DOMAIN SPECIFICITY IN LEARNING PHONOLOGY

by

Yeeking Regine Lai

A dissertation submitted to the Faculty of the University of Delaware in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Linguistics

Fall 2012

DOMAIN SPECIFICITY IN LEARNING PHONOLOGY

by

Yeeking Regine Lai

Approved: ____

Benjamin Bruening, Ph.D.
Chair of the Department of Linguistics and Cognitive Science

Approved: ____

George H. Watson, Ph.D.
Dean of the College of Arts and Sciences

Approved: ____

Charles G. Riordan, Ph.D.
Vice Provost for Graduate and Professional Education

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: ____

Jeffrey Heinz, Ph.D.
Professor in charge of dissertation

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: ____

Arild Hestvik, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: ____

William Idsardi, Ph.D.
Member of dissertation committee

I certify that I have read this dissertation and that in my opinion it meets the academic and professional standard required by the University as a dissertation for the degree of Doctor of Philosophy.

Signed: ____

Irene Vogel, Ph.D.
Member of dissertation committee

ACKNOWLEDGMENTS

I would like to express the deepest gratitude to my committee chair, Professor Jeffrey Heinz, who has the attitude and the substance of a genius: he continually and convincingly conveyed a spirit of adventure in regard to research and scholarship, and an excitement in regard to teaching. Without his guidance and persistent help this dissertation would not have been possible.

I would like to thank my committee members, Professor Arild Hestvik Professor William Idsardi, and Professor Professor Irene Vogel for their input and guidance.

In addition, a big thank you to my fellow students at UD for their help. Finally, I would like to thank the financial support provided by the Linguistics and Cognitive Science Department and the Graduate Office at UD and the National Science Foundation.

TABLE OF CONTENTS

Appendix

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

The main goal of this research is to find evidence for or against the claim that humans employ a distinct learning mechanism for phonology. The learnability of logically possible phonotactic patterns was examined in experimental settings using the artificial language learning paradigm. The findings of these experiments support the idea that there is a distinct learning mechanism for phonology (as opposed to syntax) as the phonological learning mechanism is subject to computational constraints that do not appear to be found in syntactic and visual pattern learning.

Three series of experiments were conducted. The first series of experiments were designed to test the learnability of a pattern found in natural language syntax, but not in the phonology of any human language. The pattern was realized over sentences (syntactic context) and over words (phonotactic context), and the results show that human subjects were only able to learn this pattern if it was presented in syntactic context. These results support the idea that the learning mechanism for phonology is different from the one for learning syntax, and the absence of this pattern in phonology is due to the computational restrictions of a phonological learner.

The second series of experiments were designed to examine and compare the learnability of a particular phonologically plausible sound pattern which is not found in any natural languages and its counterpart which is found in natural languages. The results from this experiment show that the unattested phonotactic pattern was more challenging to learn than the attested phonotactic pattern. These results suggest that

the phonological learning mechanism is subject to even stricter computational constraints than those observed in the first series of experiments.

The third series of experiments were carried out to investigate whether the same learning restrictions revealed by the results of the second set of experiments also apply to the non-linguistic domains. The same pattern tested in the second series of experiments were embedded in sequences of shapes (visual context), and in sequences of drumbeats (auditory context). The results show that the computational constraints of a phonological learning mechanism were also observed in the non-linguistic auditory learner, but not in the visual pattern learner. This could imply that the phonological learner and the non-linguistic auditory share some but not necessarily all properties.

The specific patterns tested are well-understood and well-motivated from the perspective of theoretical linguistics and theoretical computer science. The results from these experiments provide insights into the properties of human's phonological learning mechanisms and advance the debate between domain-specific and domain-general mechanisms for human learning.

Chapter 1

INTRODUCTION

Chomsky hypothesizes that humans are endowed with distinct learning mechanisms for different cognitive functions (Chomsky, 1980). Supporters of the domain-specific hypothesis believe that the learning mechanisms used for language are specially designed for the purpose of learning language only.

Opponents to this theory hypothesize that there is a more general mechanism responsible for all cognitive functions (e.g. Skinner, 1957, McClelland & Rumelhart, 1986).

Heinz and Idsardi (2011) suggest that there are at least two different learning mechanisms which are responsible for two different subdomains of language. They propose that syntactic and phonotactic patterns are learned via two different learning mechanisms and because of this, attested syntactic patterns are inherently more complex than attested phonotactic patterns in terms of computational characterization.

In this research, domain specificity in phonological learning is investigated in two ways: 1) whether it is a distinct mechanism for phonology, i.e. whether it shares a general learning mechanism with the syntax learner, and 2) how restricted the phonological learning mechanism is, i.e. whether it is restricted by certain computational constraints.

The main goal of this research is to find evidence for or against the claim that humans employ a distinct learning mechanism for phonology. The learnability of logically possible phonotactic patterns was examined in experimental settings using the artificial language learning paradigm, and the complexity of the patterns are characterized by Formal Language Theory. The findings of these experiments support the idea that there is a distinct learning mechanism for phonology as the phonological learning mechanism is subject to computational constraints that do not appear to be found in syntactic and visual pattern learning. The method of employing the artificial language learning paradigm to explore the learnability of patterns of different complexity that is defined by the Formal Language Theory has gained popularity among linguists and psychologists in the recent years as it provides a good starting point for researchers to understand human's computational powers (Fitch and Friederici, 2012).

Three series of experiments were conducted. The first series of experiments were designed to test the learnability of a pattern found in natural language syntax, but not in the phonology of any human language. The pattern was realized over sentences (syntactic context) and over words (phonotactic context), and the results show that human subjects were only able to learn this pattern if it was presented in the syntactic context. These results support the idea that the learning mechanism for phonology is different from the one for learning syntax, and the absence of this pattern in phonology is due to the computational restrictions of a phonological learner.

The second series of experiments were designed to examine and compare the learnability of a particular phonologically plausible sound pattern which is not found in any natural languages and its counterpart which *is* found in natural languages. The results from this experiment show that the unattested phonotactic pattern was more challenging to learn than the attested phonotactic pattern. These results suggest that the phonological learning mechanism is subject to even stricter computational constraints than those observed in the first series of experiments.

The third series of experiments were carried out to investigate whether the same learning restrictions revealed by the results of the second set of experiments also apply in non-linguistic domains. The same pattern tested in the second series of experiments were embedded in sequences of shapes (visual context), and in sequences of drumbeats (auditory context). The results show that the computational constraints also played a role in the non-linguistic auditory learner, but not in the visual pattern learner. This could imply that the phonological learner and the non-linguistic auditory share some but not necessarily all properties.

The specific patterns tested are well-understood and well-motivated from the perspective of theoretical linguistics and theoretical computer science. The results from these experiments provide support that the human's phonological learning mechanisms are restricted by computational constraints and suggest that the phonological learning mechanism is domain-specific rather than a domain-general one.

1.1 Overview of the Dissertation

This dissertation is organized as follows: Section 2 discusses the computational background of attested phonotactic and syntactic patterns. Also discussed are the observed differences in computational characterizations of these patterns, along with hypotheses of how specialized the phonological learner is. The previous literature on artificial language learning paradigm, which is the general method of testing these hypotheses, is also discussed in this section. Section 3 focuses on the series of experiments that was carried out to examine the question of whether phonological learning is distinct from syntactic learning. The details of methodology, results and implications of the results of these experiments are presented in this section. Section 4 introduces the series of experiments designed to test the hypothesis of phonological learning is subjected to the subregular constraints; in this section, the methodology, the stimuli and the implications of the results are discussed in detail. Section 5 was designated to the non-linguistic experiments, which were designed to explore the role of the computational constraints, shown to restrict phonological learning, in the non-linguistic domains. The visual domain was first examined, and then followed by the auditory non-linguistic domain. A general discussion and conclusion are provided in Section 6.

Chapter 2

BACKGROUND

This chapter focuses on the background of the Formal Language Theory framework which defines the complexity of the phonotactic patterns tested in this dissertation, the experimental paradigm as well as the hypotheses of this dissertation.

