functions, and set the scene for a discussion of how these nets learn through experience. It should be clear by now that the type of net architecture adopted is dictated by the demands of the task at hand. It may be less clear, but it is nonetheless true, that it is also dependent on the internal dynamics of the system as well as the parameter adjusting procedure adopted.

Dynamical systems and system dynamics

Some of the roots of connectionist or PDP models are to be found in the work of the Russian psychologist and neurologist Luria who put forth the notion of the brain as a *dynamic functional system*. According to McLelland, Rumelhart, and Hinton (1986), by this he meant that "every behavioural or cognitive process resulted from the coordination of a large number of different components, each roughly localized in different regions of the brain, but all working together in dynamic interaction" (p. 41). In much the same way, a network can be considered a *dynamical system* made up of individual components (the units), all working in a coordinated fashion to give the final output, a motor activity perhaps, or recognition of an object. In a dynamical system, "the activation values of the units are seen merely as variables that assume various values at various times.... [and] The 'knowledge' contained in a PDP model lies in its connection strengths ... or weight matrix" (Smolensky, 1986, pp. 397-398). To repeat for the sake of clarity, then, *knowledge* in a connectionist system is encoded in the *set of connection weights*.

The beauty of conceptualising a connectionist model as a dynamical system lies in the analytical resources that this viewpoint carries with it, the mathematics and physics of dynamical systems. Any dynamical system can be described in terms of its *kinematics* and its *dynamics*. For each dynamical system, there is a *state space*, the set of all possible states of the system, and a *set of trajectories* through the state space, where a particular trajectory is a path taken from an initial starting state by the sequence of states that ensues as the system evolves or changes over time in obedience to the algorithm which governs its evolution. (See Smolensky, 1986, pp. 397-399, for more information on dynamical systems, and P.M. Churchland, 1989, chap. 5, as well as Churchland & Sejnowski, 1992, pp. 167-174 for discussions of state spaces.) The system's kinematics describes the geography of the state space, while the dynamics deals with how the system changes with time.

In connectionist networks, therefore, the *system dynamics* is dependent on both the type of units employed (linear, threshold, or sigmoid) and the parameter adjusting procedure or algorithm, as well as the connectivity pattern or architecture of the system. The "parameters" in the case of connectionist networks are quite simply the weights on the connections between units. For any system, different kinds of state space can be constructed depending on which particular aspect of the behaviour of the system is to be examined. Some state spaces that have been found to be useful in monitoring the effects of changes on systems with hidden units are

(1) "synaptic" or connection weight space, the space of all possible combinations of weights in the network; (2) hidden unit activation-vector space, the space of all possible activation vectors across the hidden units, and (3) weight/error space, the space that measures the percentage error associated with any particular weight configuration (see Churchland & Sejnowski, 1992, p. 169 for representations of (1) and (2), and P. M. Churchland, 1989, p. 166 for (3)). The trajectories traced through the connection weight state space would show the paths of how the weights change during learning. These changes are reflected in hidden unit activation-vector space by the partitioning of the space into regions which represent similar vectors, with degree of similarity being determined by spatial proximity. With extensive training, each region of similarity can develop a "prototype" vector region which acts as an attractor for vectors which fall into its similarity space. This is a useful feature of the network because, in developing the prototype, the system has essentially learned to categorize, and any vector falling into the region, that is, any vector which is relevantly similar to the prototype vector, will then be categorized as the prototype. Thus, the system will be able to categorize samples of the same category which were not in the training set (see P. M. Churchland, 1989, p. 206). This ability, of course, is comparable to that of the human brain.

A trajectory through the weight/error space, on the other hand, would show how the error is minimized through learning. In most connectionist systems, this path is not linear, nor steep, but winds tortuously in miniscule shifts down the "error gradient" to its optimal position, that is, the global error minimum. This process of error minimization, therefore, is termed "gradient descent" on the error surface, where the "error surface" of the system is defined as follows: "For each combination of weights, the network will have a certain error which can be represented by the height of a point in weight space. These points form a surface called the 'error surface'" (Hinton, 1989, p. 194).

In some cases, as in nets with linear output units and no hidden units, this error surface forms a bowl, which guarantees that the global error minimum will be reached, since a bowl has only one minimum. This means that the system will come to represent accurately what it is being trained to represent. In other cases, however, for example, nets with hidden units,

> the error surface may contain many local minima, so it is possible that steepest descent in weight space will get stuck at poor local minima. In practice this does not seem to be a problem.... One reason for this is that there are typically a very large number of qualitatively different perfect solutions. (Hinton, 1989, p. 199)

In short, though it is highly likely that the system will choose a poor solution to a given problem, if the rate of learning is slow enough, this does not in fact happen because, as the old adage has it, "there are many ways to skin a cat."

Considerable progress in problem solving in connectionist models has been gained by designing nets to mimic the dynamics of thermodynamic systems which are well known in theoretical physics. Hopfield, for example, had the intuition that interactions between representations in nets may be governed by the same laws as govern the behaviour of spin glasses. He later developed a net to prove his case. This was an important development in model design for several reasons: first, all of the theoretical framework of the physical system was available to modellers; second, the theoretical framework brought with it new ways of visualizing the dynamics of the networks, those of the dynamical systems mentioned above; and third, insight was gained into the capabilities of these networks and how improvements could be designed. Indeed, an improvement to the original Hopfield net, later called the Boltzmann machine because the network dynamics reflected Boltzmann's contribution to statistical mechanics, was developed. This type of net was based on the physics of the annealing of metals which, if it is carried out gradually enough, will reach a "global" energy minimum. In net terminology, this is equivalent to finding a global or optimal solution to a problem, rather than a local or non-optimal one as would be found by the Hopfield net (for a detailed discussion of these nets, see Churchland and Sejnowski, 1992, chap. 3). Further modifications to these designs are presently being explored by connectionist modellers in an effort to optimize performance (see Hinton, 1989, for a recent overview).

To summarize briefly, then, the internal dynamics of a connectionist system can range from relatively simple to extremely complex, and different models can have very different dynamics depending on the task given the model, a task which it should be remembered is generally dictated by the observed behavioural performance of living organisms. The guiding principle appears to be how best to design the dynamics so that the error of the system is minimized. Since error reduction is intimately dependent on the parameter adjusting procedure employed in the net, an examination of such procedures is next on the agenda.

Learning modes and parameter adjusting procedures in connectionist models

Learning in connectionist nets is accomplished by adjustment of connection weights. These weights, therefore, are the parameters that are adjusted in the parameter adjusting procedures or learning algorithms. In associative nets, the weights altered are those of the connections impinging on the output units only, since these nets generally have only two layers: the input layer and the output layer. In nets with hidden units, however, the weights which are adjusted are not only those on the connections from the input layer to the hidden layer (or layers), but also those on the connections from the hidden layer to the output layer. Since the process of learning is considered to be a process of error minimization of the system as a whole, a decision to adjust a weight must be made on the basis of a judgement of whether or not a change, positive (increasing the weight) or negative (decreasing the weight), would result in an improvement in overall performance of the system.

There are two basic modes of learning in networks: *supervised* learning, in which case an evaluation of the weight setting job is supplied from the external environment, and *unsupervised* learning, in which no such report is available. With both these modes of learning, true to the interactive nature of the "ordered triple" (architecture, dynamics, parameter adjusting procedure), performance is dependent on the architecture and internal dynamics of the net. As was indicated above, supervised learning, however, is additionally dependent on an external evaluation for its performance (see Churchland & Sejnowski, 1992, p. 97). Within each of these modes, learning can also be monitored, which is the case when there is an internal measure of error, or unmonitored, in which case no measure of internal error is employed. Thus, there exists in actuality, a taxonomy of learning procedures (see Churchland & Sejnowski, 1992, p. 98). In the supervised learning mode, then, there is an external source of knowledge or "teacher" which has access to the "correct" answer to guide the development of the network during learning. Monitored networks, on the other hand, use acquired knowledge to improve performance. For example, some monitored nets can make use of past experience to make predictions about new input, the discrepancy between the prediction and the actual input being fed back into the system to improve further predictions. But what happens in unsupervised, unmonitored nets? How do they learn? It would appear that even in these cases, which are thought to reflect some aspects of development in real nervous systems, there is some control over what is learned because there is "an implicit objective function that is optimized during the learning process" (ibid., p.98). In both unsupervised and supervised learning, then,

> the point of the learning algorithm is to produce a weight configuration that can be said to represent something in the world, in the sense that, when activated by an input vector, the correct answer is produced. Nets using unsupervised learning can be configured such that the weights embody regularities in the stimulus domain. For example, if weights are adjusted according to a Hebb rule, then gradually, without external feedback and with only input data, the net structures itself to represent whatever systematicity it can find in the input, such as continuity of boundaries. This means that unsupervised nets are useful in creating feature detectors, and consequently unsupervised nets can be the front end of a machine whose sensory input must be encoded in some perspicuous fashion before it is sent on for use in such tasks as pattern recognition and motor control. (Churchland & Sejnowski, 1992, p. 97)

The learning algorithms themselves can be very simple, as for example a "Hebbian" procedure conceptually based on the seminal insight of Donald Hebb (1949) who proposed that, in the nervous system, synaptic weights could change according to the activities of the neurons at both the pre-synaptic and post-synaptic sites. In a Hebbian algorithm, therefore, the change in weight is dictated by the product of the activities of both the input unit and the output unit. This simple type of learning procedure, however, is only viable for certain types of input vectors: for Hebbian learning to be effective in connectionist models, the input vectors must be orthogonal. If they are merely linearly independent, the algorithm is not optimal. For more complex cases, even in simple linear associators, versions of the "least squares" learning procedure is more efficient. In this procedure, a measure of the error for each input-output pair is calculated using the difference between the actual state of the output and its desired state. Due to the fact that the "desired state," that is, the answer, is known, this is a supervised learning procedure. A function of the calculated difference then determines by just how much the weight of the connection between a given pair of units should be changed[21].

The best-known of the connectionist learning algorithms, also a supervised procedure, is backpropagation of error, which has been described as an automated multi-layer least squares procedure. That

[21] See Hinton (1989) for more details on the calculations involved.

is, it "is a generalization of the least squares procedure that works for networks which have layers of hidden units between the input and output units" (Hinton, 1989, p. 198). How the procedure works in very general terms is as follows: during the training process, each weight must be changed by an amount that is proportional to the rate at which the error changes as the weight is changed, a procedure whose level of complexity depends on the type of units involved. Once the appropriate changes have been made for the output layer, the same procedure must be carried out for each layer in the network, "moving from layer to layer in a direction opposite to the way activities propagate through the network. This is what gives back propagation its name" (Hinton, 1992, p. 146). The backpropagation algorithm has proved efficient at training networks to perform a wide variety of tasks accomplished by humans, and has even found some useful applications in the field of neuroscience. For example, Andersen and Zipser (1990) trained up a network to respond to visual stimuli using the backpropagation of error algorithm, and discovered that the response properties of the units mirrored very closely those of neurons in the posterior parietal cortex thought to be involved in spatial representation. Thus, as advocates of the co-evolution of neural network modelling and neuroscientific research have maintained, this model may have helped further neuroscientific knowledge by elucidating something of the function of these neurons.

In spite of the many documented successes of the backpropagation algorithm, however, many have protested that it is hardly likely that the brain itself would employ such a method of error minimization. And, indeed, though some of the objections put forward have been shown to be relatively trivial, for example the mode of relaying error feedback[22], others are considered more weighty (see P. M. Churchland, 1989, pp. 181-188; and Hinton, 1992). The most serious objection to biological validity for supervised learning models is the need for a "teacher" to control what is learned. As Hinton himself points out, most learning in natural settings must do without a detailed description of what to represent and how to go about representing it. In the interest of approximating more closely biological systems, therefore, modellers have developed several unsupervised learning procedures.

> All these procedures share two characteristics: they appeal, implicitly or explicitly, to some notion of the quality of a representation, and they work by changing the weights to improve the quality of the representation extracted by the hidden units. In general, a good representation is one that can be described very economically but nonetheless contains enough information to allow a close approximation of the raw input to be reconstructed.... Almost all the unsupervised learning

[22] In nets, information must travel in both the forward and backward directions on the same connections. This does not happen in neural networks, but the many feedback connections from "higher" layers to "lower" ones nullifies the significance of this objection.

> procedures can be viewed as methods of minimizing the sum of two terms, a code cost and a reconstruction cost. The code cost is the number of bits required to describe the activities of the hidden units. The reconstruction cost is the number of bits required to describe the misfit between the raw input and the best approximation to it that could be reconstructed from the activities of the hidden units. (Hinton, 1992, p. 148).

Two simple strategies for minimizing these two costs are *principal components learning* and *competitive learning*. In principal components learning, a small number of hidden units cooperate with one another to collectively represent the input pattern, whereas in competitive learning, a large number of hidden units compete against each other to represent a particular input pattern, the winning hidden unit being "the one whose incoming weights are most similar to the input pattern" (ibid., p. 149). In these two learning strategies, then, we have examples of distributed representation of the input (principal components), and localized representation (competitive). Though both of these do mimic some aspects of representation in the brain, neither one is considered maximally biologically valid. However, some learning strategies have been developed which lie between the extremes of purely localized and purely distributed, and these, it would seem, do have their counterparts in brain activity. The brain appears to use the strategy of "population coding," so called because information is represented by the activity of a whole population of neurons grouped closely together: the information in population codes is both localized to the particular population, and distributed across its members, so that even if some of the members of the population are damaged or slow in responding, the group still works to give an output. Furthermore, since population coding can be viewed as a "convenient way of extracting a hierarchy of progressively more efficient encodings of the sensory input" (ibid., p. 150), it may be a method developed by the brain to minimize code costs (see Hinton, 1992, p. 151).

Commentary

There seems little doubt that connectionist modellers are acutely aware of the need to develop biologically realistic algorithms if they wish to model how the brains of living organisms function. They are constantly incorporating important new biological constraints into their theorizing. There also seems little doubt that, whatever procedures they arrive at, they will only be approximations of actual brain functioning given the impossibility of exactly duplicating the physical hardware. According to Churchland and Sejnowski, however, this may not pose as serious a problem as some would suggest. They posit that, even in networks with identical architectures and dynamics, whether biological or artificial, if a large number of parameters are involved, it is extremely unlikely that the learning algorithm used will endow the trained up nets with identical sets of weights. This would be just as true in the case of homozygotic twins as in the general population. Indeed, it has been shown that neurons from identical locations and with identical functions from identical clones of the species Daphnia magna (the water flea) vary substantially in structure, which would suggest that, at a minimum, they form different connections with different weights (Edelman & Reeke, 1990, p. 215). Nevertheless, if learning algorithms can be considered on an abstract level as "efficient devices for searching parameter space for combinations of values that optimize some input-output function" (Churchland & Sejnowski, 1992, p. 134), it is likely that, in two systems of similar architecture and dynamics, the algorithm will delineate two overlapping *regions* of parameter space. "Thus when a model net is highly constrained by neurobiological data, the probability is nontrivial that the region of parameter space defined by the model and the real neural net will overlap" (ibid.). If they do overlap, not only will the artificial net have been shown to be biologically valid, but it can be manipulated experimentally in a manner which would be unthinkable in real nervous systems, and thereby become a source of *new* information on the functional aspects of the system under study.

The take-home point that Churchland and Sejnowski (and presumably all connectionist modellers) wish to make is this: model nets that are constructed to be relevantly similar to biological networks may be useful in discovering features of the model system which suggest avenues of further research within neurobiology. Once these ideas are tested, the information gained can be fed back into new models initiating a mutually beneficial co-evolutionary spiral. This approach, they contend, is eminently more sensible than two separate disciplines, both bent on discovering how the brain functions marching along side by side in disregard of each other's findings.

Attention in connectionist models of learning

It should be recalled at this juncture that the action of attention at the cellular level is to *selectively* enhance and/or inhibit the activity of neurons involved in information processing. Since it is the level of activity of the cells which appears to determine the subsequent changes in synaptic weights which are the result of the learning process, it is by affecting cellular activity that attention has an effect on learning. The set of synaptic weights in place within a system, then, is the embodiment of the knowledge of the system at any particular time, and has been determined in part, at least, by which information has been allocated attention. The emphasis placed on the "selective" action of attention is important because the primary function of the attention system in the brain is "to control the amount and the temporal order of information passing through the system" (Mozer, 1991, p. 65). As has been noted previously, however, attention rarely acts unilaterally, usually interacting with expectation and intention emanating from "higher centres" in the brain to control the nature of information passing through the system in addition to the amount and selection of what gets processed in what order.

With these characteristics in mind, it is interesting to examine how attention might be represented, as it should be if biological realism is sought, in connectionist models with no explicit attentional mechanism. To consider the selective action of attention first of all, it would appear that, in both supervised and unsupervised models of learning, the amount of information processed is selected in large part by the experimenter, in that the machine is exposed to only an amount of information suited to its capacity, and presumably all of the information presented is processed. It would also appear that little consideration is given to temporal order since the input is processed in parallel. This apparent fixing of two of the important functions of attention may be a natural consequence of the focus of the model, namely learning, a focus which presupposes an overriding interest in the nature of the information processed. In supervised learning models, the selectivity of attention on the kind of information processed is incorporated into the "teacher" function of the algorithm. The system is essentially "told" what to pay attention to, and is "corrected" until it produces the required answer. In these models, then, the external feedback serves as a "higher centre" source of selection which controls the nature of the information attended to. It is the particular parameter adjusting procedure employed in the model that performs the specific activation/inhibition action of attention at the "cellular" or unit level according to the directions given from on high, but no explicit directions as to what information to encode at the hidden unit level is given. Though, in unsupervised learning models, there is no external "teacher" function, there does appear to be some direction to the system with regard to the nature of the information encoded. Remember that, in these models, there is an "implicit objective function" that is optimized during the learning process. However, since unsupervised systems work to enhance some "quality of a representation" at the hidden unit level, and since this "quality" must be dependent on the input, the selective action operating within the system would appear to be driven by "bottom-up" considerations rather than "top-down" as in the case of supervised learning. The optimization process, of course, is once again embodied in the learning algorithm or parameter adjusting procedure.

The comments of Churchland and Sejnowski with respect to external feedback being necessary to teach the model what to "care" about are of interest here, because they speak to the issue of the influence of the "value" of the information on attention. In supervised models, the system is given *a priori* what is of value in the information in that the required output is worked towards systematically. What is not given, however, is what is of value at the hidden unit level in order to produce the final valued answer: the system must extract this by means of the application of the learning algorithm. And, indeed, it is in discovering the features of the various input samples which are of value in every case that the model learns to generalize to new samples. This process, then, can be viewed as a seeking out of what is "familiar" in the input samples, given that only input features that are reliably present can be considered of value in a decision on whether or not a new input sample belongs to a certain category. Thus, learning to generalize itself may be seen as a method of determining what features of an input set are valuable for categorization.

So much for value in supervised learning systems. But what of unsupervised models where no directions as to what to care about are given? It would appear that, even in these cases, the system seeks out what is of value in the input in a manner similar to that employed in supervised models. It should be recalled that, in the absence of external feedback, an unsupervised net seeks out regularities or, as Churchland and Sejnowski (1992) put it, "whatever systematicity it can find in the input" (p. 97). This is just as true for principal components learning models as for competitive learning models. In principal components models, the sifting out of useful features at the hidden unit level becomes a task of allotting a unique feature of the input to each hidden unit. In models of competitive learning, on the other hand, each hidden unit comes to represent a kind of prototype of a cluster of input samples, and hence must have embodied in its weights a set of features necessary to typify a small category. Thus rather than the features of value for categorization being partitioned among the hidden units as in principal components models, they are partitioned among the weights of a particular hidden unit in competitive learning. So, in these unsupervised models, as in the supervised ones, once they are trained up, what is "familiar" in the input could be said to be "capturing" the attention of the system, but the driving force of the mechanism appears to be more bottom-up than top-down. It seems fair to conclude, therefore, that, in spite of the control exerted on the selection on the amount and temporal order of the information processing, many important aspects of attention are represented implicitly in connectionist models without an explicit attentional mechanism.

There are connectionist models of brain function in which attentional mechanisms are incorporated explicitly, however, and in these models much of what is known about attention is built in to the model. For instance, in an elaborate model of visual perception, Kosslyn and Koenig (1992) have included an attentional window within a visual buffer module: this window pinpoints the location in the buffer in which the information is preferentially passed on for further processing. Moreover, the attention window may be shifted by either a stimulus-based mechanism which operates on the basis of significant novelty in the input, or by a "higher level" shifting mechanism dependent on the output of an associative memory and a recognition module which initiates a "disengage-move-engage" mechanism based on the attention shifting model of Posner (see Kosslyn & Koenig, 1992, p. 105 for a diagram of the system described). In a connectionist working model labelled MORSEL (multiple object recognition and attentional selection) developed by Michael Mozer, an attentional mechanism which acts as a retinotopic map of locations can be driven in a stimulus-based manner, and/or by higher level expectations and task demands. This model, of course, reflects very closely the premises set out by Van der Heijden. It serves four distinct functions: controlling order of readout, reducing crosstalk, recovering location information, and coordinating processing performed by independent modules, all of which have been identified as among the primary functions of attention in the brain (see Mozer, 1991, chap. 4). Interestingly, Mozer discovered that "highly familiar stimuli outside the focus of attention can work their way through the system better than other stimuli" due to the fact that the letter and word recognition module of the system "is able to recognize familiar stimuli based on fewer perceptual features than less familiar stimuli" (pp. 68-69). Mozer goes on to comment: "If this result is typical, it would appear that focal attention is less critical for highly familiar stimuli. Although this property is unlikely to distinguish MORSEL from other theories of attention, it is reassuring to verify that the computational model does indeed behave as expected" (p. 69). As this last comment highlights very well, there is clearly a concern on the part of connectionist modellers to represent attention as accurately as possible when it is explicitly represented in a model of brain function.

2.3.4 Conclusions on the validity of connectionism

The detailed and extensive discussions of connectionism and its underlying constraints elaborated in this chapter have been entered into with but one purpose in mind: to serve as a basis upon which to come to some conclusions on the validity of connectionism as a method of modelling of how the brain learns. Before making any decisions on these matters, however, it is necessary to have a clear understanding of what aspects of "validity" are appropriate to this situation. There are two separate though mutually dependent variations of what it means for a concept to be "valid" which are pertinent here: (1) sound, just, well-founded, and (2) producing the desired effect. Based on these two understandings of validity, then, the questions which must be answered in this section are:

(1) To what extent can connectionism be considered a sound, just, well-founded model of learning in the brain?

(2) Do connectionist models reproduce effective learning?

Connectionism as a sound, just, well-founded model of learning

As has been argued at length in this chapter, connectionism is solidly founded on neuroscientific research at every level of brain organization. Although it can always be contended, and is frequently so among members of the neuroscientific research community, that the results of such research do not speak "the truth, the whole truth, and nothing but the truth" about how the brain functions, it does seem safe to say that the insights gained by modern day neuroscientists are as detailed and precise as modern technology allows. Moreover, the results underpinning these insights are under constant review with respect to their veracity and coherence with the enormous body of data that has been garnered on brain function to date, not only in the neurosciences but also in psychology. Suffice it to say, then, that the data base on which connectionism is founded can hardly be faulted for a lack of comprehensiveness or fine detail.

Of course, in order to build its working models, connectionism abstracts from the biological data physical and mathematical descriptions of brain function, and in so doing translates the languages of anatomy, physiology, and biochemistry into those of architecture (as it relates to artificial intelligence models), dynamics, and mathematics. None of these could be considered a poorly developed discipline lacking a sufficiently sound empirical data base or well-tested theoretical premises. With respect to the "translation" of biological data into a mathematically grounded model of a system, however, questions could be posed as to which elements of the biological system it is most critical to emulate in a model. But, as has been mentioned in the discussion, these questions are posed and hotly debated by connectionist modellers themselves, different decisions leading to different models of the same system which are then compared for their efficacy in mimicking the system under study, as well as for their deficiencies. Once the architecture, dynamics, and parameter adjusting procedure of a model have been decided upon, therefore, the soundness of these decisions is checked against the accuracy of the performance of the model. Thus we come to a discussion of the second criterion of validity dealing with the effectiveness of the model.

Do connectionist models effectively reproduce learning as it occurs in the brain?

Before getting to the meat of this question, it is useful to remind ourselves that the basic organization or architecture of the brain is geared towards the formation of systems, both microscopic and macroscopic. Antonio Damasio (1994) encapsulates this notion in these words:

> In short, then, the brain is a supersystem of systems. Each system is composed of an elaborate interconnection of small but macroscopic cortical regions and subcortical nuclei, which are made up of microscopic local circuits[23], which are made up of neurons, all of which are connected by synapses. (p. 30)

Hence, learning in the brain is essentially learning for the purpose of producing an effectively organized functional unit. In spite of the fact that knowledge of the changes that occur during the learning process at the molecular level is available and is

---

[23] Here Damasio is distinguishing between the terms "system" and "circuit" (or "network") which are often used interchangeably in the literature. "Circuits" he considers to denote microscopic systems, while "systems" are considered macroscopic.

incorporated into theories of learning at the neuronal level, the lowest level of simulation of biological learning attempted by connectionist models is that of the neuronal group or microscopic module. At the upper end of the scale would be models of entire sensory systems constructed of a number of interacting macroscopic modules, all based on connectionist architectural and dynamical principles at the microscopic level.