The theoretical background including the computational characterizations of attested phonological and syntactic patterns are discussed in Section 2.1. The definitions of Strictly Local, Strictly Piecewise and tier-based Strictly Local classes are provided in Section 2.2. The hypotheses of the dissertation and their predictions are outlined in Section 2.3, and finally a literature review of the artificial language learning paradigm is provided in Section 2.4.

2.1 Computational Characterizations of Attested Language Patterns

The Chomsky Hierarchy (Chomsky 1956) classifies logically plausible languages and language patterns between patterns that can be learned by a Turing machine and Finite Sate Automaton. It outlines as follows:

finite $\subset$ regular $\subset$ context-free $\subset$ context-sensitive

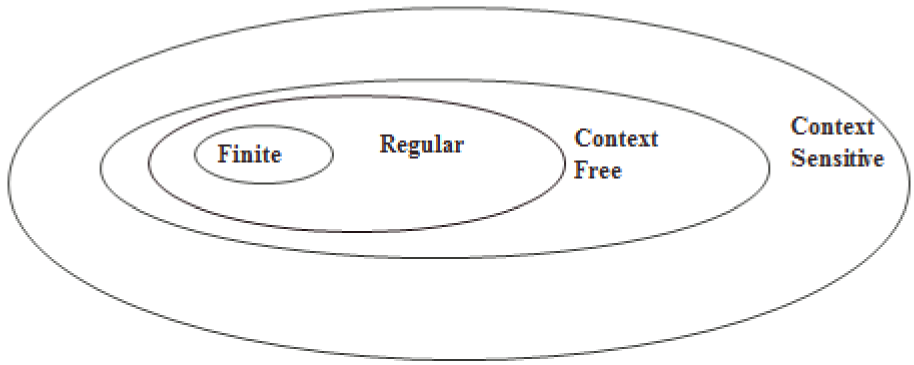

Figure 1 Chomsky Hierarchy

Some syntactic patterns, for example, have been classified as context-free (Chomsky 1956), or context-sensitive (Kobele 2006, Schieber 1985). But phonological patterns are widely recognized as at most regular (Johnson 1972, Kaplan and Kay 1994). There is evidence to support the hypothesis that phonological patterns are actually subregular (Heinz 2007, 2009, Magri 2010, Heinz 2010), meaning not all regular patterns are phonological.

There are several possible reasons why some patterns are not attested. First, it could be an accidental gap, meaning that, those patterns could be present in theory, and a learner would not have any problems with learning such patterns. Secondly, it is possible that these unattested patterns indeed exist in natural languages, but they are yet to be uncovered by linguists. Lastly, the absence of certain types of patterns could be due to their unlearnability. This means humans are incapable of learning those patterns. One method to determine whether an unattested pattern is simply accidental or unlearnable, is to test if human subjects can learn it in an experimental setting.

Although learning a phonological pattern in an experimental learning does not resemble how languages are learned in natural settings, it is a more feasible method than inducing an unattested pattern to be learned in natural settings over a long period of time.

Furthermore, these computational classes are merely classification labels of attested language patterns, and it may or may not be the case that human's learning mechanisms are indeed restricted by computational properties. One of the main goals of this dissertation is to understand whether computational properties play an important role in pattern learning. *Patterns*, in this dissertation, are defined by a set of strings of symbols, therefore, a phonological pattern referred here is a phonotactic pattern characterized by a set of legal strings of sound segments rather than a phonological alternation which involves input/output forms or underlying/surface representations. A study as such will inform us whether the typology of phonological patterns is shaped as such is at least partly due to computational limitations on learning. If so, where do the computational powers of the phonological learning mechanism fall w.r.t. the Chomsky's Hierarchy.

2.1.1 Is Phonology Regular?

Regular languages are ones describable by finite state automata, and according to computational analysis of attested phonological patterns, virtually all of them are regular (Kaplan and Kay 1994). A finite automaton has a finite number of states, and consequently, finite memory. A recursive expression requires a potentially infinite stack space to do the recursion, thus it is not possible to recognize it with a finite automaton, therefore it is not regular. If recursive patterns are defined as non-regular is accurate, then this would imply that phonological patterns are not recursive.

Some researchers (e.g. Itô and Mester 2009, Van der Hulst 2010) propose that there is recursive structure at the prosodic word level. Itô and Mester (2009) propose that prosodic words can be embedded within prosodic words based on evidence they find in English r-insertion instances. They propose a recursive structure (for prosodic words as well as for phonological phrases) as shown in Figure 2.

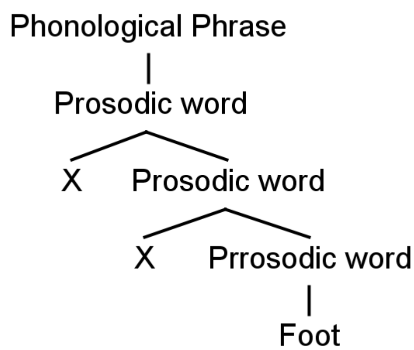

Figure 2 Prosodic word projections under Itô and Mester's analysis.

This analysis deviates from the standard proposal of prosodic hierarchy (Nespor and Vogel 1986, Selkirk 1984), which assumes that each level of constituents must be dominated by a different level of constituents, and hence argues against the occurrence of recursion at word level. More recent proposal by Vogel (2009) suggest that recursion occur at the domain which is between word and phrase. In addition, Vogel (2010) also finds compound words lacks recursion. The debate of whether recursion occurs at the prosodic word level is still going on, and whether it is necessary to assume recursion at this level remains an open question. In this dissertation, we will adopt Nespor & Vogel and Selkirk's proposals that recursion does not occur at the word level.

Moreover, most of these proposals argue that recursion occurs at the prosodic level rather than the segmental level, which is the focus of this study. Therefore, even if such claims are valid, they would not affect the assumption for the present study.

Another possible instance of recursion is reduplication. But we follow Inkelas and Zoll's (2005) analysis and consider reduplication to be a morphological operation rather than a purely phonological one. Therefore, the view that phonological patterns are not recursive is still maintained.

If phonotactic patterns at the word level is regular (not recursive), and these computational limits are psychologically real, then logically possible sound patterns that fall outside these limits ought not be learnable by humans. So if "being regular" is a firm boundary, then logically possible, non-regular sound patterns should not be learnable.

2.1.2 Non-Regular Language Patterns

While attested phonotactic patterns are regular, attested syntactic patterns go beyond this boundary. English nested embedding (NE) (Chomsky 1959), and Swiss German crossing dependencies (CD) (Schieber 1985) are context-free and context-sensitive respectively, and they are computationally more complex than regular patterns. The NE structure can be exemplified by (1).

(1) The rat [the cat ate] was brown.

An NP can embed a relative clause that modifies its head noun, and an NP that resides within the relative clause can again also embed another relative clause, for example,

(2) The rat [the cat [the boy chased] ate] was brown.

Embeddings can occur recursively, and in theory, there is no upper bound on the number of embeddings. The three nouns are the subjects of the three predicates in (2). The first noun rat is the subject of the last predicate was brown, the second noun cat is the subject of the second verb ate, and the third noun boy is the subject of the first verb chased. The nouns and verbs establish long-distance dependencies. An illustration of the long-distance dependencies found within an NE structure is shown in (3).

(3) The rat [the cat [the boy chased] ate] was brown.

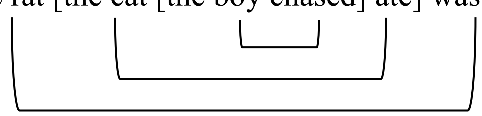

The NE can be abstractly represented by $A_1A_2A_3B_3B_2B_1$, where A represents the noun category, and B represents the verb category. The index numbers indicate the dependencies between the A and B elements. Nested embedded constructions have been reported as difficult to process in natural language (Bach, Brown, and Marslen-Wilson 1986, Blaubergs and Braine 1974, Foss and Cairns 1970).

Another non-regular construction found in human languages is crossing dependencies. An example from a subordinate clause of a Swiss German sentence is given in (4).

(4)	…mer em Hans	es	huus	hälfed	aastriiche
	…we Hans-DAT	the	house-ACC	helped	paint
	'… we helped Hans paint the house.'				

(Shieber 1985:334)

In (4), there are 3 different NPs, but only two of which are subjects of the 2 verbs. The first NP *mer* is the subject of the first verb *hälfed*, the second NP *em Hans* is the subject of second verb *aastriiche*. The dependencies between the subject NPs and the verbs are long-distance. An illustration of this type dependencies found in a CD structure is shown in (5).

(5) we Hans the house helped paint

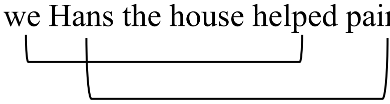

An abstract representation $A_1A_2A_3B_1B_2B_3$ can be used to describe CD (the example (5) show only two dependencies $A_1A_2B_1B_2$), where A and B belong to the noun and verb classes respectively, and the index number indicates the dependency between A and B. In terms of string language, both NE and CD are described by the language A^nB^n.

The more complex patterns NE and CD are attested in syntax but not in phonology. If it is true that phonology and syntax are both learned via the same learning mechanism then this mechanism must be able to learn the NE or CD regardless of which context it is presented in. On the other hand, if there are distinct learning mechanisms for syntax and phonology, and if the phonological learner is only capable of learning regular patterns, NE or CD patterns should not be learnable if they are embedded in phonological context. This is the basis for the set of experiments in Chapter 3.

2.1.3 Subregular Boundaries

Computational analyses of phonological patterns have made significant claims regarding what constitutes a possible phonological pattern. For example, virtually all phonological patterns have been argued to belong to the regular class within the Chomsky Hierarchy (Kaplan and Kay 1994).

However, Kaplan and Kay's work does not imply that all regular patterns are possible phonological ones. Heinz (2010) argues that phonotactic patterns actually belong to specific subregular classes, namely, the Strictly Local (SL), Strictly Piecewise (SP), and Tier-based Strictly Local (TSL) classes (McNaughton and Papert 1971, Rogers and Pullum 2011, Rogers et al. 2010, Heinz 2010, Heinz 2011a, 2011b, Heinz, Rawal, and Tanner 2011). These classes of languages are all proper subsets of the Regular languages. Informally, Strictly Local patterns refer to local dependency patterns, Strictly Piecewise patterns refer to long-distance dependencies and Tier-based Strictly Local patterns are essentially local dependency patterns operating over abstract phonological tiers (these are more formally defined in section 2.2). Figure 3 provides a schematized representation of these kinds of constraints, and classifies

Sibilant Harmony (SH), which is an attested long-distance dependency pattern, Nasal Place Assimilation, which is an attested local dependency pattern, and First-Last Assimilation (FL), which is an unattested, non-SL, non-SP, and non-TSL pattern (for details, see section 2.2). Although terms like "harmony" and "assimilation" usually indicate alternations, these terms are used throughout to refer to the valid phonotactic (surface) generalizations resulting from such alternations.

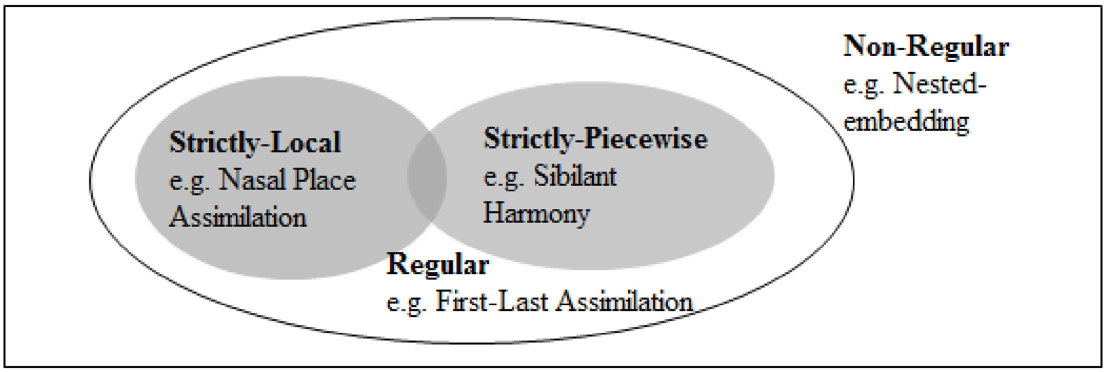

Figure 3 Subregular Boundaries. TSL is a proper superset of SL but a proper subset of regular and that it is unknown if it properly includes SP or not.

If the learning mechanism for phonology can only learn phonotactic constraints that are SL, SP, or TSL as Heinz (2010) suggests, then the absence of patterns such as FL from the attested languages can be explained: the regularities present in patterns of this type cannot be extracted by humans' phonological learning mechanism. As explained below, the specific patterns tested are well-understood and well-motivated from the perspective of theoretical linguistics and theoretical computer science. The next sections develop these ideas in more depth.

2.2 Strictly Local, Strictly Piecewise and Tier-based Strictly Local

SL patterns are those which can be described in terms of a finite set of forbidden (contiguous) sequences of symbols of length k (and is thus called Strictly k-Local). The set of forbidden contiguous sequences can be interpreted as OT-style markedness constraints such as *ab or *xy.

On the other hand, SP languages make distinctions on the basis of (potentially discontiguous) subsequences of length k (Heinz 2010, Rogers et al. 2010). A string is a subsequence of another string iff its symbols occur in the other string in order. For example, both [ʃʃ] and [oa] are subsequences of [ʃokiʃaʃ] but [ao] is not. Strictly 2-Piecewise languages are those which can be described by grammars which are sets of forbidden subsequences of length two. To illustrate a Strictly 2-Piecewise language, consider Sibilant Harmony as an example. SH requires all sibilants within a word to agree in anteriority; therefore words obeying this pattern do not contain subsequences of 2 disagreeing sibilants (i.e. [sʃ] and [ʃs] are forbidden). The set of forbidden potentially discontiguous sequences can be interpreted as OT-style markedness constraints such as *a…b or *x…y.

Tier-based Strictly Local patterns are essentially SL ones which operate on an abstract tier projected from the segmental tier. SH is also a TSL pattern because it can be described as forbidding agreeing contiguous sequences on a sibilant tier. A pattern belonging to TSL but neither to SL or SP is long-distance disharmony pattern with blocking (Heinz, 2010). Readers are referred to Heinz (2010) for a more detailed discussion on SL, SP and TSL patterns and to Rogers et al. (2010) and Heinz et al. (2011) for mathematical details.

An example of a logically possible non-SL, non-SP, non-TSL but regular pattern would be one that requires every word to have an even number of sibilants (i.e. words with an odd number of sibilants are disallowed). This pattern cannot be described by a finite set of forbidden sequences or subsequences, not even with any type of phonological tier projection. It follows that phonological learning models which can only learn phonotactic constraints that are SL, SP, or TSL will fail to learn this pattern.

Another, less bizarre, regular pattern that is non-SL/-SP/-TSL is First-Last assimilation. Words obeying this pattern require the first and the last sound segment of a word to agree in some feature. As explained further in Section 4.1.1, this pattern is phonologically plausible. Therefore, First-Last assimilation plays a central role in this dissertation, especially in Chapters 4 and 5. Henceforth, in this dissertation the term “subregular” will be reserved specifically to mean patterns belonging to the SL, SP and TSL classes.

2.3 Hypotheses

By examining the learnability of a context free phonotactic pattern, one can find out whether the phonological learner is domain specific or domain general (meaning phonology and syntax share one learner). If the context free phonotactic pattern is unlearnable, one can conclude that there are at least two distinct learners responsible for learning patterns in these two aspects of language. Hypothesis 1 (H1) and Hypothesis 3 (H3) are hypotheses which support the domain-general hypothesis Hypothesis 2 (H2) and 4 (H4) support the domain-specific hypothesis. The details of these 4 hypotheses and their hypothesis space are provided in the subsections below.

The examination of the learnability of a regular but non-subregular pattern provides evidence which either support or provide counter-evidence for the claim that the phonological learning mechanism is subject to the subregular computational constraints. If the phonological learner respsects the subregular boundary, then H4 is supported. On the contrary, if the phonological learner does not respect the subregular boundary, H4 is supported.

2.3.1 H1: Domain General

There is only one general mechanism responsible for learning phonology and syntax. Since the least restrictive language pattern attested in natural language is context-sensitive, this learner is sufficiently powerful to learn any context sensitive patterns. Following this logic, even though a non-regular phonotactic pattern is not attested, it could still be learned. Figure 4 represents the hypothetical boundary of this general learner.