Since the artificial intelligence literature abounds with information on such models, and since many of them have already been referred to in this chapter, an overview of connectionist accomplishments would seem to be sufficient here[24]. Beginning at the lowest level, then, the physiological "binding" effect of oscillations on neuronal activity thought to be due to attention and therefore implicated in learning has been successfully mimicked; simple forms of learning such as the development of the response properties of cells in the visual cortex, and ocular dominance columns, both thought to be the product of competitive interaction between neurons, have been modelled, as have aspects of general development in the brain such as the formation of topographical maps; there are models of object recognition and categorization; the ability to encode spatial relations using the principle of population coding so as to recognize complex objects

---

[24] Readers are referred particularly to Rumelhart, McLelland, & the PDP Group (1986); McLelland, Rumelhart, & the PDP Group (1986); Churchland & Sejnowski (1992); Kosslyn & Koenig (1992). References to specific models can be found in these texts.

like faces rather than just a group of unrelated constituent parts can now be artificially reproduced; models of sensory motor integration of the type that occurs in the vestibulo-ocular reflex are now working; the operation of whole sensory systems such as the visual system has been modelled, and through careful lesioning of specific modules, the effects of brain damage in these areas have been reproduced. Perhaps these few examples suffice to give at least a glimpse of the achievements of connectionist modellers at what might be considered the "lower end" of the intellectual scale, though it must be emphasized that none of these forms of learning is lacking in complexity nor employs principles which would be foreign to learning at the "upper end" of intellectual ability.

Accepting for the time being that there is a difference between these forms of learning and "higher" learning, what has been achieved at this "upper end" of the learning scale? To give but a cursory overview, there are models of the formation of schemata and general thought processes, the encoding of memories, speech perception, the learning of past tenses of verbs in English, word recognition, sentence processing, and reading. Part of the success of many of these models is that they show organizational principles, and "deficits" or "developmental" errors comparable to those documented in observations of human learning in the same areas. Take, for example, NETtalk, a system which learns to read text aloud without being given any guidelines whatsoever with respect to the phonetic significance of standard English spelling:

an analysis of the partitioning of the hidden unit vector space revealed that it had been split into 79 subspaces each one corresponding to one of the distinct letter-to-phoneme associations in English. Furthermore, the subspaces were divided into two distinct classes which effectively distinguished between consonants and vowels. "Looking further into this hierarchy, into the consonant branch, for example, we find that there are subdivisions into the principle consonant types, and that within these branches there are further subdivisions into the most similar consonants" (P. M. Churchland, 1989, p. 175; see also a schematic representation of these divisions, ibid., p. 176, or Churchland & Sejnowski, 1992, p. 119). This type of result is extremely encouraging, since the organization the machine memory exhibits appears logical and, moreover, is consistent with what might be expected on the basis of a rational, critical thinking approach in humans. As a consequence, it speaks for the possibility that "critical thinking" might be the natural order of things in brains, rather than a technique that must be learned.

Another method of measuring the validity of connectionist models which involve "higher" forms of learning is to examine to what extent they explain psychological findings in the same area, whether dealing with the performance of "normal" brains (that is, undamaged) or those of brain-damaged patients. Connectionist modellers regularly undertake such explanatory attempts[25]. In his discussion of his object recognition model MORSEL, for instance, Mozer (1991) devotes a chapter to "Psychological Phenomena Explained by MORSEL", a chapter which deals with topics covering a range of categories: "basic phenomena," "perceptual errors," "other perceptual and attentional phenomena," and "neuropsychological phenomena." Furthermore, connectionist modellers evaluate the successes *and* failures of their models in light of how well psychological phenomena are elucidated, and plan extensions or corrections to the models based on these evaluations (see, for example, Mozer, 1991, chap. 7). The willingness with which connectionists embrace comparisons of this sort speaks well for the matter of the validity of their models.

Nevertheless, connectionists themselves would be the first to admit that the models are not perfectly valid representations of the workings of biological systems. Indeed, many discrepancies in the performance of the models have been marked. In some cases these are temporal in nature: the models are too slow. In others, an inordinate amount of effort is necessary to train up the system as the number of hidden becomes large: this is referred to as the "scaling" problem. In still others, the performance of the model is very good on a certain task, but the algorithm used is not biologically plausible. Even when the architecture, dynamics, and

---

[25] In Kosslyn and Koenig (1992), for example, there are many comparisons of the performance of models to psychological data.

parameter adjusting procedure are maximal in a model, the performance may not be adequate due to a lack of pertinent data on the particular brain system under study.

All of these problems are treated seriously by connectionist modellers, happily, and progress is being made on all these fronts and many more. Because of this assiduous attention to detail by connectionists, not to mention the ever burgeoning body of neuroscientific data now available, the number of successful models is growing as is their sophistication. There is little doubt that connectionism has achieved feats of brain modelling that have evaded conventional techniques. There is also little doubt that it has done so through the careful abstraction and translation of critical features of brain structure and function into machine language. Consideration of these facts leads inevitably to the conclusion that connectionism, though still in a relatively early developmental stage, has already uncovered some tangible realities about learning in the brain, and must, therefore, be considered a valid methodology to model how the brains of humans and animals learn.

As was mentioned in the introduction, however, the promise of connectionism is not just to elucidate the mechanisms behind the plasticity of the brain, but also to explicate "its intricate dependence on the extended cultural surround." This issue has not been broached directly in this chapter, but it will be dealt with in some detail in the next, at least as far as its impact on language learning is concerned.

2.4 Summary of Chapter Two

The aim of connectionism is to elucidate the nature of conscious intelligence. Because of this, its theoretical premises fall partially into the territory of philosophy, specifically into the materialist camp due to the fact that its underlying concepts are drawn from the disciplines of neuroscience, cognitive psychology, and artificial intelligence. Connectionism has been severely criticised by philosophers of the non-materialist persuasion who believe that the workings of the mind will never be reducible to those of the brain. It has also received criticism from within the connectionist camp with respect to what are perceived to be its overblown claims to explanatory power. What connectionists are seeking is explanatory unification of theories through intertheoretic reduction. However, they point out that if the theory to be reduced is fundamentally flawed, successful reduction by a new and more correct theory cannot be expected, and elimination of the erroneous theory should result. Many connectionists believe that this will be the case with folk psychology, our common sense understanding of how the mind works. Other theorists disagree, noting that if folk psychology is eliminated, this will mean the elimination of all of its premises, many of which are undoubtedly correct.

Smolensky has established that, within the study of cognition, connectionism falls into a subsymbolic paradigm rather than the symbolic paradigm adopted by traditional artificial intelligence research. He also points out that it does not operate at the neural level of neuroscientific research, and therefore connectionist models should not be termed neural models. Since for traditional artificial intelligence models, the preferred level of analysis is the conceptual level, connectionist models can be considered to be operating at a subconceptual level, the activity of a unit or a group of units having no recognizable meaning for us. A lack of absolute realism in models does not overly concern connectionists, however. They reason that models do not have to be 100% realistic to be good and useful models. It suffices that care is taken to include critical elements of the function of a system under investigation in a model. The goal of connectionist models is not to achieve an accurate description of a cognitive system, but to give an account of the causes of intelligent behaviour. This it has done through the construction of successful models of many aspects of brain function relevant to intelligence in biological systems. An examination of the appropriate classification of connectionism found it to be a methodology which leads to the production of models, most importantly working models, of brain functioning. The purpose of the chapter is to evaluate connectionism in its capacity to model how the brain learns.

Since the aim of connectionist models is to depict learning in the brain, attention must be paid to the biological constraints imposed by the structure and function of the particular system being modeled. Special consideration should be given to the level of function the model is attempting to elucidate. Churchland and Sejnowski have identified seven levels of brain organization: a model aiming at a certain level must take into account information pertinent not only to this level, but also to the levels above and below so that overall theoretical consistency is maintained. In connectionist theory, learning is instantiated in the changing of the strengths or weights of synapses between neurons, and this is seen to be directly related to cell activity. The design of a model must therefore incorporate elements of the anatomy, physiology, and processes governing synaptic change of the system under study. As the neuron is the basic functional unit in the brain, its general structure and function must be considered first and foremost in models of brain function. Nevertheless, because neurons are arranged into networks, other levels of brain organization such as the network and map levels must also be taken into account when designing models, as must organizational characteristics within these levels, for example, layers and columns. In spite of the fact that the brain has many levels of organization, it should not be considered well-designed in an engineering sense. It appears to be more of a "kludge", something that works in a messy and inefficient way, but gets the job done. The kludge-like character of the brain is the result of the evolutionary process which generally works to improve the anatomic structures already in place rather than start from scratch.

Another important feature of brain function to be represented in a model is the massive parallelism of information processing. This arises out of the formation of specialist subsystems which work with relative autonomy while still allowing integration of information to influence output. The main benefit of parallel processing for an organism is speed of output. Besides these major considerations listed above, other biological constraints that have emerged from psychology must be kept in mind: a few of these are graceful degradation of brain function, ability to cope with conflicting or inadequate data, and ability to recover function after damage.

Attention is recognized as a major force in learning. Its action is manifested at the neuronal level by either an activation or inhibition of cellular activity. A clear definition of attention is hard to come by due to the plethora of opinions as to the exact nature of its action and anatomy, and to the complexity of the phenomena associated with what has been called the "attention system." The mechanisms by which this system operates are currently under intense investigation and are the subject of hot debate in attention research circles. In this thesis, the selective phenomena associated with attention are taken to be more important than intensive phenomena such as arousal and alertness. The selective mechanism of attention can be either stimulus (world) driven, or actively directed by the person (or animal). The former is labelled passive attention, the latter, active or voluntary attention. The ability of a stimulus to capture attention even when voluntary attention is in force is related to its "importance" to the organism, and is referred to as the saliency of the stimulus. This importance or value is thought to be determined largely by an evaluation of the relevance of the stimulus to the scenario at hand, an evaluation which is dependent on the emotional state brought on by the stimulus.

There are several dichotomous concepts associated with attention: active and passive attention; parallel and serial processing of information; automatic and controlled processing; limited and unlimited capacity. Under close scrutiny, however, the dichotomous representation of these concepts does not hold water. As a result, such conceptualizations are considered merely heuristics, rather than representations of reality. The difficulties encountered in describing the nature of attention has led to the drawing up of several metaphors. Of these, the most widely used are the "filter", "spotlight", and "zoom lens" metaphors. Some newer conceptualizations are the "gradient", "attention gating", and "gradient filter" models. However, none of these metaphors or models is considered adequate at this stage of investigation. A "system internal" conceptualization of attention is put forward as affording valuable insight into the workings of attention within a sensory system. The particular system described is the visual system. In this model, selective attention is considered to be stimulus driven, and separate from the effects of expectation and intention to act, both considered voluntary aspects of attention. It is postulated, however, that the three work together to integrate bottom-up and top-down effects on the attention mechanism. This model distinguishes itself from others in that attention is not depicted as being external to the sensory system. Because the model describes only attention within a sensory system, however, it does not account for the ability of attention to select a particular "channel" to attend to. Nevertheless, similar mechanisms can be postulated to be active in attentional selection of a channel. The channels attended to may be an individual ear, or information delivered by the visual or auditory systems alone, to give but two examples, and higher centres in the brain important to expectation and intention to act are thought to exert control over the choices made. In spite of this voluntary control, stimuli which have a high importance weighting may capture attention. The high importance weighting is seen to be closely related to the familiarity of the stimulus to the subject. A working definition of attention is adopted: Attention is a facilitatory mechanism for the selection of information for further processing. It operates in close cooperation with expectation and intention to act, and can be driven by voluntary- or stimulus-initiated mechanisms.

The basic design of a connectionist model is that of a net of neuron-like units. The learning process in these models, as in biological systems, is seen as a process of error minimization. In connectionist models, the architecture, dynamics, and parameter adjusting procedure (or algorithm) are chosen to mimic the anatomy, physiology, and procedures governing synaptic change as closely as possible. However, each of these features of a net is only a gross approximation of the circumstances in the real biological system under study. With respect to how it is that biological systems can compute, it is proposed that a system in the brain is a computational system if there is a revealing mapping between the system's physical states and the elements of the function computed. Typical nets employ electrical units modeled both structurally and dynamically on neurons. As in the brain, these "neurons" are arranged into highly interconnected layers. Each unit in a particular layer processes the same kind of information simultaneously before passing on a message with respect to its activity in a distributed manner to the next layer. The simplest net architectures consist of an input layer and an output layer, and are capable of only simple associative learning. As in more complex models, nevertheless, the knowledge they contain is distributed across many units. It is because of the simultaneous or parallel processing of information as well as the distribution of the knowledge that connectionist models also go under the name of PDP (parallel distributed processing) models. Nets with hidden units, units placed between the input and output layers, form more powerful models in that they endow a global perspective on the net in spite of its local connectivity. The function of the hidden units is to discover important regularities in the information contained in the input. If the information flow in nets is only from the inputs to the outputs, the net is called a feedforward net. If on the other hand, information can flow back to a unit from other units in the net, the net is said to be recurrent. Recurrent nets are thought to model a form of short-term memory. Some nets are actually systems of specialist mininets, an arrangement which finds its counterpart in the integration of specialist subsystems in the brain.

The brain has been described by Luria as a dynamic functional system. In the same way, connectionist nets are dynamic functional systems made up of individual units all working in a coordinated fashion to give the final output. The knowledge of such systems, biological or machine, is encoded in the sets of connection weights between layers of units. The benefit of describing a net in these terms lies in the theoretical power inherited from the physics of dynamical systems. Nets with interesting properties can be designed based on knowledge of the behaviour of the dynamical model adopted. For similar reasons, nets have been designed to mimic the dynamics of well-known thermodynamic systems. As has been noted, learning in nets is accomplished by changing the connection weights between units and this is carried out by the learning algorithm. There are two basic learning modes in connectionist nets: supervised learning and unsupervised learning. In the former mode, an external evaluation of the weight setting job is available, while in the latter it is not. Models with an internal measure of error are also possible within both these modes. The goal of any parameter adjusting procedure is to produce a weight configuration that represents something about the world. Thus, the complexity of an algorithm can range from very simple (e.g. "Hebbian" learning, an unsupervised algorithm) to very complicated. Some recent connectionist models employ combinations of simple and complex algorithms to attain their learning goals. One well known algorithm using fairly complex mathematical calculations is backpropagation of error, a supervised learning algorithm which has been highly successful in modelling learning in biological systems. The backpropagation algorithm, however, as other supervised learning algorithms, is not considered biologically plausible, whereas unsupervised learning algorithms have much greater biological validity. Principal components learning and competitive learning are two strategies that have been developed to mimic learning in some biological systems. Furthermore, even more powerful learning systems can be constructed using algorithms that combine elements of both these strategies, and they too find resonance with what is known of the function of real systems in the brain. The main point that connectionist researchers wish to make, then, is that the disciplines of artificial intelligence, neuroscience, and psychology should not just co-exist side by side, but should co-evolve, each one informing the other for the purposes of furthering knowledge of the brain.

Given that attention has been shown to be important in the learning process, it seems necessary to consider where its action is represented in connectionist models. In models without an explicit attentional module, attention is seen to be active in the selection of the nature of information processed. In supervised models, this selectivity is incorporated into the correction or "teacher" function of the algorithm; in unsupervised models, it is the optimization of the "implicit objective function" of the system which controls the nature of the information encoded. Thus, what is of value to the system is seen to be imposed explicitly in supervised learning models, while it is imposed implicitly in unsupervised models. However, in both types of model, what is of value at the hidden unit level is left to be discovered by the hidden units themselves, rather than being spelled out explicitly. It is argued that what is important for the organism at the hidden unit level is information that is regularly present or familiar in the input set. In models in which attention is represented by an attentional module, many of the results of attention research on information processing in the brain are repreoduced, including the ability of familiar information to capture attention.

Finally, the validity of connectionism is evaluated on the basis of two questions: (1) Is connectionism a sound, just, well-founded theory of learning? (2) Do connectionist models effectively reproduce learning as it occurs in the brain? On both these counts, connectionism is jugded to be valid, though not entirely so.

Chapter 3. CONNECTIONISM, ACADEMIC ACHIEVEMENT, AND SECOND LANGUAGE LEARNING IN IMMERSION SITUATIONS

3.1 Introduction: The Challenges Facing the Connectionist Perspective

It was established in chapter 1 that a good deal of the criticism that has been levelled at Cummins' Threshold Hypothesis by the second language research community is well justified, in spite of Cummins' evident intentions to aid second language learners and the benefits that his theorizing has brought to them. Of particular concern among the points raised by critics were the following:

1. Cummins' conceptualization of the nature of cognition, especially its tight association with language proficiency, which led to accusations that the Threshold Hypothesis is in actuality a cognitive deficit hypothesis of language learning for second language learners at the low end of the language proficiency scale.

2. A chronic lack of definition of the variables degree of bilingualism (or language proficiency), cognitive ability, and threshold of linguistic competence making the hypothesis impossible to test empirically.

3. The dependence of the hypothesis on studies which have been

shown to be methodologically inadequate.

4. The exclusion of sociological variables either within the constructs of the hypothesis itself or as variables impacting on its explanatory power.

Because of these flaws, as well as others not listed here, the Threshold Hypothesis has been judged to be, at best, useful in only an heuristic fashion, or, at worst, unhelpful. The basic premise of this dissertation, on the other hand, is that connectionism can provide useful insight into any cognitive consequences of second language learning in immersion situations by virtue of the fact that it approaches the issues from the more finely resolved perspectives of neuroscience, cognitive psychology, and artificial intelligence, and now that the foundational tenets of connectionism have been elaborated upon, the question that must be posed is "Can it indeed provide such insight?".

While the answer to that question will not be forthcoming until a connectionist perspective on second language learning has been fully developed and discussed, the differences between Cummins' perspective and the connectionist one which touch on the above criticisms can be made clear here. On the issue of the nature of cognition, for instance, it would appear that, because of the importance he accords to language as an influence on cognition, Cummins envisages cognition as active at a conceptual level. It is almost certainly for this reason that he proposes that learners who do not have an adequate command of their second language may be cognitively disadvantaged. Connectionism, however, would question the level at which cognition acts as well as the specialness of language's ability to sculpt it. On the connectionist view, cognition is thought to act primarily at the subconceptual level rather than, or perhaps as well as, the conceptual; connectionism, then, opens up the possibility that many, and even all, kinds of information processing can influence cognition and therefore reduces the importance of language processing as a means of altering cognitive abilities. Furthermore, the connectionist focus on learning as central to cognition coupled to the ubiquity of learning capability (plasticity) in the brain suggests that there is no locus of cognition per se, though there may be certain brain structures that are necessary to our conscious perception of cognitive processes (see Crick & Koch, 1992). It also suggests that there is little empirical support for the notion of a "basic" core of cognitive abilities, as was proposed by Cummins, if these are conceived of as independent of the influence of experience. It was demonstrated in the first half of this century, for example, that even an ostensibly "non-cognitive" ability such as seeing must be learned through prolonged experience (see Hebb, 1949; see also Crick and Koch, 1992, for insight into seeing as a constructive process). Thus, if perception itself involves learning, then it seems highly unlikely that any aspect of mental processing does not. One area in which Cummins' approach and the connectionist one do agree, however, is the importance of focused attention in facilitating the processing of unfamiliar information. Since the effect of attention will be detailed in the following sections, it will not be dealt with here, and a more in-depth examination of the differences between the two views on cognition will be undertaken in chapter 4.

With respect to the impossibility of testing the Threshold Hypothesis due to the lack of definition of the variables degree of bilingualism, cognitive ability, and threshold of linguistic competence, the aim of the connectionist approach is not only to refine notions on which aspects, if any, of the first two variables may be covarying[1], but also to explore the nature of "linguistic thresholds" with an eye to clarifying whether or not they are real phenomena, and if so, how they may be more accurately conceptualized. More precisely, it will be proposed in a new information processing hypothesis developed on the basis of connectionism that "cognitive ability" can be more accurately represented by "visual processing ability", and "degree of bilingualism" by "comprehension of second language". It will also

---

[1] It should be noted that the position being taken here is that, since all aspects of brain processing are almost certainly influenced by experience, all are cognitive to some degree, even those that have not traditionally been considered to be so. Hence, to suggest that some aspect of cognitive ability may be changing with the acquisition of a second language is not to beg Reynolds'(1991) question of whether or not there is a relationship between degree of bilingualism and cognitive ability at all. What Reynolds was considering as cognitive ability does not appear to be based on connectionist theorizing and is presumably similar to Cummins' conceptualization.

be proposed that, during the process of second language learning, the developmental pattern of certain changes that are likely to occur in brain processing may appear to have threshold-like characteristics if they become information processing strategies. The upshot of this focusing in on only certain features of mental processing and language competence should be to render the new hypothesis more readily testable empirically than Cummins' hypothesis, though it is conceded that decisions as to which aspects of visual processing ability and second language comprehension should be measured and how these should be measured may be contentious.

To speak to the third criticism of the Threshold Hypothesis listed above, that of its dependence on data from studies which are methodologically inadequate, the connectionist approach to be taken here will avoid (or at least minimize) this problem by constructing a model of the second language learner from a bottom-up perspective dependent largely on empirical studies from neurobiology, language psychology, and artificial intelligence research. Though the data gathered in these studies are by no means immune to experimental error or philosophical biases (see, for example, Hoyningen-Huene, 1993), they tend to be more reliably reproducible than those emanating from educational studies, in part because the subjects of study are less complex than the total human being (as is the case with neurobiology and artificial intelligence), and in part because attention to experimental design is considered of primordial importance in these disciplines. As a result, they should form a sounder foundation upon which to base a model of brain functioning.

The last major point of criticism of the Threshold Hypothesis is that it does not take into account sociological variables as factors in determining the cognitive consequences of bilingualism, and indeed, attempts to screen them out altogether by focusing on "child input" variables exclusively. From the connectionist vantage point, however, the intricate inter-relationship of all aspects of brain functioning and experience would render any attempt to screen out environmental influences futile. For connectionism, the environment is all pervasive and is being encoded by the brain consciously and unconsciously in an integral fashion along with the learning of a second language (or mathematics, or art, or science, and so on). Moreover, though the information processing hypothesis forwarded by the connectionist viewpoint, like Cummins' hypothesis, will deal with only two variables, neither of which would appear to be related to environmental influences, motivational, attitudinal, and experiential factors stemming from the environment can all be integrated into the premises of the hypothesis by means of their interaction with attentional allocation to give an accounting of why different learning circumstances lead to different learning outcomes. If the hypothesis resists falsification upon testing, educators can then use the insight it provides to construct optimal learning environments for different groups of second language learners.

In this chapter, then, the challenges to the connectionist perspective that will be taken up will be the following:

1. The bottom-up development of a hypothesis on the changes in brain processing brought on by second language learning in immersion situations.

2. An examination of the explanatory power of the connectionist hypothesis of second language learning developed.

3. An examination of the relevance and implications of the connectionist hypothesis for issues surrounding the interaction of second language learning and academic achievement, as well as second language teaching in immersion situations.

In the final chapter, chapter 4, Cummins' Threshold Hypothesis will be re-examined in the light of the connectionist hypothesis in the interest of clarifying

4. The differences between the models of cognition underpinning both hypotheses.

5. Whether or not the Threshold hypothesis is a deficit hypothesis of second language learning.

6. Whether the connectionist hypothesis can give insight into the reality and nature of linguistic thresholds.

3.2 Development of a Connectionist Hypothesis on the Changes in Brain Processing Initiated by Second Language Learning in Immersion Situations

3.2.1 Development of a general model of learning in the brain

In developing a connectionist model of learning that will be of use to educators, the importance of maintaining a theoretical coherence between levels cannot be overemphasized. This priority falls naturally out of the theoretical unification goal of connectionism, what one might call connectionism's "prime directive." Because of this, the most logical starting points are the synaptic (subcellular) and neuronal levels, which have been extensively studied and are now well understood, while studies of learning at the more fundamental molecular level are still highly controversial and obscure. As has been noted previously, it is largely for this reason that connectionist models of learning take the "neuronal" unit as their base level.

At the subcellular level, then, learning is believed to be manifested by changes in the weights or strengths of synaptic connections between neurons, and is a product of neuronal activity or firing rate. Under normal circumstances, the higher the activity, the more likely it is that changes in synaptic strength will occur. Neurons, of course, do not act in isolation but in association with many other neurons connected into networks.

Eventually, with persistent association of activity between neurons, a more stable synaptic state is reached such that further change becomes difficult, though, of course, change will inevitably occur if the activity of the cells decreases[2]. More often than not, this stable state will be a ceiling condition brought on by the association of the activities of certain neuronal clusters. However, as in the case of anti-Hebbian learning, it may be a bottoming-out of the ability of some neurons to affect others. As this stabilization of synaptic strengths occurs, a system (subsystem or circuit) representing the input data is formed. Once formed, the output of these systems can furnish the input, or part of the input, for neuronal networks further downstream which will themselves be transformed into new knowledge systems given adequate experience. An important point to remember here is that, since learning is a product of neuronal firing, and since neurons are always firing even if only at a baseline level of activity, learning is always going on, though at a rate commensurate with neuronal activity. Because learning is enhanced by high activity in the cells, a state of arousal or alertness, a so-called "intensive" state, will promote learning due to its general excitatory effect on the brain. Any state which brings about a reduction in the activity of the cells, such as drowsiness, would have an opposite effect on learning capability.

---

[2] A prime instance of the type of change being referred to here, and one with which we are all familiar, is forgetting.

For similar reasons, emotional state, whether positive or negative, is also interactive with the learning process, and appears to be highly influential in directing selective attention to certain stimuli in preference to others by assigning a higher value in terms of relevance to the context at hand to these preferred stimuli. According to Damasio (1994), this evaluation of the input data is encoded in the brain as a "somatic marker," a readout of the body state experienced on the reception of the stimuli. Once established, a somatic marker and a representation of input data are stored in tight association with one another such that any new processing of the input data activates simultaneously the corresponding somatic marker. When selective attention is focused, it provides further activation of the cells whose activity represents the selected stimuli, thus increasing the likelihood that learning will occur in association with them. If Damasio is correct in his theorizing, the learning that occurs would take place under a similar body state and therefore would be associated with a similar somatic marker, though emotional reactions and hence body state and somatic markers can be influenced by new experience as can all learned reactions.

Selective attention, it may be remembered, can be stimulus driven or voluntarily driven (expectation and intention)[3] at both the system internal (local) and system external (global) levels. Though

[3] These are sometimes termed *exogenous* and *endogenous* attention respectively.

these aspects of selective attention are discussed separately, it is unlikely that they operate entirely autonomously. This notion is in keeping with the conclusions of the previous discussion with respect to the falseness of dichotomous representations of attentional function. In any given complex situation, therefore, it is highly probable that both stimulus driven and voluntary attention are acting cooperatively. On either the local or global level, there is no doubt that expectation and intention to act are dependent to a high degree on experience, and, as a consequence, what has been learned in the past must have some influence on what is learned in the present, particularly in situations in which attention is voluntarily allocated. However, stimulus driven attentional selection may also be dependent on learning for its allocation. Damasio (1994) has described how learning is associated with emotional response and how such emotional responses are individualized in these words:

> At a nonconscious level, networks in the prefrontal cortex automatically and involuntarily respond to signals arising from the processing of ... images. This prefrontal response comes from dispositional representations that embody knowledge pertaining to how certain types of situations usually have been paired with certain emotional responses, in your individual experience. In other words, it comes from *acquired* rather than *innate* dispositional representations, although, as discussed previously, the acquired dispositions are obtained

under the influence of dispositions that are innate. What the acquired dispositional representations embody is your unique experience of such relations in your life. Your experience may be at subtle or at major variance with that of others; it is yours alone. Although the relations between type of situation and emotion are, to a great extent, similar among individuals, unique, personal experience customizes the process for every individual. (pp. 136-137)

On this view, then, stimulus driven attention, which certainly operates "automatically and involuntarily" and "at a nonconscious level," could quite conceivably be subject to the influence of experience, even though the latter is developed under the control of innate mechanisms, because of the steady accumulation of emotional experience associated with all stimuli. Thus, for any organism including human beings, the influence of innate mechanisms on stimulus driven attention would be strongest in the very young, and would gradually be tempered by experience with age. But when Damasio speaks of individual experience in humans, he is speaking in effect of cultural and social experience. Hence, it would seem a logical conclusion that, by means of their relationship to emotional responses, cultural and social experience are able to direct and maintain attention and thereby be a force in guiding further learning. Obviously, experiences or stimuli which have had positive emotional associations (pleasure or reward) in the past will encourage learners to approach the learning situation in a positive way, and vice versa for situations with negative associations. On a less obvious level, however, it is entirely plausible that attentional allocation itself can be associated with positive or negative emotions with regards to the success of the allocational outcome. Remember that an emotional state is an evaluation of the relevance of input data, and input data can just as easily be the results of a decision made consciously or unconsciously with respect to where attention was allocated. This is, of course, a "cognitive" strategy since it is dependent on learning, and so it is to the effect of this strategic allocation of attention on learning that it seems appropriate to turn next.

As was mentioned above, selective attention can be active at both the local (system internal) and the global (system external) levels, and within both attentional mechanisms can bring about activation and/or inhibition of cellular activity. On the local level, the effects will be felt largely within a small area of the brain, while at the global level whole channels, sometimes as comprehensive as entire sensory systems, can be activated or inhibited. In cells within the systems being activated, the possibility of learning is enhanced, whereas in cells within inhibited systems it is suppressed. In general, these effects are short-lived due to the fact that the focus of attention is constantly moving from one area of the brain to another. If, however, attention is focused more regularly on some brain areas, which means that the information processed by these areas is processed preferentially to other information, then learning in these locations and locations associated with them will occur at a higher rate than in areas not selected by attention. This would be equally true whether the effect of attention is to enhance the activity of the selected area, or to inhibit the surround. The net effect of either scenario would be a relatively higher activity in the selected areas. Now, if a learner consistently arrives at a successful outcome by selecting a certain information channel, the visual system or the auditory system, for example, it is highly likely that allocation of attention to this channel will develop into a persistent strategy unless the circumstances which render this strategy efficacious change. A persistent strategy of this sort would by no means preclude the allocation of attention to other channels. It would simply draw attention to certain brain sites more often than to others. A similar situation could arise, of course, as a result of genetic predispositions in the individual learner, especially if these are encouraged by environmental circumstances. At these persistently selected brain sites, however, learning should carry on at a higher rate than at those selected more rarely, all other intensive effects such as emotional state and arousal being equal, and if, as has been suggested by some studies on attention, the selection of one channel involves the inhibition of other channels, learning in these other channels should go on at a lower than "normal" (for the person) rate in these systems while the strategy persists. Even if, on the other hand, there is no active suppression of activity outside of selected areas, learning in these non-selected areas should be slower than that in selected ones. During the time that this global selection is in action, however, local selective attention will still be carrying on within systems such that learning will be enhanced in these "pockets" of activity whether the area is globally selected or no.

So far, all the reasoning put forward on learning in this section has its roots in biological research rather than machine models. But as has been discussed in depth in the previous chapter, machine models have turned up some interesting insights into how neural networks learn. It would seem useful, therefore, to return to these concepts to see how they might dovetail with and perhaps extend the picture of learning elaborated to date. To recap briefly what has been discussed about learning in these machine models, then, learning in connectionist systems occurs in two basic modes: supervised and unsupervised. In supervised modes, an external assessment of the performance or weight setting job of the model is made and the results fed back into the system. Unsupervised modes lack this external evaluation and depend instead on the ability of the system to maximize an implicit objective function particular to its structure. Within both these modes, nevertheless, learning can be monitored, which means that it can have an internal assessment of error. It can be seen, then, that there are four possible modes, and since within these the possible combinations of architecture, dynamics and parameter adjusting procedure are many, the number of potential model types is high. Although these learning modes are commonly discussed as separate and distinct from one another, systems with combinations of the two parent modes are often constructed, the unsupervised learning unit serving as the front-end of the model. The function of the latter is to sort information into basic components before supplying them to a supervised learning unit which then picks out the combinations of these components that are most important to the task at hand. It seems more appropriate, therefore, to conceptualize learning mode in these models, and even more so in biological systems, as varying along a continuum between fully supervised and fully unsupervised.

In the previous chapter, the extent to which attentional mechanisms could be said to be involved in learning in these machines was discussed, and it was decided that the degree of "attentional focus" that could be ascribed to them varied according to whether the learning algorithm was supervised or unsupervised, supervised learning showing a higher level of attentional focus than unsupervised. If connectionist models do offer real insight into learning in biological systems, it should be possible to reverse the perspective of that question to ask to what extent the supervised/unsupervised learning continuum is reflected in attention directed learning in, for example, humans. It was also noted that attention can be focused or divided (less focused), and presumably it can vary between the two extremes of these states such that there is a continuum of attentional states between fully focused and fully unfocused. So, since attentional focus may be driven voluntarily, that is, by expectation and intention to act, or by the stimuli impinging on the system(s), and since expectation and intention to act are highly experience dependent, it can be said that the degree to which attentional focus is driven by expectation and intention equates roughly to the degree to which there is knowledge external to the system guiding its learning, and the degree to which it is stimulus driven can be equated roughly to the degree to which external knowledge is absent[4]. Thus, learning carried out under voluntary attention conditions can be considered "supervised" learning, while learning acquired under stimulus driven attention may be termed "unsupervised" in the connectionist senses of these words.

Pushing this analogy a little farther, in a voluntary attention learning situation, the "teacher" or evaluation mechanism which determines how the synaptic weights will change is the experience base of the learner herself. That is, the teacher resides inside the head of the learner. Reiterating the point made above in discussing Damasio's insights into attention and learning, then, to say that the evaluation mechanism which guides further learning is the experience base of the learner is to suggest that personal history involving cultural and social milieu as well as

---

[4] To say this is not to contradict previous arguments with respect to the involvement of experience in stimulus driven attentional allocation, but merely to emphasize the difference in the magnitude of the contributions of experience in stimulus driven and voluntary attentional allocation.

characteristics of past learning situations are in constant interaction with what is currently being learned. In fact, it would be interesting, if difficult, to assess to what extent attention may be directed by external teaching sources, especially if these directions go against those of the "inner teacher." Presumably, if this inner teacher has a strong voice (in other words, if the learner has lots of experience with the task), it will have a commensurately strong activating effect on the processing of the information normally associated with the task such that the latter will capture attention promptly and be processed very quickly and without the need for consciousness, or, to use the term common in attention research circles, automatically. If the inner teacher is not well informed, however, many sources of information may be approximately equally activated and, as a result, attention will not be strongly drawn in any one direction, and should moreover be largely stimulus directed. Thus the processing of information would be slower in all instances and more susceptible to voluntary control either through strategic decision-making or through directions given by an "external" advisor such as a teacher or a computer help program. This external direction, however, would still be interactive with the internal directive forces of the learner. Because the information is unfamiliar, it would also be passing through a limited capacity system, so that too much new information might cause an overload. This scenario is in keeping with much of attention research on these matters and particularly with the research of Raichle et al. (1994) which found that practised tasks are processed not only more quickly, but by a different route in the brain than novel/naive ones. It seems reasonable to assume that the difference in routes taken by these two categories of tasks is mandated by the need for thought, whether conscious or not, with respect to which information is most relevant to the situation.

In order to clarify these concepts, consider what might happen if a person is taken out of their normal social milieu, and put into one in which the terminology and manners are unfamiliar, though the language and culture are the same. With respect to what to say and how to behave, because expectation and intention to act are ill-informed, information processing would be slow and laborious (in terms of resources available), such that the behaviour and speech of the person would probably show uncertainty due to the slowness of processing of information, and/or inappropriateness as a result of false priming of information for processing. Because the language and culture are the same, however, much information would still be processed by the fast, capacity unlimited route, making it unlikely that a major change in information gathering or attentional strategy would result, though the unfamiliarity with the terminology (causing difficulty in understanding what is being said as well as in what to say) and manners (a visually observable behaviour for the most part) might bring about a shift in attention away from the auditory system and towards the visual system. That is, the visual system might be allocated attention more often than usual.

Consider, now, a situation in which a person is unfamiliar with the language, whatever the social milieu. In such a circumstance, auditory speech input would provide a minimum of useful information even in the case of modifiers of semantic content such as intonation and pitch. Thus, in spite of the fact that there would be some familiar speech sounds within the foreign language, expectation and intention to act (as far as speech actions are concerned) would be entirely uninformed on some input, or inappropriately informed by the native tongue. Attention would be drawn to all linguistic stimuli according to past experience, and decisions as to which sounds are more important and should be allocated attention would be made on an information base which is not applicable to the new situation. Hence, the outcome of the attentional allocation based on external stimuli (world driven) would not be highly informative. Furthermore, the quantity of unfamiliar auditory data would almost certainly overload the novel/naive information processing route. In short, little or no information which could be said to be useful in decision making with respect to acting on the world (the main purpose of the attentional system according to many attention researchers) would be forthcoming from the auditory system. The most likely result of this problem would be a strategic (voluntary) allocation of attention to systems which *would* afford information that would lead to good decision making. Moreover, if expectation and intention to act are well informed for these other systems, attention may be automatically drawn to them as a result of the sheer speed of processing. In other words, focused attention may be "captured" by these familiar stimuli. The results of such an attention capture could be many and varied, and would depend on the idiosyncrasies of the individuals involved: more attention might be allocated to all other systems in the brain instigating a state of high alertness in these systems; some systems might be allocated attention preferentially due to the high relevance of the information they provide; one system might receive almost all of the attention. Whatever the details of the change, one thing that seems certain is that the auditory system would *not* be allocated its normal share of attention while attention is focused elsewhere. It would therefore suffer from a reduction in neuronal activity as a result, and learning would not be going on at a normal rate within this system[5]. This is not to say that *no* learning will be going on through the auditory system. As has been discussed, it is unlikely that there is ever a time when learning does not proceed in the brain at some rate. It is to say, however, that the rate of learning will be reduced in this system while its neuronal activity is reduced. On the other hand, learning in the systems being allocated an increased amount of attention will be accelerated. The consequences of these matters will be dealt with in more detail

[5] This statement should hold true whether there is suppression of activity in the auditory system or not, since the latter would not be receiving its normal share of activation due to attentional allocation.

below.

The point of these examples is to demonstrate that where attention is allocated is determined by factors such as the experience base of the person and the situation at hand, and that these factors interact with information processing systems within the brain to try to make sense out of a particular circumstance, or complete a given task. Thus, purely "voluntary" (in the sense of being dictated by internal directions or choice whether externally initiated or not) control of attention is rare. As has been noted previously, attention is always subject to "capture" by stimuli which already have a high relevance weighting for a person. The interest of this dissertation, of course, is to try to figure out what happens within the information processing systems of the brain, and with it attention allocation, in a situation of immersion in a second language. Hence, a more detailed exploration of the dynamics of language learning will be taken up in the following section.

Before turning to language learning, however, a summary of the salient points of this particular *connectionist general model of learning* would seem in order[6]:

---

[6] It should be noted that the theorizing on learning put forward here represents a piecing together after connectionist principles of information on learning and attention by the author, and in no way can be taken to represent a definitive connectionist model of learning.

1. Learning is manifested by changes in the weights of synaptic connections between neurons, and is a product of neuronal activity or firing rate. Under normal circumstances, the higher the activity, the more likely it is that changes in synaptic strength will occur.

2. Neurons are extensively connected to one another in networks. These networks form systems of stored knowledge on scales ranging from the extremely micro to the extremely macro (in brain terms) after adequate exposure to classes of input data. The output of a particular system is input for systems further down the information processing system.

3. Emotional state, whether positive or negative, is interactive with the learning process. It appears to be highly influential in directing selective attention, and is thought to act by assigning value to stimuli according to past experience. This past experience is dependent on social and cultural milieux, as well as predispositions innate to the learner.

4. Selective attention can be stimulus driven or voluntarily driven at both the local and global levels. When voluntary, attention allocation is guided by expectation and intention which are sculpted to a high degree by experience. Stimulus driven attention also appears subject to the influence of experience, however. Thus, what is learned in the present must be guided to some extent by

what was learned in the past. Attentional allocation itself may be trainable with experience.

5. The action of selective attention at both the local and the global levels is to enhance and/or inhibit the activity of neurons. Where the cellular activity is enhanced, learning should carry on at a faster than normal rate; where it is not enhanced by normal attentional allocation or inhibited, the learning rate should be slower than normal. If a learner develops an attentional allocation strategy such that certain systems in the brain are persistently selected more often than others, then learning through these systems should occur more readily than through systems not so selected.

6. Learning in connectionist systems occurs in two basic modes, supervised and unsupervised, depending on whether an external assessment of the output, or "teacher", is available to guide learning or not. Complex learning systems, both biological and machine, generally use combinations of these two modes, however, so that learning may be more usefully conceptualized as occurring along a continuum of modes between fully unsupervised and fully supervised. Highly supervised modes can be considered to have a greater degree of attentional focus than less supervised or highly unsupervised modes.

7. Since in voluntary attention attentional focus is driven by

expectation and intention, and since the latter can be considered to provide an evaluation of the performance of the system, learning carried out under predominantly voluntary attention allocation may be thought of as largely supervised learning, while stimulus driven attention may be considered largely unsupervised. In the absence of external directions, the "teacher" which determines how the synaptic weights will change is the experience base of the learner. Thus, the teacher resides inside the head of the learner. If the inner teacher has lots of experience, attention will be allocated promptly and information processing will be automatic; if the inner teacher is not well informed, attention will not be strongly focused and information processing will be controlled. In the case that direction external to the learner is available (for example, a real teacher), it will always be interactive with the predispositions of the experience base of the learner.

3.2.2 Development of a speech-based connectionist model of language learning

Some general considerations

In this discussion of the learning of a first language, speech learning is given primacy since most first languages are learned as speech, and also because the model of second language learning towards which this thesis is progressing is one of second language learning in "immersion" situations in which the learner is surrounded by the speech sounds of the second language. It is also assumed that speech comprehension must be mastered before speech production proper, although it is recognized that the two processes will be interactive once speech production has begun, and even as early as the babbling stage. Hence, the processes involved in speech comprehension will be emphasized in the following discussion. The role of reading in language learning will be dealt with only insofar as it impacts on educational issues surrounding second language learning. Because the focus of the discussion is on speech learning, then, only the function of the auditory system will be described in any detail, all other sensory systems being taken to be subsidiary, though complementary, to the auditory system in speech learning.

Architecture of the auditory system

When an auditory signal impinges on the hair cells on the basilar membrane of the cochlea (the inner ear), it causes them to send a signal through the auditory nerve fibres to the cochlear nucleus corresponding to each ear. From there, the signal is transmitted binaurally (that is, with signals from both ears present) through several brain centres (the superior olive, the inferior colliculus, and the medial geniculate body being the most important) to the primary auditory cortex (A1). There are two such routes, one on each side of the head, which operate in parallel in auditory processing (see Churchland, 1986, p. 109; or Kosslyn & Koenig, 1992, p. 216, for diagrams of auditory system connections). Thus, there are two primary auditory cortical areas, one in each of the superior temporal lobes. As is the case for primary visual cortex (V1) and the retina, A1 preserves the organization of the basilar membrane, the major sensory receptor area of the ear, which is linearly arranged. In fact, both the primary auditory cortices and the inferior colliculus are *tonotopically* organized; in other words, they preserve the frequency sensitivity pattern of the hair cells. Another way in which the auditory system is similar to the visual system is that it has extensive feedback connections from higher to lower centres, allowing information to be passed as far back through the system as the hair cells themselves. "Although the function of the recurrent connections is not fully understood, they are thought to be important for regulating selective attention to particular sounds" (Kandel & Schwartz, 1985, p. 408). It has been shown that some critical functions, such as the localization of sounds, are carried out by the intermediate centres of the auditory system. However, the system as a whole works to identify incoming signals which may be of importance to the organism:

> The auditory system is equipped with a mechanical device, the cochlea, that performs a Fourier analysis on incoming sounds, breaking them down into simpler frequency components. The hair cells of the cochlea are tuned both mechanically and electrically, properties that enable them to transduce particular frequencies of sound into activity in auditory nerve fibers. The auditory pathways of the brain provide us with the ability to detect sounds of different pitch,

loudness, and points of origin in space. We also have the ability to attend selectively to sounds of particular interest. Thus, the auditory system provides us with all the sensory capacities necessary to analyze complex sounds such as human speech. (Kandel & Schwartz, 1985, p. 408)

Since it is in its ability to analyze and selectively attend to speech that the auditory system is of interest here, notions of attention allocation and the formation of informed subsystems within the auditory system will be given priority in the following discussion.

Dynamics of speech processing

In recent years, Stephen Kosslyn and Olivier Koenig have described (and modelled) in great detail how language is processed in the brain. Since these researchers are connectionists, and since their views are highly regarded among connectionists themselves, it seems appropriate to examine their vision of speech learning, and then to relate it to the more general picture of learning outlined above. According to Kosslyn and Koenig (1992), then, the primary auditory cortex acts as an auditory buffer which "stores relatively unprocessed acoustic cues. These cues organize the amplitudes of different frequencies at different points in time, representing changes in loudness, pitch, pauses, patterns of rising and falling amplitudes, and the like" (p. 215). Because the auditory buffer can hold more information than can be processed in detail, however, a selection of the most important information to process further is necessary. This procedure is carried out by selective attention according to the criteria of saliency detailed in chapter 2, many of which, it may be remembered, are dependent on experience. The structure upon which auditory attention allocation depends is the inferior colliculus, and it is of interest that some of the neurons in the inferior colliculus project into the superior colliculus to cells that are involved with motor production. Since the superior colliculus is a visual reflex centre, this allows auditory information to produce a visual fixation of the source of the sounds, "[h]ence, one typically pays attention to a single object, registering its appearance and sounds at the same time" (ibid., p. 217). Indeed, it has recently been suggested by Driver and Spence (1994) that there are spatial synergies between auditory and visual attention rendering it difficult to attend to visual and auditory information simultaneously presented in different locations. The information in the auditory buffer passes through an auditory preprocessing subsystem that extracts stable patterns which underlie combinations of such cues. In the opinion of Kosslyn and Koenig (1992), this system is not devoted to language per se (pp. 217-219), a point which will be discussed in more detail below. The subsystem becomes "tuned" by experience such that stimulus properties that distinguish between words are encoded downstream in the system, and feedback connections reinforce the encoding of those properties in the auditory buffer thereby enhancing their perception. As Kosslyn and Koenig note:

This idea allows us to understand why it is so much easier to distinguish the sounds speakers make when they are speaking a familiar language. As an aside, an intuitive understanding of this effect may be the reason why people tend to talk loudly to someone who does not speak their language very well; the speaker knows that the person will have trouble encoding the sounds, and unconsciously tries to give him or her every opportunity to hear them. (p. 220)

More formal corroboration of the difficulties encountered by second language learners in encoding the speech sounds of the target language can be found in a study by Danhauer, Crawford, and Edgerton (1984). They discovered that predominantly Spanish speaking subjects (those with only a minimal knowledge of English) performed significantly more poorly on an English language based nonsense syllable test than either English monolinguals or English/Spanish bilinguals, making them susceptible to being judged hard of hearing.

It has been demonstrated that even in early infancy, humans can distinguish phonemes as categories, hence it would appear that experience is not necessary for their perception (Kosslyn & Koenig, 1992, p. 221). What experience with a first language does, therefore, is to enhance the perception of phonemes and complex groupings of phonemes which are of value in the learner's environment. If a particular phoneme is not used in the language a child is exposed to, it later becomes difficult to perceive, and this process of the weeding out of valueless phonemes takes place within the first year of life (Werker & Tees, 1984; Werker, 1991). Looking at this issue of value from an interesting perspective, Kosslyn and Koenig (1992) comment that "[a]pparently, our auditory systems develop certain *insensitivities* (which lead us not to make discriminations; i.e., to group within a category) because this is useful" (p. 222). A possible mechanism for this type of suppression of certain synaptic connections may be long-term depression (LTD). Churchland and Sejnowski (1992) speculate that homosynaptic LTD may "have a role in culling out low-grade or 'don't care' information" (p. 291). Jenkins (1991) gives a brief and coherent overview of what is happening in a child's brain during this early period of exposure to a first language[7]:

> The baby does not have to "break in" to the acoustic stream in some magical fashion to get into speech; the child is already sensitive to the phonetic universals. Then, very rapidly in the first year of life, the sensitivities are selectively enhanced or suppressed to yield a pattern that selects the speech sounds that the child hears in the community. Either

---

[7] The fact that the culling out of sensitivities to phonetic universals is occurring in the first year of life is important in that this is a period of very high plasticity in the brain, and whatever sounds are given a high importance rating during this time will probably be designated as such for the lifetime of the child. Nevertheless, it is evident that the development of insensitivities is by no absolute since it is possible to learn other languages, even those using very different sets of phonemes, into adulthood.

> peaks are forming, or valleys are being pushed down, to use Pisoni's metaphor. There is, overall, a sharpening of distinctions that count, and a grouping together or collapsing of the sounds that are not important in the child's environment. (p. 436)

Whatever the mechanism, then, as a result of linguistic experience, certain combinations of sounds are stored as valuable while others are not, and when activated auditorily, these stored representations can subsequently be perceived preferentially and understood through the semantic associations that have been simultaneously picked up during their encoding and storage. Kosslyn and Koenig (1992) posit that if a word is familiar, it will trigger activation of the largest possible unit stored, that is the word itself. If the word is unfamiliar, however, it will be matched to its component sounds in a search for recognition. In the case of an unfamiliar language, even intermediate groupings, not to mention certain phonemes, may not be available, and the preprocessing subsystem will be forced to learn new groupings or to encode previously suppressed phonemes (pp.223-226). Since second language learners would have difficulty in recognizing the speech sounds of the second language, it can be said that, during the learning period, they have temporary phonological deficits. This would explain the results mentioned earlier of Danhauer, Crawford, and Edgerton (1984). In spite of the fact that speech processing subsystems are specialized to process speech sounds, Kosslyn and

Koenig consider that "there is sufficient continuity between speech and nonspeech sounds for the same subsystems to encode both" (ibid., p. 226), an opinion which is supported by speech perception research (see Darwin, 1991; Lindblom, 1991). This is an important point, since it implies that what happens to the speech processing system may affect to some degree the auditory processing system in general.