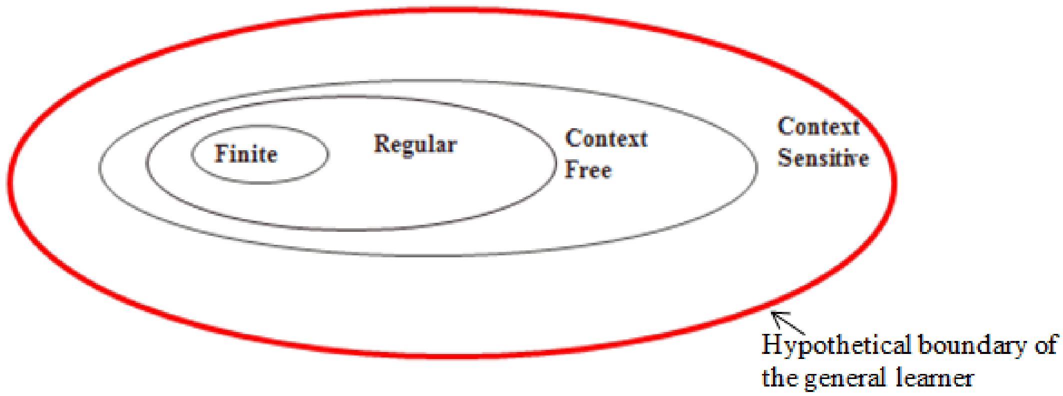

Figure 4 Hypothesis space of H1.

2.3.2 H2: Domain Specific

This hypothesis predicts the learning mechanism responsible for learning phonology is different from the one responsible for learning syntax. Figure 5 below illustrates that the respective boundaries for the phonological and syntactic learners.

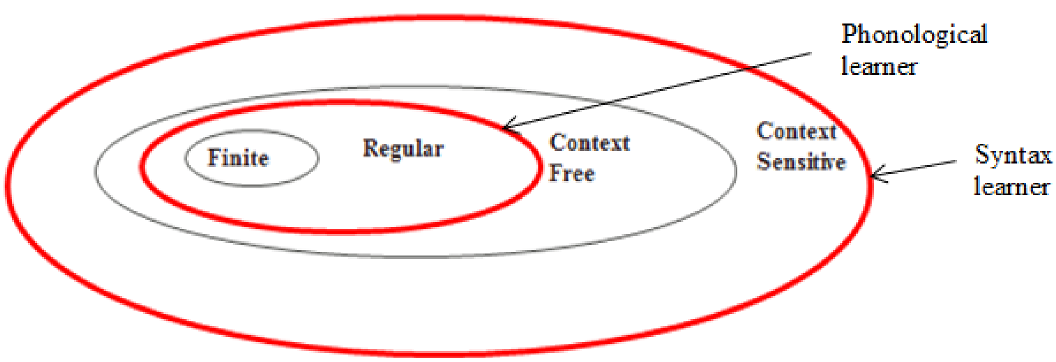

Figure 5 Hypothesis space of H2.

2.3.3 H3: Domain General (ignores subregular boundaries)

There is only one learner that is general enough to learn all attested language patterns. However, this learner is *not* restricted according to the traditional boundaries of the Chomsky Hierarchy. This learner can learn attested language patterns that are attested, but not necessarily any context-sensitive ones.

If this hypothesis is true, it predicts the boundary of what is learnable/unlearnable cuts across the computational boundaries, and implies if a context free pattern is embedded in phonological context, it should be learnable even though it is only attested in syntax not in phonology. Figure 6 below illustrates the hypothetical learning boundary of this learner.

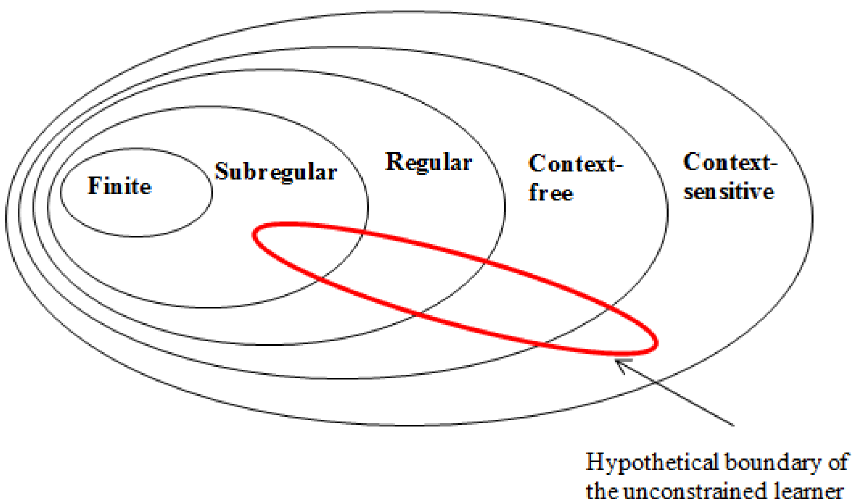

Figure 6 Hypothesis space of H3.

2.3.4 H4: Domain Specific (respects subregular boundaries)

If it has been shown that the context free phonotactic pattern is unlearnable, this suggests that the phonological learner is different from the syntax learner. However, but it would still be of interest to know, how restricted the phonological learner is. The attested phonological patterns belong to specific subregular classes (Heinz, 2010). If the phonological learner is only restricted to learning these subregular classes of phonotactic pattern, one should expect a regular pattern that resides outside of these subregular classes are unlearnable. This hypothesis can be tested by examining the learnability of a regular but non-subregular pattern, and the results can differentiate the hypothesis space assumed in H2 and H4. Figure 7 below is the representation of the hypothesis space of H4.

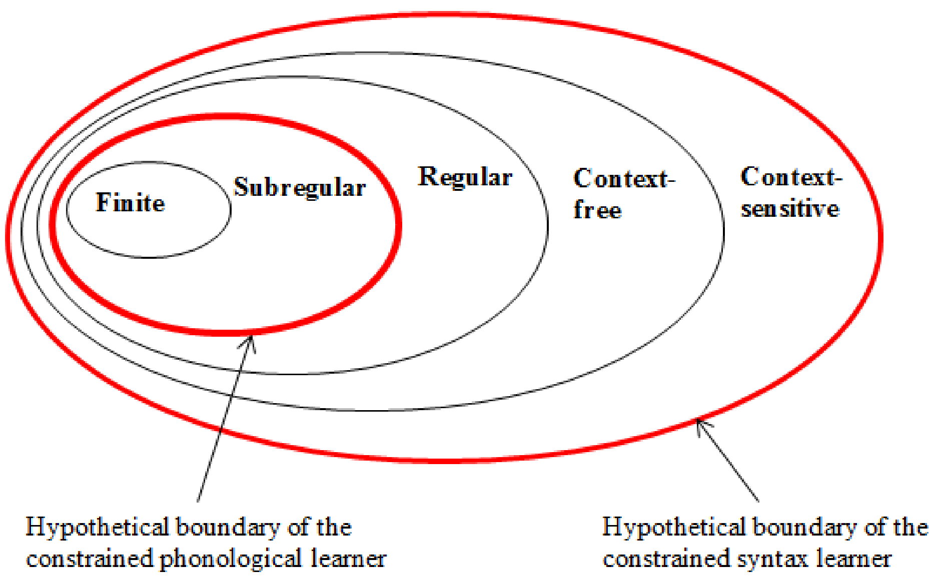

Figure 7 Hypothesis space of H4.

2.4 Artificial Language Learning Paradigm

The above mentioned hypotheses are tested by experiments using the artificial language learning paradigm. The advantage of this type of experiment is its ability to tease apart the intertwined factors of language learning that one usually sees in a natural setting. Artificial languages give experimentalists better control over what is provided to the subjects during the training phase of the experiment than they have using natural data. As detailed below, the artificial language learning paradigm (ALL) has been used in the research of language learning in both children and adults in the areas of word segmentation, phonotactics, and syntactic structures.