As was touched on briefly above, however, the auditory processing system is not the only sensory system implicated in speech learning. It has been shown that the vibratory and tactual information of the neck and face during articulation, as well as the visual information of the face and mouth, can communicate language (see Jenkins, 1991). Visual information appears to be particularly important to speech perception. For instance, Danhauer and his group (1984) found that there was an improvement in score on the nonsense syllable test for all groups when the test was administered in a face-to-face presentation rather than a controlled auditory only presentation. Significantly for this dissertation, this was particularly so for the predominantly Spanish speaking group; whereas the improvement in scores for the monolingual English and bilingual Spanish/English speakers was of the order of 6%, the improvement for the Spanish speakers was approximately 14%. It will be argued later (as was intimated in the last section) that second language learners in the early stages of language learning become more dependent on the visual system for processing information, including speech information, than fluent speakers of the target language.

In other cross-language studies on adults, it was discovered that visual information not only affects speech perception, but can actually override auditory information in situations in which the visual and auditory information do not correspond appropriately. Furthermore, the alteration of auditory perception in second language learners resulting from such uncorrelated visual information is influenced strongly by first language prototypes (see Werker, 1991). The weight given to visual information in speech perception should not be surprising in light of the above-mentioned reflexive mechanisms in force correlating visual and auditory attention to a sound source. Once a sound is well learned, therefore, it is to be expected on the basis of neural connections that either auditory input alone, or visual input alone (i.e. a degraded informational input) might be sufficient to complete the pattern and reproduce what I will call a "speech" image, due to the fact that there is considerable debate in speech research circles about what is actually perceived in the process of speech perception[8]. The main point being made here, is that, under normal

[8] In the opinion of many speech researchers, what are perceived are "phonetic gestures" which have been defined by Fowler and Rosenblum (1991) as "organized movements of one or more vocal-tract structures that realize phonetic dimensions of an utterance" (p. 34). However, in the opinion of Liberman and Mattingly (1985), these gestures "must be seen, not as peripheral movements, but as the more remote structures that control the movements. These structures correspond to the speaker's intentions" (p. 23).

circumstances, the auditory and visual systems work in concert in locating and perceiving sounds, including speech sounds[9], and, given this fact, it seems reasonable to propose that, when auditory information is deficient in some way (low intensity, or garbled, or even not previously encountered and stored), corroborating visual information would be sought to boost the signal leading to perception and identification, and vice versa in the case of poor visual information. Moreover, in every circumstance, the stored representation which is activated will be strongly dependent on the balance of informational input from the two systems. Fowler and Rosenblum (1991) have commented in this regard:

> Subjects' identifications of syllables presented in this type of experiment [a video display of a face mouthing sounds synchronized with an acoustic signal of a speaker] reflect an integration of information from the optical and acoustic sources. Furthermore, as Liberman ... pointed out, the integration affects what listeners experience *hearing* to the extent that they cannot tell what contribution to their perceptual experience is made by the acoustic signal and what by the video display. (p. 37)

Indeed, Driver and Spence (1994) report that, in an experiment of

---

[9] With respect to speech, Fowler and Rosenblum (1991) note that "both sources of information, the optical and the acoustic, apparently provide information about the same event of talking, and they do so by providing information about the talkers' phonetic gestures" (p. 37).

the type mentioned above, there was "a powerful effect of Speaking-versus Chewing-lips" (p. 318) in a word recognition test, and that the benefit was evident even when the auditory and visual speech information were coming from different locations. Given the interdependence of the visual and auditory systems in the processing of speech, it becomes comprehensible that, under certain circumstances, especially those in which the information from the visual system registers strongly with a well learned stored representation of a speech sound while that from the auditory system does not, visual information might override auditory in speech perception as Werker and her colleagues discovered. This would be particularly likely to occur if there was a predisposition towards allocating attention to the visual system in the person attempting to perceive the sound.

Integration of attention and language learning model

The description of speech learning just detailed gives some notion of what changes and processes are occurring biologically in a language learner during the early stages of learning of a first language. It is useful, however, to couch these concepts in terms of the general model of learning already outlined, that is to say, in terms of where attention is going and whether the learning is supervised or unsupervised. Given that, in the very early stages, a basic sorting of sounds into meaningful "chunks" is taking place, a job for which unsupervised learning with its ability to seek out "systematicities" in the input data would be eminently suited, it seems reasonable to suggest that, during this period of language learning, the predominant learning mode would be the unsupervised one, though, of course, supervised learning would always be active to some extent in the form of genetically encoded behavioural predispositions. Focused or directed attention to language based on expectation and intention to act, then, would be only weakly active in these early stages[10], but would gradually build up as the experience base of the learner becomes better informed, and would be observed to be active once unused phonemes are no longer properly heard by the learner. As was mentioned above, attentional allocation during this time would be directed largely by genetic predispositions more concerned with survival than language. Thus, the language learning mode would pass gradually from being generally unsupervised (or supervised only by genetically encoded biases) to generally supervised as experience is gained and systems of knowledge develop. The more knowledge gained, the more supervised the learning becomes, and therefore, the more incoming information is subjected to screening. Aiding and abetting this process, would be information from the external environment which would provide not only the necessary linguistic input, but a more global form of evaluation of prelinguistic actions in terms of emotional and linguistic feedback: active teaching from caregivers, successful and unsuccessful attempts to mimic speech sounds, and

[10] Research has shown that as early as a few days after birth, infants will show a preference for their native language, indicating that language experience begins in utero (see Werker, 1991, p. 103).

environmental reactions to these could all be considered types of supervised learning in which attention is being guided more directly to favour some sounds over others. While the processing of language is taking place, of course, information from the senses which relates to the linguistic information, that is, general auditory, visual, tactual, taste, olfactory, and proprioceptive information, to mention only the better known of the sensory modalities (see Ganong, 1991, p. 107 for a more complete list), will be stored in tight association with the linguistic information and will be valuable in both the development of meaning, and the concomitant evaluation of the importance of the latter for the future guidance of attention.

The result of all of this is eventually a highly structured language system in the brain, one which continues to be refined throughout life according to how well one learns the language, whether one learns two or many languages, and by what means. Fortunately, neurological research has reached a stage of development such that information on the organization of language in the brain is now available. Hence, it is to a brief survey of this information that we turn next.

The microarchitecture of language in the brain

The processing of language input produces a highly organized system of knowledge in the brain made up of specialist subsystems or modules at both cortical and subcortical levels. Indeed, as has been mentioned above, the speech processing system itself would appear to be a module within the auditory processing system. In a recent review of data from various experimental sources investigating language organization in the brain, George Ojemann (1991) has concluded that, in light of this new data, some revision should be made to the traditional model of language organization in which "language is processed serially from decoding in the posterior temporal cortex (Wernicke's area) to motor expression in the posterior inferior frontal lobe (Broca's area)" (p. 2281). He goes on to comment (ibid.) that

> While the newer data agree with the older on the importance of the perisylvian cortex of the left hemisphere (of most individuals) for the generation of language, they suggest that models of cortical language organization should include separate systems for different language functions. Each system includes localized frontal and temporoparietal areas as well as neurons widely dispersed elsewhere in the cortex, with the entire system being activated in parallel. In many respects, these new findings related to language in man are similar to those described for organization of other functions in the primate association cortex ... However, two aspects of cortical language organization differ substantially from other findings made in the primate cortex: the long-known lateralization of language to one hemisphere (designated as "dominant"), and recent findings of a substantial variance in

the individual pattern of organization within that hemisphere.

Several aspects of Ojemann's conclusion support connectionist conceptions of language learning and organization: first, the compartmentalization of language learning into separate systems of specialized knowledge, for example, naming in general, naming different categories of objects, reading in general, syntactic aspects of language, all of which have been found to be subserved by separate systems to some degree, though it is assumed that there is some manner of interaction between them which is as yet undetermined; second, the parallel activation of all areas in a system rather than serial processing; and third, the finding that there is substantial variance in the organization of language from individual to individual, a feature which may speak to the importance of experience, though it may equally well be dependent on genetic predisposition. The two possibilities are not mutually exclusive, however, and it seems reasonable to assume that, as is suggested by Damasio's work, they interact. That experience can have a major impact on brain organization can be seen from data demonstrating that, after major injury to the left hemisphere of the brain in infancy, language will develop in the right hemisphere, and that language deficits may change over time after a patient has suffered a particular brain lesion (Ojemann, 1991, pp. 2283-2284).

Another finding that is of interest to connectionist theory is that, with respect to naming, the total cortical surface area devoted to this specialization was smaller in subjects with a higher verbal IQ than in those with a lower verbal IQ. Thus, "[t]here is a relation between organization of the language cortex and language abilities" (ibid., p. 2284). Speculating on the role of experience in accounting for individual variances in organization, Ojemann notes:

> Although these findings do not indicate a topographic effect of linguistic experience, patterns of localization of two languages in the same subject do suggest some such effect. Essential areas for naming in a later-acquired language in which the subject was less fluent were larger than those for naming in an earlier acquired, more fluent language. (p. 2285)

Thus it would appear that, as is the case for the sensory systems, related speech information is stored in close proximity, and that the proximity becomes greater the more strongly related the information. The findings on the organization of language in bilinguals being of particular interest to this thesis, explaining this lower level of organization of a less fluent second language is important, as is explaining how it is that, in polyglots, different languages form discrete systems to the extent that, as a result of brain lesions, "only one of their multiple languages [may be left] intact, one that may not be the mother tongue, the language most frequently used by the subject, or the language of the subject's environment" (ibid., p. 2282). In the next section on second language learning, then, an attempt will be made to elucidate mechanisms by which such separation of function might come about.

3.2.3 Learning a second language in immersion situations: Towards a connectionist hypothesis of the cognitive consequences of bilingualism

Since second language learning can take place in a myriad of different venues, each having its own particular set of contextual factors and thus spawning a particular learning route, a choice must be made as to which second language learning circumstance to examine here. In this dissertation, only second language learning in immersion situations will be considered. There are three major reasons for this specificity of focus: to begin with, the Threshold Hypothesis of Jim Cummins, which provided the initial impetus for this thesis, was based largely on research in language immersion situations; secondly, the dispute as to the cognitive consequences of bilingualism centres around second language learning in immersion situations - no one suggests that cognitive deficits will result from second language learning in non-immersion situations; thirdly, the term "immersion situations" constitutes a class of learning situations in its own right, given that there are early and late immersion situations, and what have come to be known as "submersion" situations, situations in which learners are surrounded by the second language in both the school and the dominant culture of the society and in which the first language of the learners is neither used nor promoted in the school, hence to develop a model of second language learning for this one group of situations will be a difficult enough task in itself. The salient characteristics of an immersion situation, then, are that the learner (a) has no or very little knowledge of the target language, (b) is surrounded auditorily by the language, (c) receives little direct instruction in the language which is expected to be learned through content, and (d) is expected to function in a "normal"[11] fashion in the learning environment.

According to the model of language learning given above and the information from Ojemann's review, the task which the brain must undertake in learning a second language in order to function in a society is the development of a second language system or module, independent of the first language system. A major question which arises here, is whether or not this would take place via the first language system or entirely independently. It will be argued that the most likely scenario would be that the development of the second language should occur, at least in the initial stages, through the first. No matter how the development of the second language occurs, the fact remains that the auditory specificity of the speech module for the first language must be broken down, or at

[11] By "normal" fashion, I mean that he or she will be expected to carry out tasks and thereby solve problems dictated by someone or some circumstances in the environment.

least opened out, in order for the speech sounds of the second language to be processed and perceived adequately. And since fluent speech comprehension generally precedes fluent production, this aspect of second language learning will be dealt with first.

As was discussed previously, the first language speech system functions such that the processing of certain phonemes, sound combinations, patterns of intonation, pitch, and pauses is selectively enhanced by attentional mechanisms if this auditory input is considered of value to speech, while other auditory input not thought to be important is selectively suppressed. Given this system of organization, a learner immersed in an unfamiliar second language will not only be unable to recognize the unfamiliar speech sounds but will be unable to hear them properly even though his auditory system is functioning perfectly adequately. In order for this situation to be corrected, the attention system operating on the speech module must be recalibrated to value phonemes which were previously devalued, as well as sound combinations and patterns of speech sounds which have not been encountered in the past. Attention may be initially drawn to the new sounds because of their very novelty, but novelty is decided upon in reference to background input. When everything, or nearly everything, is novel, then, novelty (except for novelty or abruptness of onset) ceases to be a good criterion for allocating attention. It must be remembered that the purpose of allocating attention is to impose temporal order on the processing of input, and that this ordering is decided upon in reference to the importance of the input signal. New input has only its novelty to give it importance, and if all input is novel, there is no basis upon which relative importance can be judged. Another aspect of information processing that must be taken into consideration is that all novel input passes by the high attention demanding, limited capacity, controlled processing route. It is highly likely, therefore, that this route will quickly become overburdened by speech input in a language immersion situation. The question that must be broached now is "Where does attention go in order to make decisions on which sounds are important in the second language input?"

To answer that question, it is necessary to consider just how closely the auditory and visual systems are linked in speech perception. As was noted above, the connections between the inferior and superior colliculus bring about a synchronizing of auditory and visual attention to the source of the speech sound, and furthermore, there is an integration of visual and auditory information in the encoding of speech. It was also pointed out that the contribution or weight of visual and auditory input in controlling output would change according to circumstances, but since each type of input gives equivalent information with respect to "speech" gestures, phonetic gestures according to many (see Mattingly & Studdert-Kennedy, 1991), the one can substitute for the other to some extent at least. An immersion situation in which the acoustic input is largely novel and, therefore, difficult to hear properly is essentially the same as a situation in which extraneous noise degrades the acoustic speech signal. In such a circumstance, then, visual input would become very valuable and may even be weighted more heavily in importance than in normal first language speech perception. Summerfield (1991) has postulated that, in noise, "[t]he fact that these [visual and auditory articulatory] changes are occurring on the same time scale in the two modalities may allow vision to direct auditory attention to the relevant aspects of the acoustical signal and thus help segregate the talker's speech from background sounds" (p. 127). In a recent paper, Niebur and Koch (1994) have proposed a model of selective attention in vision based on just such a temporal correlation of firing in neurons in the visual system. It seems reasonable to suggest as Summerfield does, therefore, that a temporal correlation of neuronal firing in the auditory and visual systems coming from the same location in space would similarly tie their activities together through the attention system. On this basis, then, it also seems reasonable to suggest that, in a second language learning situation, vision would direct the learner's auditory attention to relevant aspects of the speech sounds thereby ensuring that the information from the two modalities is encoded synchronously and consequently associated in the brain. In other words, visual attention to articulatory gestures could aid the overall speech learning process.

The visual information available to an observer of a speaker's face, however, is not restricted to articulatory gestures. Mood, expression, intent, sense of goodwill towards the listener, all of these and more can be gleaned from a face. Normally, of course, similar information can be garnered from the auditory input of speech. But the distinction between the accessibility of the information from the two modalities, is that knowledge of the complex code upon which the visual information is dependent is less culturally specific than that of the auditory information. Though visual gestures divulging information about the person do differ from culture to culture, these differences are nowhere near so diverse nor as complicated as in speech systems. What this means is that the information from the speaker's face gathered by the visual system will be much more familiar (i.e. have more meaning and thus value) to the second language learner than the auditory. Since this is so, much more of it will pass by the practised information processing route, thereby lightening the load on the novel/naive route taken by unfamiliar information and making the information processing task more manageable. Moreover, the familiarity of the visual information will tend to capture attention through speed of processing and draw it away from auditory information, especially given the poor communicative value of the latter. The same arguments with respect to communicative value and familiarity can be made for general visual information relating to a given situation or task. If one is shown what to do, words are often superfluous, and body gestures and "language" in general tend to be more familiar across cultures than language systems. So, for reasons of

a) helping to encode articulatory gestures

b) gathering meaningful information

c) reducing the information processing load

it is postulated that, in a second language learning situation, attention should be drawn to the visual system more often than it would be if the learner's first language is being spoken, and in the beginning stages of learning much more often, for as long as the second language is difficult to understand.

As has been discussed in the section on general learning, this should mean that visual information processing would be stimulated preferentially, while auditory processing would be understimulated and might be suppressed during this low comprehension period. This in turn would predict that learning via the visual system would be enhanced, while learning through the auditory system would be inhibited. Since other aspects of visual information, not just those related to the encoding of speech, would be of value under such circumstances, information related to recognition of objects and semantics for example, it seems likely that several areas within the visual system would be selectively activated. In the case of the auditory system, however, it is possible that only the speech module would be affected directly by withdrawal of attention, and that other auditory information will be processed as usual. Music, for example, could be processed in the normal way due to its familiar content and discreteness from language. But given that most immersion language learning situations take place in a classroom, and that in most classrooms, speech noises are by far the most prevalent, it seems likely that, over time, a systematic strategy of allocating attention to the visual system would be built up because of the higher value of the information obtained through it. This is assuming, of course, that the capacity of the visual system to process information is not already at maximum under normal first language classroom circumstances[12]. In other words, attention itself will have been trained to strategically choose the visual "channel" over the auditory because it is more energetically economic to do so.

But, it may be protested, if this is so, how do learners ever learn the new speech sounds well enough to form a coherent and discrete system as brain research has indicated they do? To answer this question, it is necessary to remind ourselves of the two modes of learning postulated by connectionism: supervised and unsupervised learning. When the attention system selects information for further processing on the basis of familiarity, it is essentially engaging

---

[12] Such a situation seems unlikely, however, since it is the capacity of the attention system that appears to be the limiting factor in processing information, and under normal first language classroom circumstances, attention capacity would be divided between several sensory and other information processing systems. In the case of immersion in a second language, therefore, it would simply be a matter of consolidating resources and directing them towards the visual system.

the supervised learning mode. That is, expectation and intention to act are well informed and prime the processing of certain stimuli over others. Thus, the learning that goes on through the visual system on the basis of familiarity of input would be positioned towards the supervised end of the unsupervised/supervised continuum. In the case of the auditory speech information, however, expectation and intention would be ill-informed in that they would be priming sounds and combinations of sounds expected in the first language, and suppressing many sounds and combinations of sounds in the second language from individual phonemes on up to words. Hence, in this case, the "supervision" is causing problems for learning and a return to unsupervised learning would appear to be mandated so that a new value system can be built up in the speech module. This may require an active withdrawal of voluntary attention from the auditory system in order that attentional selection of information for further processing can become more stimulus driven for a time in order to aid the sorting through by the auditory system of new systematicities in the second language input. Of course, this withdrawal could be accomplished automatically by a switch to the visual system as the primary information processing mode, hence a move of focused attention to the visual system could kill several birds with one stone. As in all unsupervised learning situations, value would once again be built up out of the new systematicities (i.e. features that are persistently present) in the input and, as knowledge of the novel input is acquired, the learning would become more and more supervised, eventually forming novel systems or networks to process the formerly unfamiliar information.

During the process of forming this new system, it is likely that the already established speech system will be shaken up. After all, previously suppressed access to phonemes may have to be reactivated, and many of the phonemes needed for the first language, as well as many of the combinations of sounds, will also be necessary to the second, thus, some sharing of resources seems inevitable. This sharing of resources, however, may mean that new connections must be made, or at least that old connection strengths must be altered, all of which should have some effect on the processing of information by the first language system, though this may be minimized by a subdivision of labour through the use of specialist mininets (see Churchland and Sejnowski, 1992, pp. 125-130). The beauty of this unsupervised/supervised learning system is that both can be going on at the same time: while attention is being allocated to the visual system, whether to encode visually accessible aspects of language learning or just to gather contextual clues to help make sense out of the language spoken, the auditory system can be working away in a low attention, non-conscious manner on discovering the phonological regularities in the input. And, of course, though posited to function separately here, it is expected that knowledge compiled through the two learning modes will cross-correlate both temporally and semantically in the brain, and in the process eventually form a unified and functionally separate speech system.

With respect to the question of whether the second language system is built upon the first language system or develops entirely autonomously, it seems likely that whatever information can be used by the second language will be used quite automatically. For example, semantic overlap of the lexicons would mandate some form of association between words since, in the beginning at least, words can be considered prototype representations of a class of objects or situations, and any new way of naming a class or member of a class would be drawn by its semantic associations towards the already formed "state space" of the familiar word. This attraction would change as the second language word became more tied linguistically to the sounds and structure of the second language, but it seems unlikely that all association between the first and second language representations of a concept could be ruptured due to the semantic links. What this automatic association would appear to suggest is that there should be a progression from a heavy dependence on translation in the beginning stages of learning to a lesser, or even non-existent dependence as fluency is acquired. According to Ojemann's findings on the relationship between verbal IQ and language organization detailed above, the level of fluency would appear to reflect the level of organization, and therefore, the degree to which the linguistic system is specialized for that language. However, an in depth investigation of these issues is beyond the scope of this dissertation.

Since the interest here is in gaining some insight into potential implications for education, it is important to consider what some of the consequences of a shift to a preference for visual processing, both language and non-language, might mean for the second language learner. Before going into these matters, however, it is equally important to remember that, though only immersion situations will be examined, within this category there are several different situations each of which will carry with it its own set of local environmental constraints that will alter the general picture. All of these will not be dealt with. Indeed, the only problem to which the hypothesis of second language learning outlined above will be applied in this thesis is that of explicating the differential results obtained in academic achievement by minority and majority group second language learners in immersion situations. That is, only an alternate explanation of the data which gave rise to the Threshold Hypothesis proposed by Cummins will be sought for reasons which should be evident: this one problem is in and of itself highly complex.

To summarize the argumentation of this section, then:

> It is hypothesized on the basis of a connectionist perspective on how the brain processes information that, during the initial stages of learning in immersion situations, learners of a second language will adopt, whether consciously or subconsciously, a shift in information processing strategy

such that visual processing of information will be favoured over auditory processing, and that this shift will bring about an increased ability to learn via the visual system while it is in effect.

This hypothesis will be termed the *shift in information processing strategy hypothesis*[13]. In the following sections, support for its premises will be sought in the second language learning literature, and it will be examined for its ability to explicate satisfactorily differential academic achievement for minority and majority group second language learners.

---

[13] Although the shift in information processing has been termed a "strategy," it should not be assumed that this shift is the result of a voluntary decision on the part of the learner. The strategy may be better conceptualized as stemming from decisions made by the brain itself and may or may not be accessible to the conscious perception of the learner.

3.3 An Examination of the Explanatory Power of the Shift in Information Processing Strategy Hypothesis

3.3.1 Introduction

In the previous section, it was postulated that, for reasons relating to the multimodal character of language[1] and information processing priorities, students in an immersion second language situation would adopt a strategy of allocating attention more frequently to the visual system in trying to make sense out of their environmental circumstances, both linguistic and non-linguistic, and that, if this attentional shift does indeed occur, learning through the visual system will be optimized at the expense of learning through other sensory systems, most notably the auditory system as it concerns speech processing. In order to lend more credibility to this hypothesis, however, empirical evidence from the second language literature which is consistent with it (as well as any which is inconsistent) should be presented. Before doing that, it is necessary to work out what the consequences of such a shift might be in information processing terms. Both of these tasks, then, will be tackled in this section.

---

[1] Jacobs (1988, p. 312) has noted: "Language depends on linguistic information and, as a *multimodal sensory system*, is not particular about the mode of input it receives; it will make use of what is available. In the absence of one mode (e.g., audition), the other (e.g., vision) will supply the necessary linguistic information. Language is subserved by the senses insofar as it depends on the linguistic information they provide, but it also transcends them because it is not modality specific" (p. 312).

3.3.2 Information processing consequences of an attentional shift towards the visual system

The basic premise upon which a shift to a greater dependence on visual processing has been proposed is that, for second language learners in an immersion environment, the auditory information related to speech does not make sense because it is to a large extent unfamiliar, the degree of unfamiliarity being dependent on the particular first and second languages involved. In other words, the general hypothesis proposes that, during the period of time that the ability of the listener to comprehend second language speech is low, it is to be expected that the dependence on visual processing for information gathering purposes would be higher than normal for that person while the learner is in the immersion environment. As the comprehension level of the learner increases, then, it is equally to be expected that dependence on visual processing would diminish accordingly. Whether it would return to pre-second language learning levels is difficult to predict. There is good reason to believe it would not, however, since a persistent learning strategy with respect to attentional allocation may have developed during this prolonged learning period, and since, due to the enormous plasticity of the brain, new information processing pathways would almost certainly have been formed, or at the very least old pathways would have been altered, to favour a higher degree of visual processing.

A corollary to this hypothesis is, of course, that persistent selection of the visual system for information processing by the brain's attention system would result in higher neuronal activity in visual cortical areas and hence learning would be expected to go on at a faster than normal rate (for the individual) via this system during the period of time that language comprehension is low, and might continue to operate at a rate higher than the previous norm for the individual due to strategic changes in information processing patterns and consequent morphological changes in the brain. Thus, for each immersion student, there should be a correlation between the ability of a learner to process information visually and his or her level of language comprehension. Though it is postulated that a shift to a greater reliance on visual processing should occur for every learner in an immersion situation, it is recognized that the dynamics of such an information processing change would vary from individual to individual. Nevertheless, it is also posited that the overall pattern of change over time would be similar in every case to the pattern shown in figure 3.1 (p.294). The general trend postulated, then, is a gradual increase in the dependency on visual information processing as the learner becomes accustomed to the consistent lack of useful auditory speech information in the second language environment, followed by a gradual decrease in dependency as language comprehension increases, then a levelling off to a steady dependency with the attainment of an adequately high level of comprehension. It is to be understood, however, that this

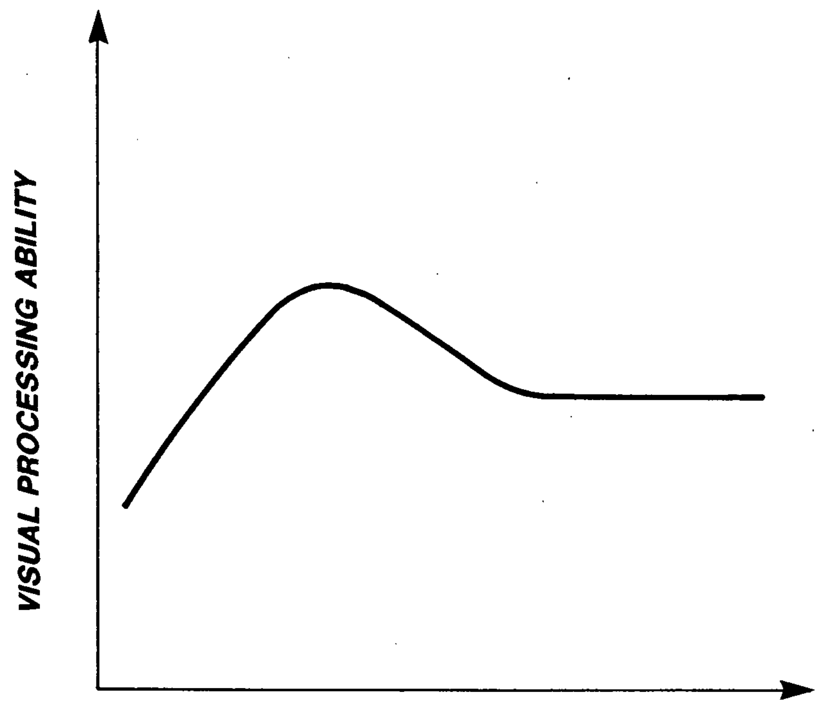

Figure 3.1 – SHIFT IN INFORMATION PROCESSING STRATEGY HYPOTHESIS

"steady" dependency would in fact be a steady average use of visual processing, as, under normal circumstances, the latter would be engaged by the brain in a highly transitory fashion, as would other sensory systems, when appropriate. Figure 3.2 (p. 296) shows the dynamics of the development of the relationship between visual processing ability and level of comprehension of the second language for three hypothetical learners. Before discussing these in greater detail, it should be noted that the visual processing ability axis in this graphic representation is divided into three sections, general deficit, normal, and benefit ranges[2] as might be decided upon by studies based on large numbers of subjects. This has been done in order to facilitate comparisons at a later stage between theories of second language learning and the development of cognitive ability. In every case shown in figure 3.2, it is assumed that the ability of the learner to process information visually is initially (i.e. at the beginning the immersion language learning process) at the mid-point of the normal range. As can be seen, the general shapes of the graphs in all cases are similar to that in figure 3.1.

Case 1 of figure 3.2, then, shows the development path of a learner whose ability to process information visually undergoes a sharp

---

[2] The designations of "deficit" and "benefit" ranges here are intended only to apply to the brain's ability to process information visually, and should not be equated with general cognitive benefits or deficits. It will be argued later that any "cognitive" benefits or deficits that might accrue are entirely task dependent, with all of the interactions that that presupposes between sociological factors and language learning.

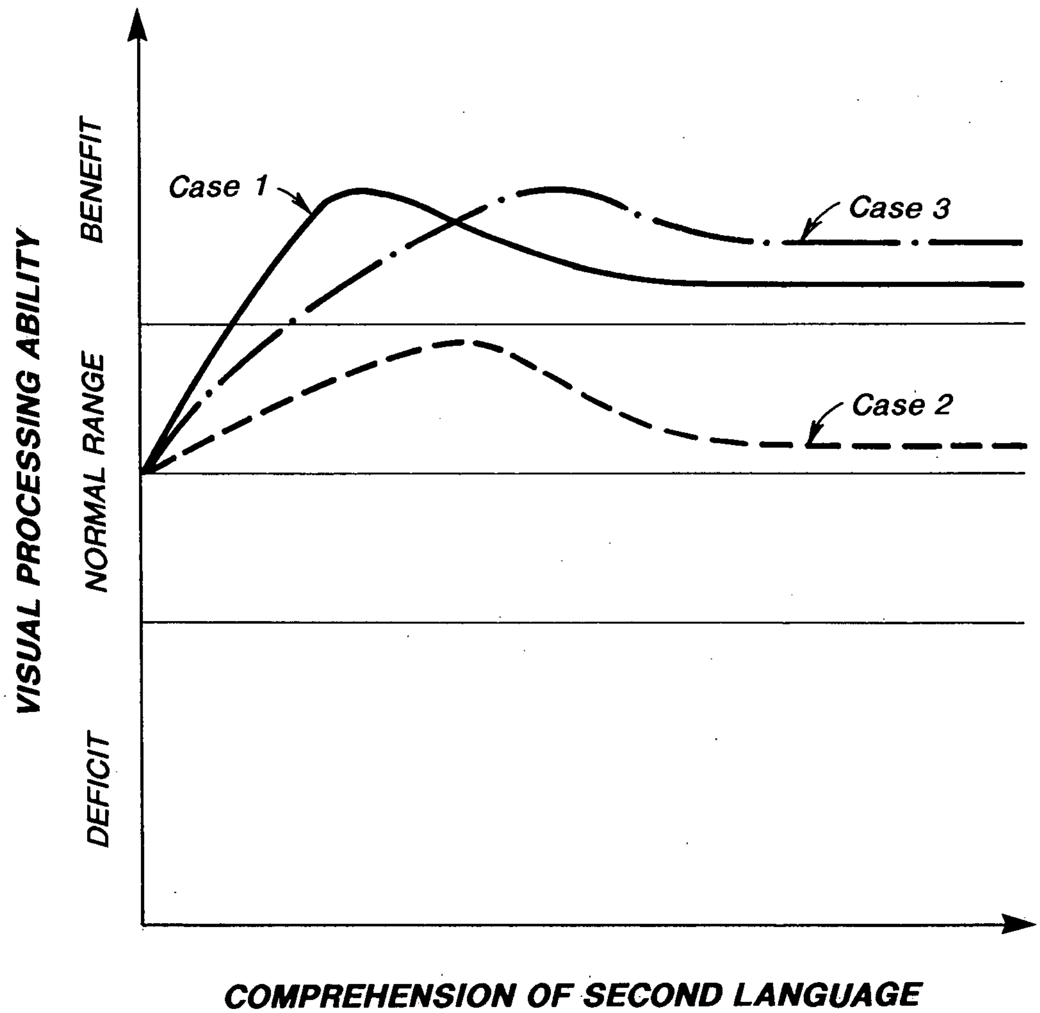

Figure 3.2 – SHIFT IN INFORMATION PROCESSING STRATEGY HYPOTHESIS : THREE SEPARATE CASES

rise, reaches into the general benefit area, and subsequently decreases somewhat before levelling out, though it still remains in the general benefit area for visual processing ability. Case 3 mirrors this progression, but the rise in visual processing ability is less sharp than in case 1. The learner in this case may be slower to take strategic decisions with respect to information processing, or may be more dependent on auditory processing of speech than the learner in case 1 to begin with. In both these cases, then, the learners sustain general benefits in visual processing ability as a result of learning a second language. The course of case 2 follows the same general pattern, but the general benefit area for visual information processing is never attained, and indeed, though the learner has made gains in visual information processing ability, that ability never reaches outside the normal range. Nevertheless, this learner can be said to be in a personal benefit situation with respect to visual information processing ability as second language comprehension increases. Of course, collapsing the graphs of all three cases (and perhaps more of the same type) into one integrated group study would produce a graph such as that in figure 3.1.

Differential pathways for minority and majority group second language learners

It was suggested above that, within a particular group of second language learners, individual developmental pathways would differ according to predispositions towards information processing inherent in the learner[3]. Essentially, these predispositions would act as constraints on the developmental sequence of information processing. In like manner, it is entirely possible, and even highly probable, that information processing pathways between discrete groups would differ according to differing group constraints. That is, given two groups of learners with equivalent ability and educational background, it is to be expected that the pathways carved out during the language learning process would be shaped differently according to experienced environmental differences which interact with the language learning. Needless to say, the greater the environmental differences, the greater the disparity between the two pathways would be. It also goes without saying that if the two groups were not equivalent in ability and educational background the differences would be exacerbated. For simplicity's sake, therefore, let us consider the disparities that might arise between two groups of learners within the same society, where the only difference between them is that one group speaks the high prevalence majority language of that society and seeks to learn a low prevalence minority language, while the other speaks a low prevalence minority language and seeks to learn a high prevalence majority language, both in equivalently structured immersion situations.

---

[3] Notice that there is no suggestion that these "inherent" predispositions are necessarily "innate," though they may be so. Such predispositions towards a particular pattern of information processing may be easily picked up through environmental contact.

Taking the case of the majority group learners first, the low status of the target language would render motivation towards learning low, though this could be aided by parental valuing of the language[4]. Thus, the effect of the status of the language on learning is uncertain in this instance. The effect of the low prevalence of the target language, however, would be to render it unlikely that the learners would have had much exposure to the (minority) second language before beginning their learning, hence the language would be almost completely unfamiliar, if not completely so. This unfamiliarity would tend to promote a very rapid switch to visual processing while learners are in the language immersion situation (i.e. the classroom and, to some extent perhaps, the school in general, though this may not be the case in schools in which immersion classes are mixed in with regular first language classes). Also, once again because of the unfamiliarity with the second language, it is to be expected that the period of intense dependency on visual information processing will be fairly protracted, thus there should be a relatively slow return to a more normal dependency on visual processing (see case 3, figure 3.2, for example). Outside the classroom, however, majority group learners will have ample opportunity to experience and use their first language. During such times, the dependency on visual processing should return to normal levels for the individual

---

[4] In the case of French immersion programs in Canada, for example, parents often accord high value to the learning of the second language for political and/or economic reasons, as was the case in the St. Lambert Experiment discussed in chapter 1.

learner. Hence, the external environment will exert a mitigating effect on the tendency to adopt a general information processing strategy in favour of a heavier dependency on the visual system. For majority language learners in this type of situation, therefore, though the experience of immersion would certainly cause some overall changes in information processing strategy due to the quantity of time spent in circumstances which promote visual processing, it is possible that the value attached to these strategic changes would be associated with, and therefore cause them to be confined to, the classroom.

Minority group second language learners whose first language is of low prevalence and for whom the target language is a majority language of high prevalence within the society can be divided into two major groupings: those who have been brought up within the majority language society (e.g. francophone students in English speaking Canada), and those who have entered the country recently (e.g. Spanish speaking students from South America who emigrate to the United States or Canada). While it is recognized that there are many subgroups experiencing varying environmental circumstances within these groupings, it is considered that some useful generalities which would have information processing consequences can be identified within each. For learners who have been brought up within the second language society, the high prevalence and status of the target language would mean that there would be a high degree of familiarity with certain aspects of the second language even if it is not familiar as a means of communication. At the very least, phonological aspects of the language would be well encoded because of the very prevalence of the language in society. It is highly likely, however, that most of these learners would have some command of the language as a communicative tool before entering the immersion learning situation due to its high status and prevalence[5]. As a result, the hypothesized shift towards visual information processing should occur at a slower pace, and, depending on the level of expertise in the second language, may not occur to any significant extent. Case 2, figure 3.2 may be a reasonable representation of the typical course taken for this kind of group.

As for majority group learners, minority group second language learners who have not been long in the country and have little or no knowledge of the target language should undergo a rapid shift to a heavy dependency on visual processing. In contrast to the majority language group, however, their first language would not be nearly as prevalent in the general society as that of the majority language learner. In fact, because of the low prevalence of their first language and high prevalence and status of the target language in the society at large, these learners would be in a

---

[5] See Bain and Yu (1987) for some insight on these matters. They comment that concerned minority parents often find "that by preschool age the *lingua communis* (the predominant community language, in this case English) has become so dominant that it is but a sentimental fiction to consider the language first spoken as the 'mother tongue'" (p. 221).

similar situation with respect to comprehensibility of language both in and out of the classroom for much of the time. Hence, the tendency to shift to and remain in a largely visual mode of processing would be exacerbated for these students. Due to the quantity of time spent in the visual information processing mode, then, the likelihood that these learners would develop a persistent strategy or bias towards processing information via the visual system is enhanced. Similarly, the likelihood that they would not pay enough attention to auditory processing is also enhanced.

What might such shifts in information processing strategy mean for these learners? To repeat a point that cannot be made often enough, what shifts in information processing strategy of the type discussed would *not* mean is that the learners are in some way, either "cognitively" or in terms of general information processing ability, deficient. As was noted by Clark (1989) in referring to the brain's "kludge-like" character, and Jacobs (1988) in comments with respect to the "multimodality" of language processing, the brain is above all a pragmatic organ which takes whatever information it can get and uses it as efficiently as possible. And, as has been illustrated by language studies, visual information and auditory information (and presumably information from other sensory modalities) can often feed into an identical computational module at the top end of an information processing system or subsystem. The hypothesized shifts should mean that learners with a pronounced modality preference for the visual system as the predominant mode of information processing would be better at accomplishing visually taxing tasks than those without such a bias. It might also mean that these students would be less adept at processing information auditorily, at least as far as language tasks are concerned, than those without such a bias. Furthermore, such students might have difficulties with tasks that are optimally promoted by a particular balance between auditory and visual processing. As has been noted previously, one of these tasks which is especially important to academic achievement is reading. It is crucial, therefore, to examine any evidence from studies of second language learners academic or cognitive performance which might be better explained by the shift in information processing strategy hypothesis and thus lend validity to its premises. If this hypothesis does have empirical support, it could be a useful tool in devising remedial educational treatments. To this end, several sources of pertinent data will be examined, the first being a study reported by Cummins himself.

3.3.3 An explanation of studies on the academic or cognitive consequences of second language learning from the perspective of the shift in information processing strategy hypothesis

(1) An examination of the nature of cognitive tests commonly used for second language learners

In an attempt to understand the reasons why minority students were seen as underachievers and were frequently recommended for assessment, psychological, reading, or speech and hearing, Cummins (1984a, chap. 2) undertook an analysis of the referral and assessment forms for 428 ESL students from a mixture of ethnic backgrounds. 46% of these students had been born in Canada, while 56% came from a variety of countries and socioeconomic situations. His quantitative analysis of the 11 WISC-R subtests given to these students showed unequivocally that they scored much better on the performance (PIQ) than on the verbal (VIQ) subtests:

> Several things emerge clearly from the patterns of WISC-R subtest scores. First, students perform much closer to the average range on Performance as compared to Verbal subtests. There is relatively little variation among Performance subtests ... In contrast to the Performance subtests, there is considerable variance among Verbal subtests. Arithmetic and Digit Span appear to be somewhat less culturally/ linguistically biased against ESL students than the other Verbal subtests. (p. 24)

On close examination, these "performance" subtests prove to be tests highly dependent on visual processing: picture completion, picture arrangement, block design, object assembly, and coding. Hence, if as has been proposed in this dissertation, these students have adopted an information processing strategy which favours the visual system, it is not surprising that they should do better on these subtests. Neither should it be unexpected that these students performed better on the arithmetic and digit span verbal subtests. Numerical symbols are also easily understood visually and would therefore be relatively easily learned and manipulated by these students, and any verbal directions given would require a minimal vocabulary for comprehension[6]. Thus, though Cummins attributes the difference in scores to the fact that some subtests are "culturally/linguistically" biased against ESL students, the information processing demands of the tasks may be an equally important explanatory factor. That the ESL students may be processing these subtests differently from their monolingual peers is suggested by the fact that, in general, there is a low correlation between arithmetic and digit span scores and the performance scale, and, for arithmetic, a high correlation with the information and vocabulary subtests (Sattler, 1982, pp. 174, 179). In the study carried out by Cummins, however, arithmetic and digit span were the highest of the verbal subtest scores for ESL students, while information and vocabulary were the lowest (Cummins, 1984a, p. 25).

Certainly, in nearly all the sample assessments cited by Cummins there was a discrepancy between the VIQ and the PIQ, the latter being higher than the former in all cases, with many discrepancies

[6] This argument with respect to the relative ease with which mathematical concepts can be understood visually will be supported below by data from another influential study carried out by Collier (1987). For more insight into the mathematics achievement of second language learners see Carey (1993c).

in the "significant" range (around 19 points) as well as several in the "extreme" category (ranging from 32 to 44 points)[7]. The point that Cummins is making in this analysis, is that if these students can achieve well on these performance subtests, they should certainly not be considered low in general intelligence. On this he and I are in perfect agreement. Something must be interfering with these students' ability to perform well on the verbal subtests, however, and, as was discussed in chapter 1, there is no shortage of potential interfering factors. As noted above, Cummins believes that cultural and linguistic factors are of primary importance and it does seem logical that they should play an important role. Nevertheless, an alternate explanation for the low scores on the verbal subtests of the WISC-R for ESL students may be that, because the students are processing information preferentially via the visual system, the activity of the auditory system is reduced and therefore the processing of information by means of this system is less efficient[8]. Some anecdotal evidence supporting this

---

[7] It should be noted that no specific ranges for the significant and extreme categories were given. The ranges shown have been taken from comments made by assessors with respect to certain discrepancies.

[8] Although this compensatory action between sensory systems could be considered to be a "balance effect" of sorts, it has no kinship whatsoever with the "balance effect" hypothesis of Macnamara (1966), which states that, for bilinguals, a development in competence in one language will lead to a decrease in competence in the other. However, the move to greater dependence on visual information gathering for the second language learner does lend some support to Macnamara's (1974) suggestion that nurseries should be used as models for language classrooms in that "using facial expressions, exaggerated tones of voice, gestures and actions" (p. 92) as a mother does automatically with her child should aid the second language learner in encoding the meaning and sounds of the

interpretation can be found in comments made by both teachers and psychologists describing the performance of students recommended for assessment. Apart from the universal discrepancy in VIQ and PIQ, specific reference was made in certain cases to problems with auditory memory:

> a) He speaks Italian fluently and English well.... His attention span is very short. He is always very easily distracted. His auditory memory and discrimination skills seem to be below average but his visual memory is better, and his visual discrimination skills seem to be fairly good. (p. 33)

> b) While she is progressing in English, she is still behind, resulting in a low overall verbal score. Performance (score) was within the average range and this may well be a measure of the girl's potential.... No real disability is obvious other than auditory memory and a rather impulsive manner of attacking her work. (pp. 52-53)

> c) His only strength was speed in copying designs. His auditory memory is quite weak which may be holding back his English development. (p. 54)

Since only a small number of excerpts from the referral and testing documents were cited, it would be interesting to find out just how

---

target language.

representative these findings with respect to auditory memory and discrimination difficulties are as they may be indicative of the fact that the root of the "linguistic" problems for these students lies in a reduced ability to process information auditorily, at least as far as speech is concerned.

Another interesting finding with respect to VIQ/PIQ discrepancies brought out by Cummins during his analysis of these results for ESL children is that

> a large proportion of children in a monolingual context who are characterized as suffering from a specific reading disability, manifest low scores on the Verbal subtests of the WISC-R although no impairment is usually detectable in the surface aspects of interpersonal communication skills. (p. 50)

Using the logic that a VIQ/PIQ discrepancy may denote a shift in information processing towards the visual, a possible explanation for the problems experienced by these monolingual students is that they too are processing information preferentially via the visual system. Undoubtedly, among the ESL students, some of the cases of the most extreme VIQ/PIQ discrepancies cited by Cummins have been referred for testing because of reading problems (see Cummins, 1984a, p.45, cases 9 and 11, and pp. 50-52). As a consequence of this well recognized link between VIQ/PIQ discrepancies and poor reading ability (see Cummins, 1984a, pp. 49-51), it seems reasonable to assume that any situation which would promote such a discrepancy would also tend to produce reduced reading ability and hence, poor academic achievement, among learners. This point will be taken up again later.

To return to the main thread of this discussion, Cummins has shown clearly that using general intelligence tests on ESL students holds problems for interpretation because of the possibility of "linguistic" demands biasing results. It is precisely for this reason that many researchers use "nonverbal" measures of cognitive ability for assessing second language learners. As Reynolds (1991) points out, although there are a number of different tests available, some of them are "used much more frequently than others, for example the *Peabody Picture Vocabulary Test* and the *Coloured Progressive Matrices*" (pp. 159-160). Closer examination of the Coloured Progressive Matrices test shows that, apart from an initial verbal explanation, no language knowledge is necessary to complete the tasks since they are based on the matching or completion of visually presented patterns. Reynolds (1991) notes that the best correlation between this test and the Wechsler Preschool and Primary Scale of Intelligence is found with the Picture Completion subtest (.43), and that "[b]ecause a respondent can adopt either a verbal/analytic approach to the test or solve the matrices through visual perceptual discovery, this is not a 'pure' test" (p. 161). The fact that many of the claims for "cognitive advantages" have been based on this test is interesting, therefore. Any putative advantages that may be emerging may simply be an ability to process information more efficiently through the visual system. Hence, whether or not this constitutes an "advantage" rests on the characteristics of the task to be analyzed.

To summarize briefly, it is clear that second language researchers have recognized that second language learners are much more adept at tests which have a minimal dependence on the target language and/or culture. What they may not have realized fully is the extent to which these tests are dependent on an ability to process information visually, and therefore the extent to which second language learners are adopting, whether consciously or unconsciously, a learning strategy which favours visual processing.

(2) A study by Rafael Diaz on "Bilingual Cognitive Development"

A prime example of a study which claims to have discovered "cognitive" advantages for ESL students based in part on CPM results, and one which relates to the debate on the Threshold Hypothesis proposed by Cummins (see Hawson, in press-a), is that conducted by Diaz (1985). In this research, Diaz was specifically interested in testing the effect of degree of bilingualism on cognitive ability, as well as exploring the cause-effect relationship between the two. Arguing from the logic of the Threshold Hypothesis, he drew up the following hypotheses:

> As suggested by previous research findings, it was hypothesized that there exists a positive relation between degree of bilingualism and cognitive ability for those children possessing relatively high second-language proficiency. No relation or a negative relation was expected for children of low second-language proficiency. It was hypothesized also that degree of bilingualism at time 1 would predict cognitive abilities at time 2, supporting a cause-effect model in which degree of bilingualism is the causal factor affecting children's cognitive abilities. (p. 1378)

The children tested in this study were Hispanic in background, and were attending kindergarten and grade one classes. They were divided into two groups, a low English proficiency (LEP) group and a high English proficiency (HEP) group. Both groups had been placed in bilingual programs due to their "inability to participate in mainstream classes conducted in English" (p. 1378). They were not, therefore, "balanced bilinguals" of the type represented in the Peal and Lambert (1962) and other studies which consistently reported positive comparisons with monolingual peers. Students in the LEP group were indeed very low in English ability, the majority of them at the beginning of the school year being able to "produce only isolated English words, as judged by their performance on a story-retelling task," while the HEP group "could produce complete sentences but only few with correct use of prepositions and verb tense" (Diaz, 1985, p. 1378).

Tests of cognitive ability were administered to the children on two separate occasions, one at the beginning and one at the end of the school year, and consisted of one subtest of analogical reasoning, three of metalinguistic awareness, and the CPM test. The results were indeed interesting and complex. However, it suffices here to report that there was support for the contention that bilingualism is causal in its effect on cognitive ability as it was defined by the tests, and that, within both the LEP and the HEP groups, English proficiency at time 1 predicted significant proportions of the CPM test performance at time 2 ($p<.001$ and $p<.01$ respectively). In his summary, Diaz commented:

> Degree of bilingualism is a strong predictor of cognitive ability for children of relatively low second-language proficiency. On the other hand, the relation between degree of bilingualism and cognitive variability seems to diminish for children of relatively high second-language proficiency. (p.1384)

He noted later (Diaz and Klinger, 1991, p. 178) that this attenuation in the predictive ability of degree of bilingualism for cognitive variance was corroborated by the results of a longitudinal study by Hakuta (1987). Thus, the predictions of Cummins' Threshold Hypothesis were not supported by Diaz's study, nor Hakuta's. In fact, the results were quite the reverse of expectations. As far as the conclusion with respect to the direction of the causal arrow is concerned, this can be explained by the dependency of the tests on visual processing ability, an argument which will be elaborated below.