The basic assumption in artificial language learning research is that some learning mechanisms are shared between artificial and natural language acquisition (Gómez & Gerken, 2000; Petersson et al., 2004; Reber, 1967). Subjects who engage in an artificial language learning experiment are exposed to stimuli (the artificial language) generated by a grammar. Thus there are neither negative examples nor explicit feedback. Based on this exposure, subjects will either be able or unable to internalize the rules that define the pattern present in the stimuli. In the testing phase of the experiment, subjects are asked to respond to some questions (which typically involve acceptability tasks) that reveal whether they have successfully abstracted the intended rule.

In the following sections, the advantages and disadvantages of ALL will be discussed, and previous works on phonotactic pattern learning, and syntactic pattern learning using ALL will follow.

2.4.1 Advantages and Disadvantages of ALL Paradigm

The study of language acquisition in natural settings are limited to languages that are attested. Artificial language learning paradigm, on other hand, allows any imaginable languages to be constructed and therefore, makes it possible for researchers to study the learnability of language patterns that are either rare, extinct or even unattested. Another advantage of ALL is that researchers have control over the input including the amount of total amount of exposure, target grammar/noise ratio, the variability of data and speaker variability etc. The ability to control the input of learning in such details can eliminate many confounding factors that are otherwise unavoidable in natural settings.

Although ALL paradigm has the qualities mentioned above, it can by no means replace methodology used in natural setting such as longitudinal developmental data analysis obtained from corpuses. Clearly, learning a phonological pattern using ALL is not the same as learning phonological rules of a first language in natural settings. Children learn their first language slowly over a period of time, and the learning is often aided by prosodic, semantic and pragmatic information. ALL paradigm typically expose participants to input language for a much shorter period of learning time, and any information that could be useful for bootstrapping purposes are usually stripped off the input, or at least are kept to a minimal.

Another criticism to ALL experiments is that ALL resembles a language game. It is indeed difficult to detect whether participants treat ALL experiments as a language learning process or a language game. The problem with this is that many language games involve rules that are not found in natural languages. If such unusual rules are learnable through games, it is then reasonable for subjects to learn unusual linguistic rules through ALL experiments if they treat them as language games. Inevitably, if this is the case, would weaken the conclusions that are drawn from any ALL experimental results. One way to prevent this problem is to instruct the subjects clearly that the objective of the experiment is to learn a foreign language. This method by no means guarantees to eliminate the problem, but given the advantages of provided by ALL, the experiments in this study still adopted the ALL paradigm as a first step to explore the learnability of phonological patterns which are unattested in natural language.

2.4.2 Phonotactic Experiments

Previous studies that employed ALL have demonstrated that adults, young children, and infants are capable of using statistical cues to detect the differences between certain kinds of phonotactic constraints (e.g. Chambers, Onishi, & Fisher, 2003; Dell et al., 2000; Finley & Badecker, 2009; Goldrick, 2004; Onishi, Chambers, & Fisher, 2002; Peperkamp, Skoruppa & Dupoux, 2006; Seidl, Cristià, Bernard, & Onishi, 2009).

2.4.2.1 Adjacent dependencies

Dell et al. (2000) tested adults' ability to learn an onset/coda constraint by exposing them to grammatical words visually. Monosyllabic words were spelled out on a screen, and subjects were asked to recite the words in the testing phase. The error productions were analyzed, and they discovered that error rates were low. They concluded that the subjects internalized the phonotactic rule with visual exposure to the artificial language. A study of Goldrick (2004) adopted Dell et al.'s methodology, but modified the stimuli to test for the effect of phonetic features rather than just segments. He found that when /f/ was restricted to the onset position of a syllable (100% chance of occurrence in onset position), but /v/, which shares the same manner of articulation with /f/, occurred in both onset and coda positions (50% chance of occurrence in each position), his subjects were more likely to make errors by putting /v/ in the onset position when they engaged in the repetition task. At the same time, subjects were also more likely to erroneously produce /f/ in coda position compared to when /v/ was not included in the training set. The results of Goldrick and Dell et al. suggest that feature-based phonotactic rules can also be learned with ALL.

On the other hand, using written stimuli may not entirely correspond to how infants and young children acquire language. Subjects of these experiments had to read from the screen, and then translate the visual symbols to phonetic information. The experiments relied heavily on the subjects' ability to convert orthographical information to phonetic information, which is not a prerequisite for acquiring language in natural settings where only auditory input is available.

Onishi et al. (2002) modified Dell et al.'s experiments and showed that English-speaking adults were able to learn a non-English onset and coda constraint in relation to an adjacent vowel after brief auditory exposure to an artificial language. All the stimuli employed in this experiment were monosyllabic CVC words. The type of onset and coda consonants that could appear in a word depended on the vowel. Subjects engaged in a repetition task after they listened to the artificial language. They found that subjects repeated words that conformed to the artificial grammar significantly faster than the ungrammatical ones, and therefore concluded that the intended phonotactic rule was learned through auditory exposure. Onishi and her colleagues then extended the study to 16.5 month-old infants (Chambers et al., 2003) and found similar results.

Both infants and adults can also make use of phonotactic regularities to segment words (Aslin, Saffran, & Newport, 1997; Gomez & Gerken, 2000; Newport & Aslin, 2004; Saffran, 2003a; Saffran et al., 1996; Saffran, Newport, Aslin, Tunick, & Barrueco, 1997).

The paradigm used is referred to as the segmentation paradigm. For example, subjects are exposed to a continuous stream of an artificial language which is made up of trisyllabic words, and these words are strung together without any pauses between them. The order of the syllables within a word always remains constant, but the order of the words can vary.

The task of the subjects is to locate word boundaries based on the transitional probabilities of syllables within-words and between-words. For example, if ABC, DEF, and XYZ are different trisyllabic words and are the only words in a language (where each letter represents a different syllable), the transitional probability of B and C is 100%, but the transitional probability of C and D will be less than 100% if ABC does not always occur before DEF (i.e. it can also occur before XYZ). If the subjects can extract the transitional probabilities from the speech stream, they will be able to learn that word boundaries are present whenever transitional probabilities are less than 100%. The test is given to subjects as a two-alternative-forced-choice test, where one choice is a part of a word, and the other is a word. The subjects are asked to judge which of the two choices is a word in the language they heard during training. If the subjects can successfully segment the continuous speech stream in training based on the transitional probabilities of syllables, they should be able to pick the word over the part of a word.

This type of experiment has shown that both infants and adults can segment speech based on transitional probabilities. Similar findings have been reported for tone sequences (Saffran, Johnson, Aslin, & Newport, 1999) and visual sequences (Kirkham, Slemmer, & Johnson, 2002).

The results of these studies have been used to support the claim that a language-specific learner is not required, since statistical learning methods were shown to be able to abstract the intended rule from both linguistic and non-linguistic input. However, the patterns which were intended to be abstracted in most of these studies could be calculated using the probability of adjacent segment co-occurrences. By keeping track of which segments could be adjacent to one another the learner would have knowledge of the possible bigrams, trigrams or 4-grams. In fact, these n-gram models are statistical counterparts of the SL class (Garcia & Oncina, 1990;

Heinz 2010). Therefore, these results are consistent with the subregular boundary that Heinz proposes.

2.4.2.2 Non-adjacent dependencies with exactly one intervening segment

Non-adjacent dependencies have been noted to be ‘more difficult’ (Gebhart, Newport & Aslin, 2009), but not impossible to learn. In artificial language experiments, only certain types of nonadjacent patterns are readily learned (Cleeremans & McClelland, 1991; Finley, 2008, submitted; Finley & Badecker, 2008, 2009; Koo & Callahan, submitted; Moreton, 2008; Newport & Aslin, 2004; Onnis, Monaghan, Richmond, & Chater, 2005; Pycha, Nowak, Shin & Shosted, 2003; Wilson, 2003).