Though the validity of these results was marred by the interference of socioeconomic variables with degree of bilingualism in between group comparisons (a confounding which Diaz attributed to differences in percentage of parental employment, length of residence in the United States, and degree of stability in this highly mobile population), the results of within group comparisons were convincing enough to lead him to formulate a new threshold hypothesis, "namely, that degree of bilingualism will predict significant proportions of cognitive variance only *before* a certain level of second-language proficiency has been achieved" (Diaz, 1985, p. 1386). Diaz summarized the two threshold hypotheses graphically as shown in figures 3.3 and 3.4 (p. 314). It is immediately evident that the shape of the graphic representation of Diaz's threshold hypothesis is very much in keeping with that of the shift in information processing hypothesis in figure 3.1. This correspondence between the two graphs may highlight an information processing answer to a question raised by Diaz in his summarizing comments, and repeated by Hakuta a few years later:

> The present findings lend some support to the claim that bilingualism fosters the development of cognitive abilities, especially during the initial period of second-language

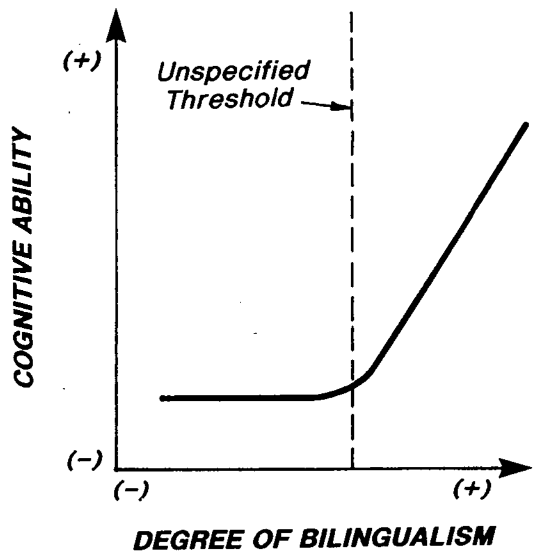

Figure 3.3 – CUMMINS' THRESHOLD HYPOTHESIS

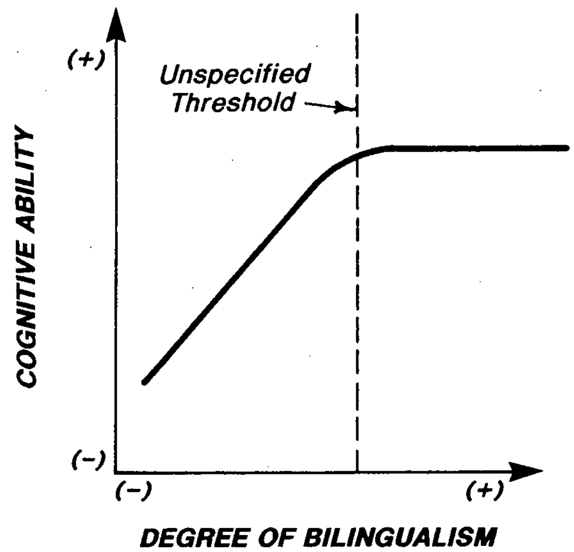

Figure 3.4 – DIAZ'S THRESHOLD HYPOTHESIS

> learning. These findings, although important, are only a first step in understanding the issues at hand. The question remains as to *how* bilingualism affects cognitive abilities, especially when cognitive ability is measured by performance on nonverbal tests such as the Raven's Progressive Matrices [CPM]. (Diaz, 1985, p. 1387)

> Why should becoming bilingual, a primarily linguistic activity, have an impact on nonverbal cognitive activities and not on metalinguistic performance? (Hakuta, 1987, p. 1385)

In response to his own question, Diaz later developed an explanatory model of the interaction between bilingualism and cognitive development (Diaz & Klinger, 1991). To that same question, however, the shift in information processing hypothesis would propose that, for students newly immersed in learning a language, the propensity to depend on the visual system for the processing of information is much higher than for those students who already have a working knowledge of the language. As a result of this, their scores on a test such as the CPM which, it has been argued, assesses their ability to analyze problems visually would be only indirectly related to their degree of bilingualism. The postulated "benefits" which have been observed would then come about in the following way: the low level of ability in (and particularly the low comprehension of) the language causes an increase in attention allocation to the visual system, and this in turn causes an apparent increase in "cognitive" ability *if this is measured using tests dependent on or analyzable by visual analysis of tasks*. Since the boost in attention to the visual system should be only temporary due to the tempering of the attention effect by the gradual increase in language ability, slowly but surely the benefits in visual ("cognitive") processing should decrease, though whether the information processing balance between the visual and auditory systems return to pre-second-language levels is uncertain. Data on ESL students' academic achievement to be examined next tend to argue against such a return (see Collier, 1987 and the argumentation in section 3), as do some data on the performance of bilinguals on a Stroop test (Biederman & Tsao, 1979; see below). Thus, the tapering off of "cognitive" benefits for the HEP group in Diaz's study may reflect at another level the tapering off of the heightened dependence on visual information processing with an increase in speech comprehension.

What is being suggested, of course, is that the observed "cognitive advantages" are evidenced because of how cognition and advantage have been defined by the test, and are therefore circular in their conception and unhelpful in their explanatory power.

(3) The "Age and Rate of Acquisition of Second Language for Academic Purposes" study by Virginia Collier (1987)

This study "analyzed the length of time required for 1,548 advantaged limited English proficient (LEP) students to become proficient in English for academic purposes while receiving instruction in English in all subject areas" (Collier, 1987, p. 617). It was an American study which was intended to test the validity of the BICS/CALP dichotomy of language proficiency proposed by Cummins, and its general conclusions supported Cummins' theorizing. What were analyzed were data on the academic achievement of these students in several subject areas at various grades, and the results were broken down to show how Age on Arrival (AOA) and Length of Residence (LOR) affected performance (see pp. 628 and 629 for graphic representations of the results). One very interesting and unexpected outcome of the study was the finding that LEP students did extremely well in mathematics. Collier noted that "Math achievement of ESL graduates in the 11th grade was much higher than 11th grade achievement in other subject areas, reaching above national averages (53rd-59th NCEs) but still lower than 4th-, 6th-, and 8th-grade LEP students math achievement" (p. 631), and later commented:

> It is encouraging that advantaged[9] LEP students can perform so well in mathematics, even when portions of the exam include

---

[9] It should be noted that the meaning of the term "advantaged" here is specific to the content of the study. Collier explained that "approximately 65% of the subset of language minority students who received special ESL instruction qualified for free or reduced-price lunches, indicating that upon entry, a majority of these students came from low-income families as measured by U.S. standards" (p. 620). The advantaged designation seems to have been decided upon in consideration of the fact that the majority of these immigrant families were in the upper- or middle-income category in their country of origin.

> math concepts and problem solving, which rely more heavily on language skills. Although the remarkably high mathematics achievement was the exception to other content-area achievement, ESL graduates' scores still followed the same pattern as that found in the other four content-area tests. (p. 637)

Even more interesting than the high achievement in mathematics was the fact that there was a correlation in the pattern of mathematics results across all grades (with the slight exception already noted for grade 11 students). Not only were the math results far higher than the results in other areas for all grades as Collier remarked, but for the 4th, 6th, and 8th grades the scores were at their highest average for students with the shortest LOR, 1-2 years. Furthermore, as the LOR increased, the average scores slowly dropped, while still remaining far above the scores achieved in other subject areas.

The explanation to the correlated patterning of these results for mathematics may be, first, that due to the dependency of mathematics on universally recognizable concepts manipulable through *visual* symbols, it is particularly suited to analysis by means of visual system processing, and second, that, in keeping with the shift in information processing towards the visual system hypothesis, students who are new to the country and unfamiliar with the language process information preferentially via the visual system, especially during the initial few years in the country but perhaps on a continuing basis. The justification for suggesting that the preference shown to visual processing continues beyond the first two years comes from the evidence that not only do these students continue to do well in mathematics, but that reading scores are consistently the lowest of all academic measures tested no matter what the LOR of the students. Given the link between VIQ/PIQ discrepancies and poor reading ability for monolingual students, it seems reasonable to propose that the poor reading achievement may be indicative of a continuing discrepancy of this sort for the ESL students. Thus the pattern of results for these students would tend to support a developmental pathway for information processing in keeping with the general pattern shown in figure 3.1.

To speak to Collier's comments with respect to the necessity of some language skills in problem solving, the amount and type of language used in mathematics problems tend to be very limited and would easily be learned given adequate motivation, the sort and amount of motivation which might very well be sparked by initial success in the subject area with its consequent teacher and parental approbation. In addition, the familiarity with basic mathematical concepts which the students would undoubtedly take into the classroom at whatever the grade level of entry into the country would reduce the information processing load, freeing up the high attention, capacity-limited information processing system to tackle both the language and the particular mathematical concepts involved. It is also intriguing to note that the content area chosen by Lambert and Tucker in their 1972 study of immersion students was mathematics precisely because it was a "nonlanguage subject matter" (p. 152). In their summary of the findings on the grade 2 classes, they commented: "What was particularly interesting about the comparison was that the Pilot and Follow-up Experimental classes scored as well as the English controls on problem arithmetic and significantly better on computational arithmetic" (p. 105).

That the experimental classes did better than the English controls in computational arithmetic would have been predicted by the premises outlined above, namely that mathematics, especially the language reduced components, is a content area which is easily analyzable visually, and that second language learners in immersion circumstances develop a heightened capability to process information by means of the visual system.

(4) A note on bilinguals' performance on the Stroop test

A search for data which might speak more directly to the visual processing abilities of bilinguals brought up an interesting study on the Stroop effect. The study itself (Biederman & Tsao, 1979) was concerned with the differences in processing Chinese ideographs and

English words under Stroop conditions[10]. In trying to explain the results obtained, the authors examined data from two experiments by different authors in which there were five distinct groups of bilingual subjects of approximately equal proficiency in both languages. Their conclusions were that "bilinguals generally show *less* Stroop-test interference than monolingual controls" (p. 128). Since this ability to perform the Stroop test more quickly than monolinguals means that the bilinguals are reading the word and then suppressing it more quickly, it may constitute a small piece of evidence that bilinguals can indeed process certain kinds of visual information more quickly than monolinguals even after becoming fluently bilingual. Likewise, the shift in information processing strategy hypothesis may bring greater insight into why it is that bilinguals have this ability.

[10] The normal Stroop test consists of trying to name the colour of a particular printed word when the coloured ink spells the name of a conflicting colour. This condition usually brings about a reduction in speed of response due to the interference effect of the printed word.

3.4 Relevance and Implications of the Shift in Information Processing Strategy Hypothesis for Second Language Learning and Academic Achievement

3.4.1 General implications

Although the evidence supporting the shift in information processing strategy hypothesis presented in the previous section is by no means conclusive, it does lend some support to the contention that second language learners in immersion situations develop an increased dependence on information processing via the visual system, and that this leads to both an enhancement of the ability to learn through visual analysis of tasks, and a reduction in information processing ability via the auditory system.

The importance of such an insight stems from the diagnostic vision it affords into information processing problems. Given the integrative nature of brain processing, any alteration in the balance of the contributions of the various sensory modalities will have repercussions on what is learned and how well that learning is encoded. Knowing which systems are likely to be affected under certain environmental conditions, therefore, gives clues as to which kinds of strengths and difficulties may occur at the behavioural level. On the positive side of things, due to the increase in activation of the system resulting from increased attention allocation, second language learners in immersion situations may be expected as a group to be exceptionally good at subjects or activities which can be accomplished by visual analysis, for example, at mathematics, computerized learning, art, and visual observation tasks of any kind. On the negative side, they should be expected to behave as if they are a little "hard of hearing" as suggested by Kosslyn and Koenig (1992, p. 222) and/or inattentive to auditory instructions; they would also be easily distracted due to the overload of the capacity limited attention route and perhaps given to "daydreaming", that is withdrawal of attention from exterior sensory input because too much of it is unfamiliar. At a more serious level for academic achievement, however, a pronounced imbalance between visual and auditory processing would predictably have an effect on reading which has been shown to depend on both auditory and visual system processing. Though the original notion put forward by Alvin Liberman, among others, in the late sixties and early seventies that the acquisition of reading requires some explicit representation of the phonological structure of utterances was received with scepticism, Klima (1991) notes that more recently "it has come to be considered self-evident that phonological representations are involved in reading and learning to read" (p. 414)[11]. The initial visual input stimulates processing in two distinct routes, called the phonological and semantic systems by Hinton et al. (1993): "Two pathways in the brain are responsible for the mental processing and

[11] It should be noted that Klima's claim would only be relevant to alphabetic writing systems.

pronunciation of written words. One (the phonological route) derives pronunciation from spelling, the other (the semantic route) from meaning" (p. 78).

Another version of the processing of written words, however, has information from both auditory and visual word form routes being integrated within the "semantic system" (see McShane, 1991, p. 289 for a graphic representation of the two routes; and, for a more complex version, see Kosslyn and Koenig, 1992, p. 192). In recent times, considerable evidence has accumulated from studies of acquired dyslexia and children with reading difficulties that the phonological route is extremely important in the development of good reading (see McShane, 1991, pp. 306-315), and it has been suggested that the process of learning to read might increase phonological awareness. Bertelson and de Gelder (1991) have noted not only that orthographic representation influences how literate people represent spoken language leading to phenomena such as "'spelling pronunciation', in which the pronunciation of words with inconsistent spelling is modified in the direction of a better match with the orthographic representation" (p. 406), but also that orthographic representation exerts an influence on auditory perception according to the direction of the writing system of the language. These examples once again illustrate that auditorily and visually encoded information are highly interactive in language perception.

Hence, learners who fail to encode the phonological aspects of a language adequately may be in double jeopardy linguistically; that is, they may have difficulties in learning to read well and, as a consequence, be less able to use their reading ability to refine their perception of language. Of course, not *all* second language learners in immersion situations will become poor readers. After all, even students with hearing impairments can learn to read well. Nevertheless, as was mentioned above, there is evidence in Collier's study (1987) that reading achievement for ESL learners is low, indeed, the lowest of all achievement measures recorded, well below school system means, across all grades even after four to five years residence in the country. Neither should it be suggested that *all* difficulties in achieving academically have their source in language. But *if they do*, at least one potential route to better performance, attempting to increase phonological awareness in these students, is sign-posted.

3.4.2 On the interference effects of sociocultural factors

Although a shift in information processing towards the visual may explain why many second language learners in immersion situations have greater difficulty achieving academically than monolingual students, it does not explain why there are differences in achievement among the former group. In order to explicate these effects, sociocultural factors of the type discussed previously for majority and minority second language learners must be called upon.

It should be evident, nevertheless, that these sociocultural factors are in intimate interaction with linguistic factors due to their determination of not only the quantity and type of input available to the learner, but also where attention is directed. In the case of majority group second language learners, the prevalence of their first language as well as the value accorded to it by the society at large would tend to counterbalance any swing towards visual processing which is automatically initiated during class time. For the minority student, however, the reduced availability combined with the experienced lack of importance or meaningfulness of the native tongue within the community would promote the allocation of attention towards the learning and use of the second, more socially meaningful language. What is being suggested, therefore, is that the learning of the high value of the majority language coupled with the learning of the lack of worth of their first language which would be absorbed through myriad societal interactions, including educational, is encoded "inside the head" of these students in a very real way, and causes them to turn their attention away from the learning of their own first language. Both these sociocultural factors, then, would consolidate the tendency to process information preferentially via the visual system, a tendency which, it has been argued, it would be better to avoid if high academic achievement is sought after. From a sociocultural perspective, therefore, it would seem particularly important for the school system to promote the upkeep of the first language and culture of minority students as a general ploy for tempering or even preventing a long-term shift in information processing strategy with its attendant linguistic influences.

Corroboration of the wisdom of this approach can be found in the academic success of the private bilingual Cuban ethnic schools in Dade County, Florida. Garcia and Otheguy (1985) found that, as well as employing "*very* extensive use of Spanish" (p. 11),

> these schools provide continuity between the home and the school by recognizing and respecting not only the child's language but also, most importantly, the parent's culture and behavioural norms. What these schools "sell" to parents, then, is a sound and *familiar* education. (p. 12)

From the information processing perspective, this would indeed appear to be a sound plan, since the abundance of familiar input and the evident emphasis on the worth of the language and culture of the students would not only enable the children to process much of the information via the low attention demanding route, freeing up the capacity limited attention route for the processing of educationally relevant material, but also predispose them to allocate attention to the learning of their own language while also learning English, thereby reducing the probability of an attentional swing to the visual system.

Is the answer to helping second language students achieve better, then, simply to load up the school and classroom with culturally familiar input of all kinds and to respect and use the students' language? It would be nice if the problem were that easily solved. Unfortunately, there is no one answer which would satisfy the great diversity of second language group situations. For one thing, in inner city schools in both Canada and the United States, there is no one language or culture common to all the ESL students. In Collier's (1987) study, for example, the data gathered identified 75 different languages and over 100 countries of origin for the approximately 1500 students included. To try to do justice to even a few of these cultural and linguistic backgrounds within a classroom or a school would be a logistical nightmare. But the basic premise of this dissertation is that, by approaching the study of second language learning from a computational neuroscientific (connectionist) perspective, a deeper level of understanding of the factors involved in second language learning can be gained so that problems which have resisted solution from a more traditional approach can be brought to heel. The question that must be posed is the following: How can the connectionist perspective help the second language immersion student in the more representative multi-ethnic classroom? Some responses to this question will be forwarded in the next section.

3.4.3 General recommendations for aiding second language learning in immersion situations based on the shift in information processing strategy hypothesis

Since the basic premise of the shift in information processing strategy hypothesis is that second language learners in immersion situations will begin to depend more on visual processing of information thereby reducing auditory information processing, and that this dependency may become a long term strategy which works against academic achievement by affecting reading ability, the essential problem to be addressed here is how to temper the shift to visual processing. The answer to that question does not lie in discouraging the use of the visual system in my opinion, because the heightened ability of these students to process information visually is a source of academic success for them, as has been seen, for example, in the case of mathematics. On the contrary, their ability to achieve in fields where visual processing might give them an advantage should be encouraged so that their self confidence is bolstered during the period in which their ability to comprehend and use the second language is low, a period which for most students will be disorienting, discouraging, and mentally very tiring due to the excessive quantity of new information the brain is required to process. Nevertheless, to avoid the build up of a long term dependency on the visual system, the students must be encouraged as much as possible to attend to auditory information. How might this be accomplished in the face of a brain whose inclination is to allocate attention elsewhere? Some suggestions for promoting attention to the auditory system are detailed below. As was stated in the general Introduction, the intent here is merely to suggest general approaches to second language teaching that might redress an auditory/visual imbalance in information processing, not to critique second language teaching practices. Such a critique is beyond the scope of this dissertation.

1. Because it is possible that the whole auditory system may be suppressed by an intense focus on visual processing, tasks which are dependent on hearing, whether linguistic or musical, or a combination of the two, should be set for these students. In a combination task, it would be helpful if the music were familiar so that the information processing demands would be reduced while, at the same time, attention would be drawn automatically to the auditory system.

2. Combination tasks in which both the visual and auditory systems must be employed to reach a successful conclusion may also aid in restoring an information processing balance. In this case, it would be helpful if the visual information were familiar so that it could be processed quickly and information processing (including attentional) resources could be concentrated on the auditory component. A very simple example of such a task would be the matching of an auditorily presented word or sentence to one of several familiar pictures or scenes presented simultaneously.

3. It should be remembered that these learners are not hearing many features of the second language speech sounds very well

(for instance, phonemes and syllables) in the initial stages, and therefore there is a need for the intensity or loudness of the speech sounds directed towards them to be increased for some time in order for them to be perceived adequately, even if not understood. The readjustment of the perceptual categories of speech which must be undergone in this initial learning period may also be promoted by presenting language input to the students in an "unsupervised learning" type of situation, that is, a learning situation in which focused attention is directed elsewhere while the auditory information is being encoded. Such a situation could be created if students were given a task dependent on visual information while being exposed to good quality speech[12] (on a tape or a radio) that they were not required to attend to. Because of the low attention allocated to the auditory information, the speech would be processed by the brain in an "unsupervised" manner such that regularities in the speech input, for example prosodic features like variations in pitch, loudness, and phoneme duration as well as other phonological aspects of speech, could be drawn out as valuable due to their high prevalence.

[12] It should be remembered that many of these students may come from family and/or societal backgrounds in which either the target language is not spoken, or is not spoken with a high degree of competence. Thus, exposure to an adequate quantity of good quality language, as far as the target language is concerned, may be minimal.

4. Specific instruction of the type given to monolingual poor readers should be provided for these second language learners to increase phonological awareness[13], thereby helping with both the encoding of speech sounds and reading. However, since phonological awareness is not a unitary phenomenon, all aspects of it may not be equally important to improved reading performance. It has been shown, for instance, that Portuguese illiterates "performed very poorly in all tasks in which they had to deal with segments, irrespective of the form of the task ... They attained more substantial levels of performance, though still inferior to those of ex-illiterates, in rhyme judgment ... and in tasks involving the manipulation of syllabic targets" (Bertelson & de Gelder, 1991, p. 399). It is possible, therefore, that more emphasis should be placed on segmental analysis in this phonological instruction of second language learners. It also seems reasonable to suggest that reading should be taught to these students using methods that are dependent on sound structure or phonics to encourage development of their weaker channel as well as by visual methods which take advantage of their stronger channel.

In concluding this section, I would like to re-emphasize that under no circumstances should students in second language immersion

[13] The term "phonological awareness" is taken as meaning an "explicit representation of the phonological structure of utterances" in accordance with general practice (see Bertelson & de Gelder, 1991, p. 395).

situations be considered "cognitively" deficient, or even deficient in information processing ability, because they may require particular teaching approaches. Any "deficiency" assigned to them is purely a product of the requirements of the education system, and it is therefore up to the education system to cater to the alterations in information processing strategy that may be underlying their inability to achieve to required standards.

Chapter 4. CONCLUSION:A RE-EXAMINATION OF THE THRESHOLD HYPOTHESIS FROM THE PERSPECTIVE OF THE SHIFT IN INFORMATION PROCESSING STRATEGY HYPOTHESIS

4.1 Introduction

Since this dissertation started with an extensive analysis of the critiques aimed at the Threshold Hypothesis proposed by Cummins, it seems only fitting that in the last chapter there should be a return to the hypothesis. The intent of this chapter, therefore, is to give a final hearing to the tenets of the Threshold Hypothesis, but to do so in the light of the knowledge gained from connectionism on how learning occurs in the brain. The analysis of the critiques in chapter 1 uncovered some serious concerns of language researchers with respect to the perspectives put forward by the Threshold Hypothesis. These were listed as a series of challenges which faced the explanatory power of any alternate hypothesis of second language learning in immersion situations. In the preceding chapter, the shift in information processing strategy hypothesis was developed along the lines of connectionist reasoning and some of the challenges facing the hypothesis were dealt with. Some important concerns of critics, and indeed, Cummins' himself, remain unclarified, however, namely whether claims that the Threshold Hypothesis is in fact a deficit hypothesis for second language learners who have acquired only a low level of linguistic ability can be upheld, and whether linguistic thresholds are real and what their nature might be. In the opinion of this author, the elucidation of these matters can be most readily accomplished by comparing the developmental pathways proposed by the two hypotheses, as well as the models of cognition which underlie them. To this end, then, the first task undertaken will be to make explicit the developmental pathway that is implicit in the Threshold Hypothesis. This will be done in terms mirroring those of the shift in information processing strategy hypothesis. The second task will be the comparison of the two models of cognition for the purpose of clarifying differences and similarities between the two, and highlighting the benefits of a connectionist vision of cognition. Finally, charges that the Threshold Hypothesis is in fact a "deficit theory in the making" (Martin-Jones and Romaine, 1986, p. 31) for second language learners at the low end of the competence scale will be reexamined, and some insight into the reality and nature of linguistic thresholds will be afforded.

4.2 An Assessment of the Validity of Deficit Hypothesis Charges Levelled at the Threshold Hypothesis

4.2.1 An examination of the developmental pathway suggested by the Threshold Hypothesis

As has been stated previously, there is no doubt that in proposing his Threshold Hypothesis, Jim Cummins was endeavouring to help second language learners by trying to make sense out of the body of conflicting data which had led to biased interpretations of their abilities. While acknowledging that socioeconomic and sociocultural factors were highly interactive with linguistic factors in determining the outcome of second language education in immersion situations, he believed that there was a psycholinguistic factor capable of being teased out of other influences that had a major role in bringing about academic success or lack thereof. His thinking appears to have run along the lines of "without an adequate grasp of the right kind or level of language, cognitive functioning will be impaired, and hence academic achievement will be compromised." Thus, he offered the notion that a lower threshold of language ability (the so-called Basic Interpersonal Communicative Skills (BICS) level) must be attained to ward off cognitive deficits, and that a second higher threshold level of language ability (the Cognitive/Academic Language Proficiency (CALP) level) must be breached before the learner will gain cognitive advantages due to improved cognitive functioning as a result of having learned two languages well. That he conceptualized the cognitive involvement necessary for BICS to be less than for CALP is inherent in his use of the label "skills" for the one and "proficiency" for the other.

Because the ins and outs of the Threshold Hypothesis and the critiques aimed at it, as well as the concepts of second language learning which have emerged as a result of this characterization of the second language learner, have already been dealt with in considerable depth in chapter 1, the intent here is not to dwell on them. Rather it is to try to see them through a new lens, a connectionist one, or, more specifically, the lens of the shift in information processing strategy hypothesis detailed in chapter 3. To facilitate this, a more detailed graphic representation of Cummins' notions of the developmental pathway for second language learners than was drawn up for his hypothesis in the last chapter has been drafted (figure 4.1, p. 338). In this graphic representation, it has been assumed that, because the acquiring of a second language is a progressive, dynamic process, the change in cognitive ability may likewise be represented as a progression. It should also be noted that, in the figure(s) dealing with Cummins' Threshold Hypothesis, the vertical axis representing cognitive ability has been divided into three sections, a deficit range, a normal range, and a benefit range, as the hypothesis proposes it can be. This mirrors the division of the visual processing ability axis in the graphic representations of the shift in information

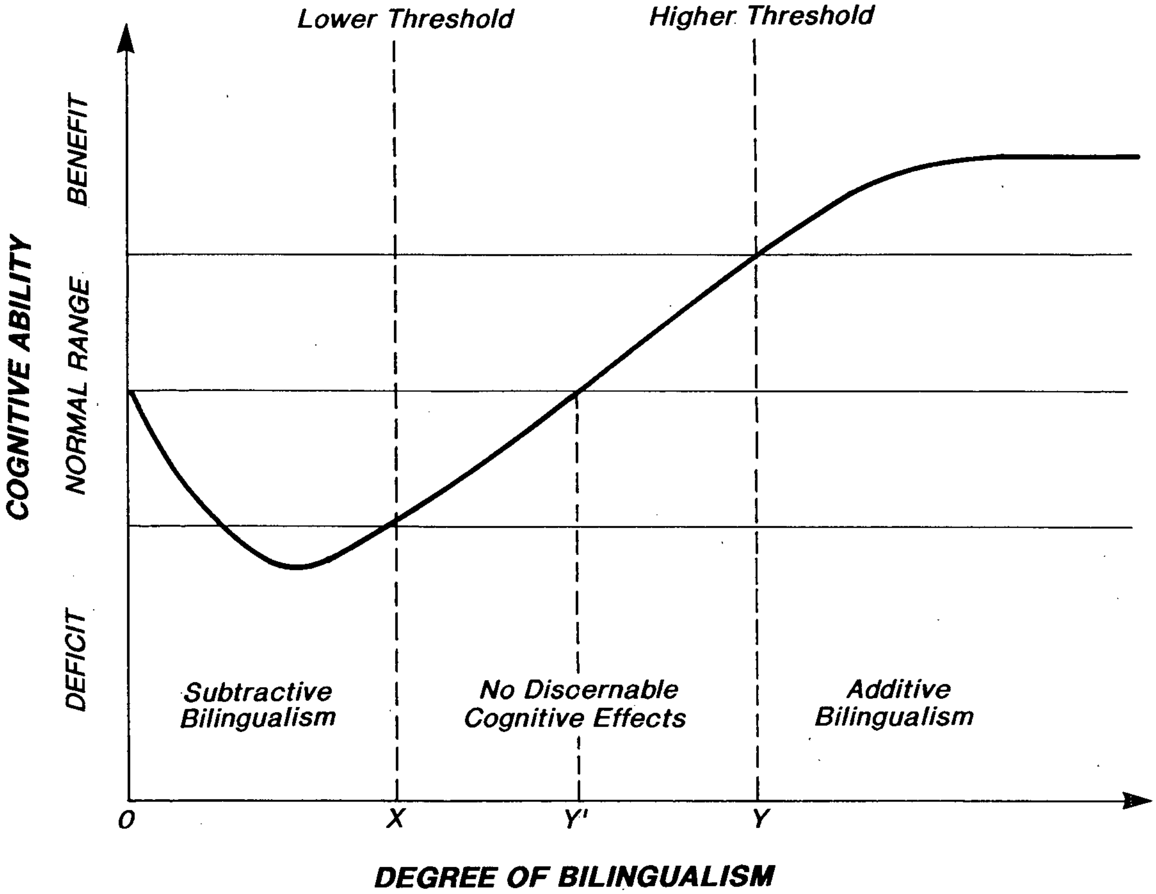

Figure 4.1 – CUMMINS' THRESHOLD HYPOTHESIS

processing strategy hypothesis shown in the previous chapter. Since it seems certain that Cummins would not suggest that second language learners from any background are in a cognitive deficit situation to begin with, and Cummins (Cummins and Swain, 1983) has even argued strongly against any assumption of this sort, the cognitive ability starting point for these students has been taken to be the mid-point of the normal range.

Starting from a position of "normal" cognitive ability, then, Cummins' Threshold Hypothesis implies that, below a certain lower threshold of bilingualism, the cognitive ability of students will drop for a period of time, dip into the deficit range, and then, as more of the second language is learned, will recover progressively until the normal range is regained with the attainment of the lower threshold of bilingualism. Of course, the actual path of this trajectory and its continuation may vary from the one shown. It is possible that, as Diaz (1985) has proposed, there would be no decline in cognitive ability for students in an additive second language learning situation (see figure 3.3, p. 314). This would constitute the upper limiting case in Cummins' hypothesis since there is no suggestion that the cognitive ability of students in any circumstance should increase during this initial period of learning. However, it is difficult to rationalize why students in an additive situation would not be subject to at least some of the same kind of intellectual confusion that is postulated to occur with students in a subtractive situation as, during this early period, their ability to interact with their academic environment would be similarly impoverished. The segment OX of the degree of bilingualism axis, therefore, delineates the range of subtractive bilingualism. Once the lower threshold is attained, an increase in cognitive ability, presumed gradual, would have to take place if the students were to reach an additive bilingualism situation in which cognitive benefits are accrued, that is, beyond the second threshold of bilingual ability. Thus XY represents the range of bilingual ability in which no cognitive benefits are seen because students are within the normal range of cognitive ability, and from the point Y on, students would experience the cognitive benefits ascribed to additive bilingualism. It is interesting to note, however, that in the "no discernable cognitive effects" range, the students must at some point cross their starting ability (the mid-point of the normal range), and hence could be considered as being in a cognitive benefit position from that point on (that is from Y' on). What this interpretation presupposes, of course, is that cognitive benefits and deficits can be related to the specific case (the individual or small group) rather than to generalities based on averaged data. This might be viewed as a "weak" interpretation of Cummins' hypothesis which its dependence on standard test data does not warrant (see Edelsky et al., 1983). It is fascinating to note, nevertheless, that if this weak interpretation is adopted, the two thresholds coalesce into one positioned at Y' on the horizontal axis, and the middle range of bilingualism in which no discernable cognitive effects are observed drops out of the picture (figure 4.2, p. 342). Another feature of this interpretation is that the cognitive effects of bilingualism need never pass out of the normal range (case 2). In other words, the students could be seen as being in a "relative" deficit or benefit situation cognitively if the basic reference point is taken to be their own ability at the beginning of the second language learning, but may not be perceived as being within the range of "general" cognitive deficit or benefit. That is to say, they may be experiencing quite normal fluctuations in cognitive ability, as it is defined by the measures used.

There are some important insights to be gained from examining Cummins' Threshold Hypothesis in terms of its developmental pathway. First of all, the gravity of the non-attainment of the first threshold level of bilingualism becomes very apparent. According to the Threshold Hypothesis, the learners may become locked into a cognitive deficit situation. However, it is equally apparent from the discussion above that the position of a language threshold is not immutable, and, moreover, is not determined by language ability at all but by how the threshold is being defined. In reality, the thresholds in this hypothesis are specified by changes in cognitive ability rather than language ability. It is only because of the importance accredited to language as a driving force in Cummins' view of cognition that the two are seen as virtually synonymous here. This will be discussed further below. Clearly, then, it becomes of the utmost importance to clarify to

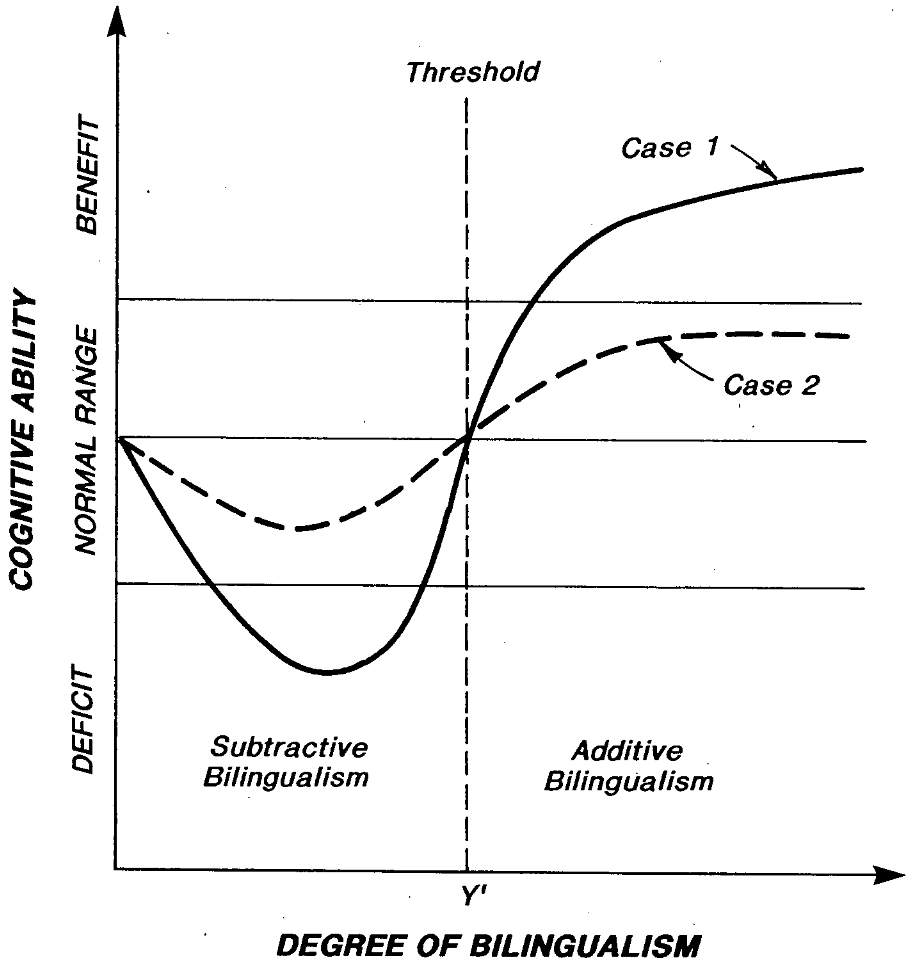

Figure 4.2 – "WEAK" INTERPRETATION OF CUMMINS' HYPOTHESIS

what extent the possibility of a cognitive deficit is real, and, in order to do this, it is necessary to examine how valid the conceptualization of cognition which underpins the hypothesis is, as well as whether or not linguistic thresholds actually exist, and, if they do, what their nature might be.

4.2.2 Comparison of the model of cognition underpinning the Threshold Hypothesis with the connectionist model

As was detailed in the first chapter (section 1.3.2), it would appear from analysis of Cummins' work that cognition was thought to consist of a set of "basic" abilities which could be applied to the learning of particular tasks, and that the knowledge gained through environmental interaction could be fed back into these basic cognitive abilities to alter them. Certain tasks were considered intrinsically more cognitively demanding than others, and cognitive abilities themselves were thought to be hierarchical in nature. The processing of familiar information was considered less cognitively demanding than the processing of unfamiliar information. Language was seen to be an especially effective force in altering cognitive abilities, with bilingual speech actions generating a qualitatively different kind of feedback for cognition to act upon. It was through the mediation of language that sociocultural, educational and political factors were thought to interact with cognition. In short, cognition was conceptualized as a global feature of brain functioning which operates relatively autonomously but can be affected by specific environmental input, even though the representation of the knowledge acquired can maintain its own functional integrity. An outcome of this view of cognition was that performance on psychological tasks which purport to test a specific ability was taken to be indicative of global cognitive competence.

Though there are several aspects of the view of cognition implicit in the Threshold Hypothesis which resonate with those of connectionists, notions on the processing of familiar and unfamiliar information, for example, or the ability of environmental input to have a profound effect in influencing cognitive abilities, there is much in the two perspectives that is at odds. On the connectionist view, for instance, a global concept of cognition is seen as potentially misleading. The perspective taken is that not enough is known about cognitive functioning to conceptualize it adequately, and so the best approach to take is one of deconstruction of brain functioning into aspects that are known to be involved, and to proceed cautiously from there. Undoubtedly, learning and memory are at the heart of cognition, hence it is believed that investigation of how these function holds good potential for elucidating other aspects of what has been defined as cognitive functioning. In this dissertation, then, since learning is basic to memory, learning was considered the best starting point for probing how cognition works. Upon examination of how learning itself occurs in the brain, it was discovered that even this one aspect of cognition is not a unitary concept, but can be deconstructed into various modes of learning, some directed by top-down and others by bottom-up mechanisms depending on attentional involvement. Attentional allocation to input was seen to be an important determinant of what was learned and how well it was learned. But attentional allocation in turn is known to be dependent on the allocation of value to input, and the encoding of value in its turn is considered dependent on present and past emotional state, relevance of the input, and its prevalence in the environment. That is to say, learning in one system in the brain cannot and should not be conceptualized as occurring independently of the function of other systems, and the function of these systems cannot be considered to be independent of learning. Furthermore, neither learning nor information processing in general can be considered free from the influences of sociocultural factors of all types quite simply because these are sources of input at large in the environment and are effective not only in directing attention but in determining which information is most readily available in the environment. Moreover, it is erroneous in the connectionist view to perceive the processing of some tasks to be intrinsically more cognitively demanding than others. What seems to determine what constitutes a high cognitive load in any task is whether or not it is unfamiliar and therefore demanding of a high quantity of attention. Thus, if one has adequate knowledge and practice in dealing with what are considered extremely difficult problems, they will not be highly cognitively demanding. That is to say, the cognitive demands of any task are relative to the knowledge base of the individual. It is also erroneous from the connectionist perspective to separate mind or mental functioning from body, as emotional state, which is influential in directing attention, is a readout of body state, and consciousness is dependent on attending to what the brain is processing.

The structure of cognition that emerges from this investigation into the connectionist vision of how learning occurs in the brain is that of an immensely interactive dynamical system of systems, sensory systems, attentional systems, memory systems, motor systems, preprocessing systems, subcortical systems, in short, systems of knowledge many of which would never have been considered to be involved in cognitive functioning in traditional models. Although it would seem that there is hierarchical organization in the brain, it does not appear to function at the level of the exclusion of the phylogenetically older structures from involvement in cognitive functioning. Indeed, the functioning of cognition cannot be considered independent of that of its constituent systems, the basic language of operation of which is better characterized in connectionism as subconceptual rather than conceptual or linguistic as it has been reputed to be in traditional models of cognition. Thus, though language as a system of knowledge is seen to be an important influence on cognition just as in the model of cognition implicit in Cummins' work, it is only in its capacity to convey condensed information to the brain that it wields its power, and not as a privileged source of data more worthy of attention than others.

One of the benefits of reconceptualizing cognition in connectionist terms is that it can be viewed as a state space of a multitude of interacting variables (systems, and within them mini-systems) all intent on proffering information for the brain to act upon. Out of this new vision, then, comes the realization that any alteration in the contribution of one variable may bring about compensatory changes in the contributions of others so that the amount of information feeding into processes such as thinking designated as cognitive may be relatively unchanged, though the conclusions come to may differ from those based on different informational input. For example, it is suggested that, in the case of second language learners in immersion situations, the change in language situation brings about an alteration of visual and auditory functioning as information gatherers, the visual system providing a greater proportion than normal, and the auditory system a lesser proportion. This does not mean that the cognitive functioning of these learners is either improved or diminished, but only that the decision making processes that are central to cognition as it has been defined are using a different mix of information to reach their conclusions. Any benefits or deficits ascribed to the performance of these students can then be seen to be constructed out of what have been declared benefits and deficits by society, whether the educational, psychological, or cultural cohorts therein or society at large.

4.2.3 Deficit hypothesis charges revisited

It is clear from the developmental pathway for Cummins' Threshold Hypothesis that this theoretical construct is, in part at least, a deficit hypothesis. In general, second language learners who have not reached the first language threshold are considered to be in a cognitive deficit position. As was discussed previously in both chapter 1 and chapter 3, this point of view has been vigorously criticised, firstly on the basis of the methodological problems inherent in the research designs of the studies used to support such a contention (MacNab, 1979; Reynolds, 1991), and secondly, as a result of the insight gained from more recent studies pointing to cognitive benefits being evident in the early stages of second language immersion (Diaz, 1985; Hakuta, 1987; Diaz & Klinger, 1991). The shift in information processing strategy hypothesis, on the other hand, does not make predictions as to what might happen to cognitive functioning per se, and zeros in instead on what might be happening to certain aspects of cognitive functioning, namely visual and auditory information processing. To be specific, it predicts an enhancement of the students' ability to process information visually, especially in the early stages of second language learning, but also a concomitant reduction in their ability to process information auditorily, particularly as it is related to the second language.

But, it may be countered, is this not also in part a deficit hypothesis? The short answer is "Yes, it is." The difference between the two hypotheses, however, is that Cummins' is predicting a non-specific *cognitive* deficit, that is one affecting the whole of cognitive functioning, while the other is predicting an *auditory system* deficit only, and one which should only be operative under certain environmental situations. In the shift in information processing strategy hypothesis, moreover, that predicted deficit in the auditory processing system is counterbalanced by the predicted enhancement of visual system processing. It is not a case of a deficit period being followed by a benefit one as in the Threshold Hypothesis, but of deficit and benefit periods going on simultaneously in two different areas of the learners brain so that a deficit in overall cognitive functioning may not occur. Hence, once again, it is suggested by this hypothesis that any deficit in performance perceived through the educational testing of these students is a construct of what is valued by educators rather than a reflection of how well their brains are working.

4.3 Some Insight into the Reality and Nature of Thresholds of Linguistic Competence

It was demonstrated in the discussion of Cummins' Threshold Hypothesis above that not only are the positions of the postulated thresholds of bilingual competence dependent on the interpretation of the theoretical constructs, but they are defined by changes in cognitive ability rather than linguistic ability. And since the close coupling between cognitive ability and language at the conceptual level has been disputed, this raises the question of how real these thresholds are. Of course, Cummins is not alone in proposing the notion of thresholds of linguistic competence: Diaz has proposed an alternate threshold hypothesis that is likewise linked to cognitive ability as has been noted, and even the shift in information processing strategy hypothesis suggests that after a certain level of comprehension of the second language has been attained there will be a reduction in the dependency of the learner on visual processing. Does this not constitute a linguistic threshold of sorts? And does this not also constitute a consensus of sorts that the concept of linguistic thresholds as a real phenomenon has some validity?

Obviously, the answer to both these questions is "Yes, it does." The problem of linguistic thresholds as it emerges from these hypotheses seems to centre not so much around whether they exist, but what their nature is (see Hawson, in press-a). In the first place, the data seem to support the existence of only one linguistic threshold rather than the two proposed by Cummins, and, according to the shift in information processing strategy hypothesis, that one would appear to be linked to level of comprehension of the second language rather than linguistic ability per se. Secondly, instead of the linguistic threshold being tied to changes in cognitive competence, it should be related in a more specific fashion to changes in information processing strategy along the proposed auditory/visual information processing continuum outlined in chapter 3. Thirdly, the attainment of a linguistic threshold may not be easily defined temporally. In the past, there has been an association of the thresholds within the schema of the BICS/CALP dichotomy with definite time periods to the onset of these thresholds (see Collier, 1987). However, the period of time necessary for any individual (or individual groups) to reach the level of comprehension necessary to alter their information processing strategy should depend on their motivation to learn, their emotional state, their familiarity with the learning environment and the culture associated with the second language, the quantity and quality of language input they receive, and the innate information processing predispositions of the learners, to name only a few of the interacting variables. Furthermore, the kind of language learned during the initial stages should be that which is most prevalent in the students' environment and which is therefore most useful in comprehending what is going on around them. If this is basic interpersonal communication, that is what will be learned; if it is academic language, that is what will be learned[1]. In most cases, of course, it will be a goodly mixture of the two. The point being made here is that, since the brain does not operate on linguistic but on mathematical principles, it is unlikely that it would adhere to the hierarchical progression of difficulty in language (or any other subject) that has been defined by linguists. It learns what must be learned for adequate functioning in the new environment, and will take its cues as to what is of value from the kind and level of language it is exposed to. As has been said many times before, the brain is nothing if not pragmatic, and its pragmatism is rooted in its function as keeper of the mind/body for survival purposes.

---

[1] For example, there was considerable variation in the levels of second language competence among adults who had entered the country at different ages and presumably had different needs and motivations in learning the target language (Johnson & Newport, 1989), and there are results showing that late immersion grade 10 students scored equally well on French language tests as early immersion students in spite of the latter being exposed to approximately three times the number of hours of instruction (see Collier, 1989).

4.4 Concluding Remarks

Although the impetus for writing this dissertation was to protest the cognitive implications of Cummins' Threshold Hypothesis for second language learners who did not manage to attain a high level of second language ability, it is to be hoped that the ideas set out here will have a longer reach than merely altering perspectives on this hypothesis, or even altering perspectives on second language learning. The concept of cognitive functioning that is emerging from computational neuroscience, cognitive psychology, and artificial intelligence as it is construed in connectionism deserves a wider audience than that. It merits being heard throughout the educational world: academics and teachers alike would benefit from at least some exposure to its basic tenets.

It is true that this newer vision of the structure of cognition is complex, straddling many so-called discrete disciplines, and demanding a willingness to learn a good measure of biological and psychological, as well as mathematical, concepts for comprehension. It is also true that as a theoretical (and experimental) approach to the study of the nature and structure of cognition it is incomplete, and moreover, is in a state of constant flux due to the ever-increasing flow of new information swelling its data banks. Does this alone not constitute good reason to avoid applying it to educational matters? Of course, my answer to this would be very strongly in the negative. My reasons are simple: firstly, no theory (or theoretical approach) can be considered 100 percent correct since it is dependent not only on the validity of the data base which supports it, but also on the correctness of the interpretation of the data; secondly, the data base which underpins connectionism is nothing if not detailed, carries with it the precision of thinking and concern for reproducibility that preoccupies scientific research, and is therefore of greater reliability than many; and thirdly, as has been pointed out by Churchland and Sejnowski, the best way for developing theories which cover overlapping territories to progress is through co-evolution, the one influencing the others. Surely no one is going to deny that neuroscientific researchers, cognitive psychologists, educators, and even artificial intelligence researchers are all interested in finding out how the brain works, how intelligent creatures learn. Should we not, then, at least be conversant with one another's perspectives on issues of central importance to us all? I firmly believe that we should, in spite of all the difficulties and hard work that such an assertion implies.

The benefits for educators that an involvement in these areas which have previously been considered "beyond our ken" would yield are not merely those that would arise from having access to analytical tools which would arm us with a higher resolving power to probe educational problems. With a clear understanding of what is being said within these other disciplines, we would be able to contribute to the debate, to use educational research data to constrain thinking in these other areas and perhaps initiate research projects which would be tailored to some extent to address issues of critical interest to educators. This would indeed mark the beginning of a new era in education.