Pycha et al. (2003) have demonstrated that subjects can learn vowel harmony and disharmony rules and an arbitrary long-distance vowel dependency rule. They considered harmony rules natural and disharmony rules to be unnatural because harmony rules are more widely-attested in natural languages than disharmony rules. They intended to investigate the difference in the effect of learning natural, unnatural and arbitrary rules. In both the harmony and disharmony rules, the value of the feature [back] of the first vowel was used to predict a value of [back] for the second vowel in the template CV_1CV_2C, where CV_1C was the root, and -V_2C was a suffix, but no single feature of V1 could be used to predict V_2 for the arbitrary rule. The methodology adopted by Pycha and colleagues included three phases: listening, learning, and testing. During the listening phase, subjects listened to both underlying (root: CV_1C) and surface forms (root + suffix: CV_1C-V_2C), which were paired up in each presentation. There were two allomorphs of the suffix, and they varied in the backness of V_2. The choice of which of these two forms depended on the backness of V_1 in the root. The listening phase was followed by the learning, and then the testing phases. During these phases, the subjects were asked to choose the right allomorph of the suffix for the root presented. The subjects received explicit feedback during training after each trial, but not during testing. The results showed that subjects could learn both the harmony and disharmony rules better than the arbitrary rule, and there was no significant difference in performance between learning the harmony and disharmony rules. One possible concern about this study is the explicit nature of the learning task. Participants were given explicit feedback after each trial during the learning phase. The explicit feedback might have resembled negative feedback, which is seldom given in natural settings of first language acquisition. This, in turn, might have induced the subjects to use a different learning strategy than the one used to acquire a first language.

Wilson (2003) also tested humans' ability to learn long-distance dependencies with the ALL, but with a slight difference in the design from Pycha et al.'s. He tested whether subjects were better at learning attested assimilation rules (assimilation and dissimilation) than arbitrary dependency rules. He presented the subjects with a list of trisyllabic words of a suffixed form (CVC_1V-C_2V) that were well-formed according to each of the rules. The suffix alternated between –na and –la depending on the rule. In accordance with the assimilation rule, -na would follow a root with a nasal C_1, and –la would follow a root with non-nasal C_1. In accordance with the dissmilation rule, the value of [nasal] in C_2 would disagree with the nasality of C_1. In addition to these assimilation and dissimilation rules, two arbitrary rules were also tested. The suffix alternation in these two arbitrary rules was dependent on the [dorsal] value of C_1. In the testing phase, subjects were presented with both stimuli heard in the listening phase, and novel words, and were asked whether they had heard the word during the training phase. The novel words consisted of forms that conformed to the rule that the subjects were trained on and forms that did not. The error rates of grammatical novel and of ungrammatical novel words were compared, and showed that the subjects were more likely to think they had heard a novel word if it conformed to the grammar in both the assimilation and dissimilation conditions, but not in the arbitrary condition. These results suggest that the subjects could learn attested long-distance dependencies (with one intervening vowel) even when they were not given explicit feedback.

Learning non-adjacent dependencies has also been studied using the segmentation paradigm. Newport and Aslin's (2004) study tested subjects' ability to segment a continuous speech stream using the transitional probabilities of non-adjacent syllables and non-adjacent segments. They synthesized the stimuli to ensure that no phonetic cues for word boundaries were present. After listening to the artificial language, subjects were asked to choose between a word and a partial-word and decide which one was a legal word in the language they had heard. In order to successfully make a correct response, subjects had to extract word boundary information from non-adjacent segment/syllable transitional probabilities. The results showed that subjects were able to learn from non-adjacent segment dependencies but not non-adjacent syllable dependencies. However, Goméz (2002) showed that subjects could learn non-adjacent syllable dependencies which resembled a syntactic structure by the aXb paradigm. The symbols a and b both represented a syllable, while X was a variable that consisted of two syllables. The element X was drawn from different sets: the set size varied across conditions (from 2 elements in the set to 24 elements in the set). The results showed that the subjects learned the long-distance rule better in the high variation conditions than in the low variation conditions. One major difference between Goméz's study and Newport and Aslin's is the training stimuli. Goméz's training stimuli included pauses which clearly marked word boundaries, whereas Newport and Aslin's training stimuli were in a continuous speech stream.

In the studies discussed so far, the distance between two dependent segments /syllables was at most one segment (except in Goméz's (2002) syntax study in which she used a bisyllabic intervening word). This relatively restricted distance is a possible concern for claiming that these dependencies are long-distance, since it may be argued that the segments were placed on different tiers (such as a consonantal tier or a vowel tier) where the assimilating/dissimilating segments would in fact be adjacent to each other. However, the use of a tier-based analysis does not explain the sibilant harmony pattern attested in Navajo. Recall that in Navajo, sibilants have to agree in anteriority regardless of their positions. For example, [sototoso], [tosotoso] and [totososo] are legal words in Navajo, but neither *[ʃototoso], *[toʃotoso] nor *[totoʃoso] are legal. In light of these patterns, a line of relatively new research has begun to examine the learnability of these patterns using ALL. These studies are so new that most of them are still under review, or have just been recently submitted. We will also discuss our own pilot study examining the same problem in a later section.

2.4.2.3 Non-adjacent dependencies with two or more intervening segments

Finley (submitted; in revision) examined the learnability of long-distance dependencies with different distances by including words with different numbers of intervening segments in her study. In Finley (submitted), there were two conditions. In the first condition, subjects were exposed to trisyllabic roots $C_1VC_2VC_3$ suffixed by either –su or -ʃu. C_1 in the root form was always either [s] or [ʃ], and the anteriority of the sibilant in the suffix alternated to agree with the anteriority of C_1. Sibilants were not present in any other positions. In the second condition, only C_2 was a sibilant, and the alternation of the suffix was the same as condition 1. After the listening phase, subjects were tested using the two-alternatives-choice design. Two words with the same root but different suffixes, one legal, one illegal, were presented to the subjects at each trial, and they had to decide which suffix was the correct form. The test words included both items that were presented in training and novel words. Half of the novel words had sibilants in the C_1 position, and the other half had sibilants in the C_2 position. In condition 1, words with sibilants in the C_2 position were never presented to the subjects in the listening phase, and in condition 2, subjects were never exposed to words with sibilants in the C_1 position. The experiment was designed to see whether subjects could generalize the long-distance harmony rule to cases they had never heard of, and if they could, the results could be used to claim that the subjects internalized the rule in an unbounded way. The results showed that the subjects in both conditions could learn and internalize a long-distance dependency.

Another interesting finding reported by Finley (to appear) was that this generalization was conditioned. Using the same design, Finley found that if subjects were exposed to sVCV-su type words, they could generalize them to CVsV-su but not vice versa. Finley provided an explanation for these findings, and proposed that the sibilant harmony instantiated by CVsV-su type words was in fact local if we consider all the sibilants to be on one tier, and the harmony pattern instantiated in sVCV-su words are long-distance. Finley claims that the local pattern is more privileged than the long-distance one in terms of linguistic processing.

The local pattern illustrated by Finley can be captured by a tier-based SL model, and the long-distance pattern described can be captured by an SP model. Her results are consistent with the claim that the usage of an SL learner is prioritized over an SP learner. This means that subjects only use an SP learner when an SL learner fails to capture the pattern instantiated in the training. This is why, when subjects were trained with words of the CVsV-su type, they did not use an SP learner: an SL learner was sufficient to capture the pattern. However, when they were exposed to the sVCV-su pattern, an SL learner failed to extract the pattern, and they had to use SP, which predicted CVsV-su to be a legal word since sVCV-su was legal.

No studies that we have discussed so far tested patterns that did not belong to either the Strictly Local or the Strictly Piecewise classes. An example of such a pattern was described earlier --- First-Last Assimilation. FL resides outside of the SL and SP classes but still belongs to the regular class. Therefore this is an excellent case to test whether the absence of such a pattern from the phonologies of the world's languages is due to its unlearnability.

Koo and Callahan's (to appear) study differs from the studies mentioned so far as the long-distance dependency pattern in their study could be interpreted as position-bound. One way to understand the pattern they studied requires the subjects to learn the probability of cooccurrences of the first and the last consonants of words with 3 consonants. Interestingly, the two critical consonants do not belong to any natural class to the exclusion of the intervening middle consonant.

All of the words in Koo and Callahan's experiments were trisyllabic with the structure of CVCVCV. The language presented to the subjects can be described by the following two rules:

(a) Whenever [s] is the onset of the first syllable, [l] cannot be the onset of the last syllable.

(b) Whenever [l] is the onset of the first syllable, [m] cannot be the onset of the last syllable.