GLOSSARY

additive bilingualism: the adding on of a second language to a well-developed first majority language with no loss to that first language.

algorithm: a systematic procedure or recipe for carrying out a computation. An instantiation of a rule specifying a computable function.

approach: the method used or steps taken in setting about a task.

attention: [folk psychological] the act or faculty of mentally concentrating on a single object, thought, or event, especially in preference to other stimuli; a state of consciousness characterized by such concentration.

attention: [as defined in this dissertation] a facilitatory mechanism for the selection of information for further processing. It operates in cooperation with expectation and intention to act, and can be driven by voluntary- or stimulus-initiated mechanisms.

cognition: [psychological] a broad (almost unspecifiably so) term which has been traditionally used to refer to such activities as thinking, conceiving, reasoning, etc.; any class of mental "behaviours" (using that term very loosely) where the underlying characteristics are of an abstract nature and involve symbolizing, insight, expectancy, complex rule use, imagery belief, intentionality, problem-solving, and so forth.

cognition: [connectionist, as developed in this dissertation] any form of mental activity that has learning (in its broadest sense) as its basis.

cognitive: [psychological] of or pertaining to mental processes such as perception, memory, judgement, complex rule use, reasoning and so forth; [connectionist, as developed in this thesis] of or pertaining to any mental processes that involve learning.

computer: a physical system computing some function where the inputs and outputs are taken to represent the states of some other system by an external advisor.

connectionism: a recently developed approach to the phenomena of human cognition that is at once a) naturalistic, b) reductionistic, and c) capable of explaining both the radical plasticity of human consciousness, and its intricate dependence on the extended cultural surround (Churchland, P. M., 1989, p. 130); parallel distributed processing (Clark, 1989, p. 83).

deficit: a disadvantage or handicap

deficit theory: [in education] a theory which posits that a learner or group of learners is educationally disadvantaged in some manner, particularly cognitively.

heuristic: an argument intended to stimulate interest as a means of furthering investigation.

hypothesis: a provisional theory set forth to explain some class of phenomena, either accepted as a guide for future investigation (*working hypothesis*) or assumed for the sake of argument and testing.

immersion language learning situation: a language learning situation in which learners are surrounded by the speech sounds of only the target language and are expected to be able to function normally. In immersion language learning situations, the second language is taught through content rather than direct language instruction, and is delivered at a level appropriate to the knowledge of the learner. That is, the second language is delivered in a simplified fashion.

information processing: the passing on of messages from one computational unit (biological or artificial) to others. The information processing approach to the study of learning is an attempt to specify the processes that operate to extract information from the sources of environmental stimulation available to the learner.

law: a ratified research hypothesis which has been shown to have generality and has been replicated in a wide variety of situations.

learning: [folk psychological] the act or process of acquiring knowledge or skill.

learning: [psychological] the modification of behaviour through practice, training, or experience; any reasonably permanent change in behaviour.

learning: [connectionist] the long-lasting changing of synaptic weights; the outcome of long-lasting changes in synaptic weights.

majority language group: a group whose language enjoys a high status and often, but not always, a high prevalence within society.

minority language group: a group whose language has a low status and usually, but not always, a low prevalence within society.

methodology: a set or system of methods, principles, and rules used in a given discipline.

model: a simplified representation of a system or phenomenon, with any hypotheses required to describe the system or explain the phenomenon.

neurocomputation: the computational function carried out by a neuron in its capacity as information processor; the computational function carried out by networks of neurons.

neurophilosophy: a philosophical stance which views the mind/brain as a single unified entity; an investigatory approach which attempts to unite basic neurobiology and behavioural sciences into a single operational framework. (adj.) *neurophilosophical*.

neurophilosopher: a proponent of neurophilosophy.

paradigm: a disciplinary matrix encompassing everything subject to professional consensus in a given scientific community; an exemplary problem solution.

positron emission tomography (PET): an imaging technique based on detection of radiation emitted by radioactive isotopes inhaled by (or injected into) the subject. Blood flow (containing isotope) increases at the sites of increased electrical activity, so a dynamic picture of neural processing is possible. Resolution is limited to a few minutes and a few millimetres.

strategy: a plan or method (not necessarily conscious) for achieving a specific goal. Strategies are essentially problem solving devices.

subtractive bilingualism: a reduced oral communicative ability and written literacy in two languages due to the learning of a second majority language after incompletely learning a first minority language.

supervised learning: [as employed in this dissertation] a characteristic of some types of learning algorithms for setting weights in neural networks, meaning that the weight change in the network is influenced by an externally generated report on its performance - an error signal. *Unsupervised* nets can monitor their own performance through internal feedback.

synaptic weight: the strength of a connection between two neurons.

synapse: a functional contact between two cells, consisting of a presynaptic terminal bouton separated by a narrow gap, called the synaptic cleft, from an area of postsynaptic membrane containing receptors. Electrical synapses are physical connections between cells, usually allowing bidirectional flow of ions.

theory: a scheme or system of ideas or statements held as an explanation or account of a group of facts or phenomena.

weight: [in machine models] the strength of connection between one unit and another in an artificial neural network - often variable; tweaking weights is the most common method of network training. It has some properties of a biological synapse.

BIBLIOGRAPHY

Andersen, R. A., & Zipser, D. (1990). A network model for learned spatial representation in the posterior parietal cortex. In J. L. McGaugh, N. M. Weinberger, & G. Lynch (Eds.), Brain organization and memory: Cells, systems, and circuits (pp. 271-284). New York, N. Y.: Oxford University Press.

Bain, B., & Yu, A. (1987). Issues in second-language education in Canada. In L. Stewin, & S. McCann (Eds.), Contemporary educational issues: The Canadian mosaic (pp. 215-225). Toronto: Copp Clark Pitman.

Bear, M. F., Cooper, L. N., & Ebner, F. F. (1987). A physiological basis for a theory of synapse modification. Science, 237, 42-48.

Bertelsen, P., & de Gelder, B. (1990). The emergence of phonological awareness: Comparative approaches. In I. G. Mattingly, & M. Studdert-Kennedy (Eds.), Modularity and the motor theory of speech perception (pp. 393-422). Hillsdale, N.J.: Lawrence Erlbaum Assoc.

Biederman, I., & Tsao, Y-C. (1979). On processing Chinese ideographs and English words: Some implications from Stroop-test results. Cognitive Psychology, 11, 125-132.

Broadbent, D. E. (1982). Task combination and selective intake of information. Acta Psychologica, 50, 253-290.

Bruer, J. T. (1993). Schools for thought. Cambridge, Mass.: MIT Press.

Calvin, W. H. (1994). The emergence of intelligence. Scientific American, 271(4), 100-107.

Canale, M. (1984). On some theoretical frameworks for language proficiency. In C. Rivera (Ed.), Language proficiency and academic achievement (pp. 28-40). Clevedon: Multilingual Matters Ltd.

Carey, S. (1987). The Francophone school-immersion debate in Western Canada. In L. L. Stewin, & S. J. McCann (Eds.), Contemporary educational issues: The Canadian mosaic (pp. 207-213). Toronto: Copp Clark Pitman.

Carey, S. (1991a). Academic achievement and reading comprehension in minority and second languages. In R. H. Short, L. L. Stewin, & S. J. H. McCann (Eds.), Educational psychology: Canadian perspectives (pp. 335-354). Toronto: Copp Clark Pitman.

Carey, S. (1991b). The culture of literacy in majority and minority language schools. Canadian Modern Language Review, 47(5), 950-976.

Carey, S. T. (1993a). "Only in Canada you say" -- Pity!: On the generalizability of French immersion results to English immersion in other countries. In L. L. Stewin (Ed.), Contemporary educational issues in Canada (pp. 183-209). Toronto: Copp Clark Pitman.

Carey, S. T. (1993b). Mothertongue and English as a second language schooling in Lesotho and East Africa. Paper presented at the Languages in Conflict and Contact Association Conference, Maseru, Lesotho.

Carey, S. T. (1993c). The threshold and interdependence concepts in minority and second language immersion programs. Journal of Canadian Association of Applied Linguistics, 15(2), 23-36.

Cheal, M., & Lyon, D. R. (1992). Attention in visual search: Multiple search classes. Perception and Psychophysics, 52(2), 113-138.

Cheal, M., & Lyon, D. (1994a). Allocation of attention in texture segregation, visual search, and location-cuing paradigms. The Quarterly Journal of Experimental Psychology, 47A(1), 49-70.

Cheal, M., & Lyon, D. (1994b). A framework for understanding the allocation of attention in location-precued discrimination. The Quarterly Journal of Experimental Psychology, 47A(3), 699-739.

Churchland, P. M. (1984). Matter and consciousness. Cambridge, Mass.: MIT Press.

Churchland, P. M. (1989). A neurocomputational perspective: The nature of mind and the structure of science. Cambridge, Mass.: MIT Press.

Churchland, P. S. (1986). Neurophilosophy: Toward a unified science of the mind/brain. Cambridge, Mass.: MIT Press.

Churchland, P. S., & Sejnowski, T. J. (1992). The computational brain. Cambridge, Mass.: MIT Press.

Clark, A. (1989). Microcognition. Cambridge, Mass.: MIT Press.

Collier, V. (1987). Age and rate of acquisition of second language for academic purpose. Tesol Quarterly, 21(4), 617-641.

Commins, N. L., & Miramomtes, O. B. (1989). Perceived and actual linguistic competence: A descriptive study of four low-achieving Hispanic bilingual students. American Educational Research Journal, 26(4), 443-472.

Corsini, R. J. (Ed.). (1987). Concise encyclopedia of psychology. New York: John Wiley and Sons.

Crick, F. (1994). The astonishing hypothesis: The scientific search for the soul. New York, N.Y.: Maxwell Macmillan Int.

Crick, F., & Koch, C. (1992). The problem of consciousness. Scientific American, 267(3), 152-159.

Cummins, J. (1976). The influence of bilingualism on cognitive growth: A synthesis of research findings and explanatory hypotheses. Working Papers on Bilingualism, 9, 1-43.

Cummins, J. (1977). Cognitive factors associated with the attainment of intermediate levels of bilingual skills. The Modern Language Review, 61, 3-12.

Cummins, J. (1979). Linguistic interdependence and the educational development of bilingual children. Review of Educational Research, 49(2), 222-251.

Cummins, J. (1984a). Bilingualism and special education: Issues in assessment and pedagogy. Clevedon: Multilingual Matters.

Cummins, J. (1984b). Wanted: A theoretical framework for relating language proficiency to academic achievement among bilingual students. In C. Rivera (Ed.), Language proficiency and academic achievement (pp. 2-19). Clevedon: Multilingual Matters Ltd.

Cummins, J. (1984c). Language proficiency and academic achievement revisited: A response. In C. Rivera (Ed.), Language proficiency and academic achievement (pp. 71-76). Clevedon: Multilingual Matters Ltd.

Cummins, J. (1987). Immersion programs: Current issues and future directions. In L. L. Stewin, & S. J. H. McCann (Eds.), Contemporary educational issues: The Canadian mosaic (pp. 192-205). Toronto: Copp Clark Pitman.

Cummins, J. (1991). Conversational and academic language proficiency in bilingual contexts. In J. H. Hulstijn, & J. F. Matter (Eds.), Reading in two languages [Special issue]. Aila Review, 8, 75-89.

Cummins, J., & Swain, M. (1983). Analysis-by-rhetoric: Reading the text or the reader's own projections? A reply to Edelsky et al. Applied Linguistics, 4(1), 21-41.

Cziko, G. A. (1992). The evaluation of bilingual education: From necessity and probability to possibility. Educational Researcher, 21(2), 10-15.

Damasio, A. R. (1994). Descartes' error: Emotion, reason, and the human brain. New York: G. P. Putnam's Sons.

Danhauer, J. L., Crawford, S., & Edgerton, B. J. (1984). English, Spanish, and bilingual speakers' performance on a nonsense syllable test (NST) of speech sound discrimination. Journal of speech and hearing disorders, 49, 164-168.

Darwin, C. J. (1991). The relationship between speech perception and the perception of other sounds. In I. G. Mattingly, & M. Studdert-Kennedy (Eds.), Modularity and the motor theory of speech perception (pp. 239-259). Hillsdale, N.J.: Lawrence Erlbaum Assoc.

Deutsch, J. A., & Deutsch, D. (1963). Attention: Some theoretical considerations. Psychological Review, 70(1), 80-90.

Diaz, R. M. (1985). Bilingual cognitive development: Addressing three gaps in current research. Child Development, 56, 1376-1388.

Diaz, R. M., & Klinger, C. (1991). Towards an explanatory model of the interaction between bilingualism and cognitive development. In E. Bialystok (Ed.), Language processing in bilingual children (pp. 167-192). Cambridge: Cambridge University Press.

Driver, J., & Spence, C. J. (1994). Spatial synergies between auditory and visual attention. In C. Umilta, & M. Moscovitch (Eds.) Attention and performance XV: Conscious and non-conscious information processing (pp. 311-331). Camb., Mass: MIT Press.

Edelman, G. E., & Reeke, G. N. (1990). Neural Darwinism. In M. A. Arbib, & J. A. Robinson (Eds.), Natural and artificial parallel computation (pp. 211-250). Cambridge, Mass.: MIT Press.

Edelsky, C., Hudelson, S., Flores, B., Barkin, F., Altwerger, B., & Jilbert, K. (1983). Semilingualism and language deficit. Applied Linguistics, 4(1), 1-22.

Enns, J. T. (1990). relations between components of visual attention. In J. T. Enns (Ed.), The development of attention (pp. 139-158). North-Holland: Elsevier Science.

Fowler, C. A., & Rosenblum, L. D. (1991). The perception of phonetic gestures. In I. G. Mattingly, & M. Studdert-Kennedy (Eds.), Modularity and the motor theory of speech perception (pp. 33-59). Hillsdale, N.J.: Lawrence Erlbaum Assoc.

Ganong, W. F. (1991). Review of medical physiology: Fifteenth edition. Norwalk, Con.: Appleton & Lange.

Garcia, O., & Otheguy, R. (1985). The masters of survival send their children to school: Bilingual education in the ethnic schools of Miami. The Bilingual Review, 12(1-2), 3-19.

Genesee, F. (1984). On Cummins' theoretical framework. In C. Rivera (Ed.), Language proficiency and academic achievement (pp. 20-27).

Clevedon: Multilingual Matters Ltd.

Goldman-Rakic, P. (1990). Parallel systems in the cerebral cortex: The topography of cognition. In M. A. Arbib, & J. A. Robinson (Eds.), Natural and artificial parallel computation (pp. 155-176). Cambridge, Mass.: MIT Press.

Hakuta, K. (1987). Degree of bilingualism and cognitive ability in mainland Puerto Rican children. Child Development, 58, 1372-1388.

Hamers, J. F. (1991). L'ontogénèse de la bilingualité: Dimensions sociales et trans-culturelles. In A. G. Reynolds (Ed.), Bilingualism, multiculturalism and second language learning (pp. 127-144). Hillsdale, N. J.: Lawrence Erlbaum Associates.

Harré, R., & Lamb, R. (Eds.). (1983). The encylopaedic dictionary of psychology. Oxford: Basil Blackwell.

Hawson, A. (in press-a). Paying attention to attention allocation in second language learning: Some insight into the nature of linguistic thresholds. The Bilingual Revue/La Revista Bilingüe.

Hawson, A. (in press-b). A neuroscientific perspective on second language learning and academic achievement. The Bilingual Review/La Revista Bilingüe.

Haxby, J. V., Horwitz, B., Ungerleider, L. G., Maisog, J. M., Pietrini, P. (1994). The functional organization of human extrastriate cortex: A PET-rCBF study of selective attention to faces and locations. The Journal of Neuroscience, 14(11), 6336-6353.

Hebb, D. O. (1949). The organization of behavior. New York: John Wiley & Sons.

Hinton, G. E. (1992). How neural networks learn from experience. Scientific American, 267(3), 145-151.

Hinton, G. E. (1989). Connectionist learning procedures. Artificial Intelligence, 40, 185-234.

Horgan, J. (1994). Can science study consciousness? Scientific American, 271(1), 88-94.

Hoyningen-Huene, P. (1993). Reconstructing scientific revolutions: Thomas S. Kuhn's philosophy of science. Chicago: The University of Chicago Press.

Jacobs, B. (1988). Neurobiological differentiation of primary and secondary language acquisition. Studies in Second Language Acquisition, 10, 303-337.

Jacobs, B., & Schumann, J. (1992). Language acquisition and the neurosciences: Towards a more integrative perspective. Applied Linguistics, 13(3), 282-301.

James, W. (1950). The principles of psychology. New York: Dover Publications.

Jenkins, J. J. (1991). Summary of the conference: Speech is special. In I. G. Mattingly, & M. Studdert-Kennedy (Eds.), Modularity and the motor theory of speech perception (pp. 431-442). Hillsdale, N.J.: Lawrence Erlbaum Assoc.

Jensen, J. V. (1962). Effects of childhood bilingualism, I. Elementary English, 39, 132-143.

Johnson, J. S., & Newport, E. L. (1989). Critical period effects in second language learning: The influence of maturational state on the acquisition of English as a second language. Cognitive Psychology, 21, 60-99.

Kahneman, D., & Treisman, A. (1984). Changing views of attention and automaticity. In R. Parasuraman, & D. R. Davies (Eds.), Varieties of attention (pp. 29-61). Orlando, Florida: Academic Press.

Kandel, E. R., & Hawkins, R. D. (1992). The biological basis of learning and individuality. Scientific American, 267(3), 78-86.

Kandel, E. R., & Schwartz, J. H. (1985). Principles of neuroscience: Second edition. New York, N.Y.: Elsevier.

Kaplan, A. (1964). The conduct of inquiry. Scranton, PA: Chandler Publishing Company.

Klima, E. S. (1990). Comment: Linguistic awareness and metalinguistic control. In I. G. Mattingly, & M. Studdert-Kennedy (Eds.), Modularity and the motor theory of speech perception (pp. 413-422). Hillsdale, N.J.: Lawrence Erlbaum Assoc.

Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. Human Neurobiology, 4, 219-227.

Kosslyn, S. M., & Koenig, O. (1992). Wet Mind. New York: Macmillan.

Kosslyn, S. M., Flynn, R. A., Amsterdam, J. B., Wang, G. (1990). Components of high-level vision: A cognitive neuroscience analysis and accounts of neurological syndromes. Cognition, 34, 203-277.

Laberge, D., Carter, M., & Brown, V. (1992). A Network simulation of thalamic circuit operations in selective attention. Neural Computation, 4, 318-331.

Lambert, W. E. (1990). Persistent issues in bilingualism. In B. Harley, P. Allen, J. Cummins, M. Swain (Eds.), The development of second language proficiency (pp. 201-218). Cambridge: Cambridge University Press.

Lambert, W. E., & Tucker, G.R. (1972). Bilingual education of children: The St Lambert experiment. Rowley, Mass.: Newbury House.

Lindblom, B. (1991). The status of phonetic gestures. In I. G. Mattingly, & M. Studdert-Kennedy (Eds.), Modularity and the motor theory of speech perception (pp. 7-24). Hillsdale, N.J.: Lawrence Erlbaum Assoc.

Liberman, A., & Mattingly, I. (1985). The motor theory of speech perception revised. Cognition, 21, 1-36.

MacNab, G. L. (1979). Cognition and bilingualism: A reanalysis of studies. Linguistics, 17, 231-255.

Macnamara, J. (1966). Bilingualism and primary education. Edinburgh: Edinburgh University Press.

Macnamara, J. (1974). Nurseries as models for language classrooms. In S. T. Carey (Ed.), Bilingualism, biculturalism, and education (pp. 91-93). Edmonton: The University of Alberta Press.

Mattingly, I. G., & Studdert-Kennedy, M. (Eds.). (1991). Modularity and the motor theory of speech perception. Hillsdale, N.J.: Lawrence Erlbaum Assoc.

Martin-Jones, M., & Romaine, S. (1986). Semilingualism: A half-baked theory of communicative competence. Applied Linguistics, 7(1), 27-38.

McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. (1986). The appeal of parallel distributed processing. In D. E. Rumelhart, J. L. McLelland, & the PDP Group (Eds.), Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations (pp. 3-44). Cambridge, Mass.: MIT Press.

McCLelland, J. L., Rumelhart, D. E., & the PDP Group. (1986). Parallel distributed processing: Explorations in the microstructure of cognition. Volume 2: Psychological and biological models. Cambridge, Mass.: MIT Press.

McLaughlin, B., Rossman, T., & MacLeod, B. (1983). Second language learning: An information processing approach. Language Learning, 3(2), 135-158.

McShane, J. (1991). Cognitive development: An information processing approach. Oxford: Basil Blackwell.

Merzenich, M. M., Recanzone, G. H., Jenkins, W. M., & Nudo, R. J. (1990). How the brain functionally rewires itself. In M. A. Arbib, & J. A. Robinson (Eds.), Natural and artificial computation (pp. 177-210). Cambridge, Mass.:MIT Press.

Montessori, M. (1955). Childhood Education (A. M. Joosten, Trans.). Chicago: Henry Regnery Company. (Original work published in 1949)

Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. Science, 229, 782-784.

Moray, N. (1969). Attention: Selective processes in vision and learning. London: Hutchinson Educational.

Mozer, M. C. (1991). The perception of multiple objects: A connectionist approach. Cambridge, Mass.: MIT press.

Niebur, E., & Koch, C. (1994). A model for the neuronal implementation of selective visual attention based on temporal correlation among neurons. Journal of Computational Neuroscience, 1, 141-158.

Niebur, E., Koch, C., & Rosin, C. (1993). An oscillation-based model for the neuronal basis of attention. Vision Research, 33, 2789-2802.

Ojemann, G. A. (1991). Cortical organization of language. The Journal of Neuroscience, 11(6), 2281-2287.

Peal, E., & Lambert, W. E. (1962). The relation of bilingualism to intelligence. Psychological Monographs: General and Applied, 76, 1-23.

Pellionisz, A., & Llinas, R. (1985). Tensor network theory of the metaorganization of functional geometries in the central nervous system. Neuroscience, 16(2), 245-273.

Posner, M. I., and Petersen, S. E. (1990). The attention system of the human brain. Annual Review of Neuroscience, 13, 25-42.

Posner, M. I., Petersen, S. E., Fox, P. T., Raichle, M. E. (1988). Localization of cognitive operations in the human brain. Science, 240, 1627-1631.

Plunkett, K., & Marchman, V. (1991). U-shaped learning and frequency effects in a multilayered perceptron: Implications for child language acquisition. Cognition, 38, 43-102.

Raichle, M. E., Fiez, J. A., Videen, T. O., MacLeod, A.-M. K., Pardo, J. V., Fox, P. T., & Petersen, S. E. . (1994). Practice-related changes in human functional anatomy during non-motor learning. Cerebral Cortex, 4(1), 8-26.

Reber, A. S. (1985). _The Penguin dictionary of psychology_. Middlesex, Eng.: Viking.

Reynolds, A. G. (1991). The cognitive consequences of bilingualism. In A. G. Reynolds (Ed.). _Bilingualism, multiculturalism and second language learning_ (pp. 145-182). Hillsdale, N. J.: Lawrence Erlbaum Associates.

Rizzolatti, G. (1983). In J.-P. Ewert, R. R. Capranica, & D. J. Ingle (Eds.). _Advances in vertebrate neuroethology_ (pp. 261-297). New York: Plenum Press.

Rumelhart, D. E., McClelland, J. L., & the PDP Research Group. (1986). _Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations_. Cambridge, Mass.: MIT Press.

Sanford, A. J. (1985). _Cognition and cognitive psychology_. London: Weidenfeld and Nicolson.

Sattler, J. (1982). _Assessment of children's intelligence and special abilities_. Boston, Mass.: Allyn and Bacon, Inc.

Schneider, W., Dumais, S. T., & Shiffrin, R. M. (1984). Automatic and control processing and attention. In R. Parasuraman, & D. R. Davies (Eds.). _Varieties of attention_ (pp. 1-27). Orlando, Florida: Academic Press.

Searle, J. R. (1992). _The rediscovery of the mind_. Cambridge, Mass.: MIT Press.

Shatz, C. J. (1992). The developing brain. _Scientific American_, _267_(3), 60-67.

Singer, W. (1990). Ontogenetic self-organization and learning. In J. L. McGaugh, N. M. Weinberger, & G. Lynch (Eds.), _Brain organization and memory: Cells, systems and circuits_ (pp. 211-233). New York: Oxford University Press.

Smolensky, P. (1988). On the proper treatment of connectionism. _Behavioural and Brain Sciences_, _11_, 1-74.

Smolensky, P. (1986). Neural and conceptual interpretation of PDP models. In J. L. McLelland, D. E. Rumelhart, & the PDP Research Group (Eds.), _Parallel distributed processing: Explorations in the microstructure of cognition. Volume 2: Psychological and biological models_ (pp. 390-431). Cambridge, Mass.: MIT Press.

Stich, S. P. (1988). From connectionism to eliminativism. _Behavioural and Brain Sciences_, _11_, 53-54.

Summerfield, Q. (1991). Visual perception of phonetic gestures. In

I. G. Mattingly, & M. Studdert-Kennedy (Eds.), Modularity and the motor theory of speech perception (pp. 117-137). Hillsdale, N.J.: Lawrence Erlbaum Assoc.

Sutherland, N. S. (1989). The international dictionary of psychology. New York: Continuum.

Swain, M., & Lapkin, S. (1982). Evaluating bilingual education: A Canadian case study. Clevedon: Multilingual Matters Ltd.

Taylor, D. M. (1991). The social psychology of racial and cultural diversity: Issues of assimilation and multiculturalism. In A. G. Reynolds (Ed.), Bilingualism, multiculturalism, and second language learning (pp. 1-19). Hillsdale, N.J.: Lawrence Erlbaum Assoc.

Treisman, A. (1988). Features and objects: The fourteenth Bartlett memorial lecture. The Quarterly Journal of Experimental Psychology, 40A(2), 201-237.

Troike, R. C. (1984). SCALP: Social and cultural aspects of language proficiency. In C. Rivera (Ed.), Language proficiency and academic achievement (pp. 44-54). Clevedon: Multilingual Matters Ltd.

Van der Heijden, A. H. C. (1992). Selective attention in vision. London: Routledge.

Wald, B. (1984). A sociolinguistic perspective on Cummins' current framework for relating language proficiency to academic achievement. In C. Rivera (Ed.), Language proficiency and academic achievement (pp. 55-70). Clevedon : Multilingual Matters Ltd.

Werker, J. F. (1991). The ontogeny of speech perception. In I. G. Mattingly, & M. Studdert-Kennedy (Eds.), Modularity and the motor theory of speech perception (pp. 91-109). Hillsdale, N.J.: Lawrence Erlbaum Assoc.

Werker, J. F., & Tees, R. C. (1984). Cross language speech perception: Evidence for perceptual reorganization during the first year of life. Infant Behavior and Development, 7, 49-63.

Zeki, S. (1992). The visual image in mind and brain. Scientific American, 267(3), 68-76.