These two rules were consistent with the FL pattern, except for it is an arbitrary dependency pattern rather than assimilation. Under these rules, the sounds [s] and [l], and [l] and [m] cannot occur at a distance, but they can be adjacent to each other on the consonant tier. Koo and Callahan used a method that was similar to Wilson (2003). Their subjects were trained with words that were consistent with the grammar, and in the testing phase, subjects were presented with both stimuli heard in the listening phase and novel words and were asked whether they heard the word during the training. The error rates were analyzed, and if subjects had learned the intended rule, they were more likely to have heard a novel word before that conformed to the grammar of the language than a novel word that did not. The results of the experiments showed that the subjects erroneously recalled that they had heard the test word in the training when they were presented with a novel legal word at a significantly higher rate than when they were presented with a novel illegal word.

Koo and Callahan suggested that these results could be accounted for by the SP learning mechanism, where the dependency is learned by ignoring the actual distances between segments. This assumption is only correct if the window size that is operated on by an SP learner has to be 3 segments or above, where the occurrences of three or more discontiguous segments are monitored. However, an alternative interpretation of the patterns learned in Koo and Callahan's study is the cooccurrences of the first and the last consonant of each word. Under this interpretation, this pattern does not seem to fall into the subregular category, and it was shown learnable under their experimental settings.

2.4.3 Syntactic Experiments

There have been a number of studies that used ALL to examine the learnability of context-free grammar, but results have been mixed. Fitch and Hauser (2004) reported that human subjects learned NE easily. Their subjects were exposed to audio stimuli during training. The stimuli were constructed by the context-free template AAABBB (A^nB^n) or the regular template ABABAB $(AB)^n$. The A syllables were produced by a male voice. The B syllables were produced by a female voice. During the test, subjects were presented with novel items, some were grammatical in the language they heard in training and some were not. They were asked whether the pattern of the items was the same as or different from the pattern of the items they heard during training. Their results showed that the subjects could learn both the context-free and the regular patterns.

Fitch and Hauser were criticized for overinterpreting their results. Perruchet and Rey (2005) argued that the subjects in Fitch and Hauser could distinguish the difference between ABABAB and AAABBB when they were trained on the AAABBB language because of the acoustic reasons such as cues based on the male-to-female voice transition rather than internalizing the context-free rule. Perruchet and Rey replicated this design, and found that participants were unable to distinguish $A_3A_2A_1B_1B_2B_3$ from $A_3A_2A_1B_1B_3B_2$ sequences if the dependencies between As and Bs in the latter, were violated but were not marked by pitch distinctions. The experimental stimuli were very similar to Fitch and Hauser's but were modified to accommodate their French subjects. They also imposed a rule between the dependencies of As and Bs so that only elements with the same index number had dependencies. For example, $A_3A_2A_1B_1B_2B_3$ would be a legal string, and $A_3A_2A_1B_1B_3B_2$ would be illegal. They also used a speech synthesizer to modify the pitch of the stimuli rather than let the subjects hear them produced by a male voice and a female voice. They found that their subjects could only detect the difference when the high pitch-low pitch transition was different from the training data, but not when the nested dependencies were violated. These results suggest that NE was indeed not learned through this paradigm.

Friederici et al. (2006) and Bahlmann and Friederici (2006) also contrasted learning of a context-free grammar A^nB^n ($A_1A_2A_3B_3B_2B_1$), and a regular grammar $(AB)^n$ ($A_1B_1A_2B_2A_3B_3$). A and B syllables were distinguished in terms of phonological properties. They observed that the processing of the context-free grammars selectively activated Broca's area (BA44/45) which is typically involved in syntactic processing (Kaan & Swaab, 2002). On the other hand, processing of the regular grammars selectively activated the left frontal operculum. Friederici et al. (2006) tested subjects' ability to distinguish $A_3A_2A_1B_1B_2B_3$ sequences from sequences where an A syllable was replaced a B syllable, or vice versa. Subjects learned to reject sequences which did not conform to long-distance dependencies, for example, $A_1A_2A_3A_4B_2B_1$ and $A_1A_2A_3B_3B_2A_4$. A and B syllables were different in that they each consisted of different vowels. De Vries et al. (2008) argued that Friederici et al.'s task could be solved by matching the number of A and B syllables, and the knowledge of the long-distance dependencies was not required. Therefore, they claimed that the task was not sufficient to conclude that subjects internalize the NE rule.

To prove their point, De Vries et al. (2008) carried out their own experiment. They first trained all the participants on the same stimuli as Bahlmann and Friederici

(2006), and required them to judge the grammaticality of new items violating the NE rule. However, the participants were tested with different types of violations, namely: scrambled (e.g. $A_1A_2A_3B_1B_3B_2$) sequences and scrambled + repetition sequences ($A_1A_2A_3B_1B_2B_1$). As they had predicted, their participants could detect the scrambled + repetition violations, but not the scrambled ones. Therefore, de Vries et al. (2008) argued that Bahlmann and Friederici's subjects were using strategies such as counting or repetition monitoring rather than using the internalized knowledge of the NE rule to solve the task.

Uddén and her colleagues (2009) found that nested and crossing nonadjacent dependencies, which corresponded to context-free and context-sensitive grammars, respectively, could be learned by adult learners over the course of nine days. They used a between subject design to test the learning of the $A_1A_2A_3B_3B_2B_1$ pattern (NE) and the $A_1A_2A_3B_1B_2B_3$ pattern (CD) in two conditions. The training sessions spanned over 9 days, and each day, subjects were trained on the same learning set presented in a different order. The stimuli they used were neither exactly sentence nor words. They employed letters as the basic unit of their stimuli. The grammatical strings of letters had the following form: prefix-AAABBB-suffix. The prefix and suffix were represented by letters drawn from the set {M, N, S, V, W, R, X}, the As were drawn from {F, D}, and the Bs were drawn from {L, P}. Among the letters in sets A and B, the letters F and L exhibited a dependency, and D and P exhibited a dependency. In the NE condition, M-FFDPLL-VS was legal, and M-FDFPLL-VS was illegal. In the CD condition, an example legal string would be M-FFDLLP-VS, and an example of illegal string would be M-FDFLLP-VS. The hyphens indicating morpheme boundaries are written here just for clarity. They were not in the actual presentation of their stimuli. The stimuli in this design showed long-dependencies between letters in set A and letters in set B, and also avoided the confounds that subjects could learn such pattern by counting and/or keeping track of repetitions. These confounds were the basis for the criticisms that De Vries et al. pointed out for earlier studies.

During training, subjects were exposed to legal sequences of words that conformed to either an NE or a CD pattern depending on which condition they were assigned to. They were presented with the letter strings, which appeared on the computer, one at a time, and then asked to type out the string using a keyboard. Before the first training session started, they were tested on their preference for legal strings of words and illegal strings of words, which they had to indicate whether they 'like' or 'dislike'. This was used as a baseline for the experiment.

After nine days of training, subjects were presented with novel strings and were tested on both their preference and their acceptability judgment. They found that the subjects were more likely to prefer legal strings over illegal ones, and they also showed a higher rate for judging legal strings to be grammatical than for judging an illegal string to be grammatical. These results showed that subjects could learn both an NE and a CD grammar with this design.

Many studies have examined the learnability of context free grammars, but results are mixed. Some researchers have shown that adults can learn context free and context sensitive patterns using ALL (Fitch and Hauser 2004, Folia et al. 2008, Forkstam et al. 2008, Friederici et al. 2006, Bahlmann, Schubotz, and Friederici 2008, Petersson et al. 2004, Uddén et al. 2009), but others remain skeptical (Perruchet and Rey 2005, de Vries et al. 2008). The types of stimuli used in experiments vary from the media of transmission (aural vs. visual), to how the dependencies are instantiated (arbitrary syllable/letter cooccurrences vs. phonological properties). Among these studies, neuroimaging researches using fMRI or ERP were conducted to examine the brainwave activities when subjects are learning a regular pattern and a context free pattern. Interestingly, it has been shown by Bahlmann, Schubotz, and Friederici (2008) that when subjects were processing context free grammar, particular parts of the Broca's area (BA 44) were more activated than when they were processing regular grammar. However, the properties they employed to instantiate the dependencies between syllables was a phonological one. Their stimuli were constructed by a series of visually presented nonsense monosyllabic CV syllables, and the dependencies between syllables were realized by agreement in the place of articulation of the consonants. See (6) and (7) for examples:

(6) ge be di tu po ko (Context free- Nested Embedding)

(7) be pu gi ku de to (Regular- $(AB)^n$)

(Bahlmann et al. 2008)

Additionally, Bahlman et al. (2009) carried out another experiment to examine the processing of context free pattern in a non-language domain. They replaced the nonsense syllables with nonsense shapes, and they found that BA 44 was also activated, but the network it was involved in was different from the one for the language task. Their results are compatible with two theories: 1) the BA 44 region is responsible for processing any context free patterns (i.e. not specialized for language). This implies that any context free patterns can be processed by BA 44 regardless of what contexts the pattern is embedded in; 2) the BA 44 region belongs to a larger specialized network that is responsible for learning language, and this network overlaps with the spatial-visual network which is specialized for learning sequences of shapes. This is interesting because even though the stimuli used in the language experiment were considered as sentences, the dependencies were instantiated via phonological properties. Therefore, it is unclear whether subjects processed stimuli such as those in examples (6) and (7) as phonological stimuli or syntactic ones, and more importantly, whether our brains are activated in the same way when context free syntactic stimuli and context free phonotactic stimuli are being processed.

The results presented by the above-mentioned studies have been mixed and not as straight-forward as the results of phonotactic learning. Nonetheless, these results provide the groundwork for our research as we are interested in finding out whether these non-regular patterns are learnable in the context of phonology. Although there is some evidence to support that these patterns are learnable (Uddén et al. 2009), it is still unclear whether such patterns can be learned from auditory stimuli that more closely resemble actual words and sentences. In this dissertation, ALL is used to investigate the learnability of these non-regular patterns both as syntactic patterns and as sound patterns.

2.4.4 Summary of Remaining Questions

The phonological experiments in the previous literature have shown that patterns that were shown to be learnable in the ALL settings belong to either SL or SP classes, perhaps with the exception of Koo and Callahan's study, in which, the pattern they tested can be interpreted as a regular, but not SL, SP or TSL pattern. However, what has not been shown is the comparison of the learnability of a subregular pattern and a regular but non-subregular pattern. The question of whether a subregular pattern is more easily learned than a pattern that is regular, but non-subregular remains unanswered.

The syntax experiments which focused on the learnability of non-regular patterns have provided support that non-regular patterns can be learned using the ALL paradigm, but whether or not these patterns are also equally learnable in a phonotactic context remains an open question.

The experiments in this dissertation were designed to answer these questions.

Chapter 3

EXPERIMENT 1: CONTEXT FREE EXPERIMENT

The aim of this series of experiments deviates from the previous studies in the literature as its main interest is to detect whether the context free pattern is learnable within the phonological context. The domain specific hypothesis predicts that context free phonological pattern are not learnable, whereas, the domain general hypothesis predicts the contrary. However, it is impossible to provide empirical proof that a particular pattern is not learnable. For example, suppose the results of an ALL experiment indicate a pattern was not learned by its subjects. This null result is insufficient to prove the pattern is unlearnable since there might be another paradigm (say, one with a longer training time) that might give different results.

Therefore, this study instead tested a weaker version of the domain specific hypothesis: context free pattern is harder to learn in a phonological context than in a syntactic context.

In the next section, details of the experimental design, methodology and predictions of the results for both the syntactic and phonological experiments will be discussed.

3.1 General vs. Specific language learners

In order to test whether one general learner is responsible for learning both phonotactic and syntactic patterns, the learnability of a context free pattern (a pattern found in syntax but not phonology) needs to be tested in these two different linguistic domains. For a fair comparison, the stimuli and training amount for these two domains must be as minimally different as possible. Inevitably, there are some inherent difficulties as syntax and phonology are made up of different basic units, but as explained below, the abstract context free pattern that is embedded in a sentence and a word is essentially the same.

The abstract representation of NE is $A_1A_2B_2B_1$ and $A_1A_1B_1B_1$. In order to test whether subjects could learn NE, some counter-examples must be included in the test set in order to see whether they could be distinguished from the NE-conforming stimuli. These counter-examples include $A_1A_2B_1B_2$, which is a CD pattern, and $A_1A_1B_2B_2$ (termed SL3 henceforth). The term SL3 is used because this pattern can be defined by a strictly-local grammar of length 3. It should be noted that $A_1A_1B_1B_1$ also vacuously conforms to CD and SL3 as well as to NE. This particular type of stimuli is termed FULL in this dissertation. Table 1 below summarizes the types of stimuli used, and their abstract representations.

Table 1 Abstract representations of four types of stimuli.

	Grammatical according to NE		Ungrammatical according to NE	
Types	NE	Full	CD	SL3
Abstract representation	$A_1A_2B_2B_1$	$A_1A_1B_1B_1$	$A_1A_2B_1B_2$	$A_1A_1B_2B_2$

Subjects should be able to show different psychological reactions to NE and CD types of stimuli if they have successfully internalized the NE grammar. However, learning through a tier-based strictly 3-local grammar, which is a proper subset of regular grammar, would also enable the subjects to differentiate NE from CD stimuli if the subjects assumed the index numbers were projected to a separate tier. Table 2 below shows the legal strings of NE, Full, CD and SL3 grammars, and their respective subsequences of length 2 and 3.

Table 2 Legal strings and subsequences generated by the NE, CD, Full and SL3 grammars.

	Possible strings	Sub-sequences of length 2	Sub-sequences of length 3
FULL	1111 2222	{11 22}	{111 222}
NE	1111 2222 1221 2112	{12 21}	{122 221 211 112}
CD	1111 2222 1212 2121	{11 22 12 21}	{111 222 121 212}
SL3	1111 2222 1122 2211	{11 22 12 21}	{111 222 122 221 211 112}

As shown in the Table 2, the SL2 sequences of NE, CD and SL3 grammar are exactly the same, therefore it is insufficient to differentiate these grammars. The SL3 sequences, on the other hand, can differentiate between NE and CD, but not between NE and SL3. In order to rule out the possibility that subjects are internalizing NE through SL3, they must demonstrate a different psychological reaction to NE than CD and SL3 types of stimuli.

3.1.1 Instantiation of NE Syntactic Pattern

NE is instantiated in a sentence by means of the agreement of suffixes. Each sentence is made up of 4 pseudowords, and each bisyllabic root of a word is attached with a suffix. The long-distance dependency is displayed via the suffixes. Two words are considered to show agreement only when the same suffix is used. If a sentence conforms to NE, the suffixes of the 1st and 4th words must agree, and those of the 2nd and 3rd words must also agree. If a sentence conforms to CD, the suffixes of the 1st and 3rd words must agree, and those of the 2nd and 4th words must also agree. If a sentence conforms to SL3, the suffixes of the 1st and 2nd words must agree, and those of the 3rd and 4th words must also agree. Finally, for the Full type of sentences, the suffixes for all 4 words must agree. See Table 3 below for an example of each type of stimuli:

Table 3 Examples of syntactic stimuli.

Types	Examples
NE	(1) dahu-k gido-m badi-m buga-k
Full	(2) tido-s gipu-s topi-s bota-s
CD	(3) hage-m tiha-s dehi-m hetu-s
SL3	(4) dube-k bigi-k taga-m tabo-m

Subjects were not informed of the syntactic classes, meanings nor the morphonological boundaries of the pseudowords used. Therefore, the only cue they could pick up on if they assume any presence of dependency would be suffix agreement.

3.1.2 Instantiation of NE Phonotactic Pattern

NE is instantiated in a word by anteriority agreement in sibilants because the full type of sibilant harmony is attested in natural language phonology (Hansson 2001, Rose and Walker 2004). By using an attested type of long-distance dependency and modifying it to fit the NE grammar increases the chance for the pattern to be learned, as it more closely resembles attested phonotactic patterns.

Each word of stimuli consists of 4 sibilants in the form of $C_1V.C_2V.C_3VC_4$. For two sibilants to be considered harmonic, they must have the same value in [anterior]. If a word conforms to NE, its C_1 and C_4 have to be harmonic, and C_2 and C_3 also have to be harmonic. If a word conforms to CD, its C_1 and C_3 have to be harmonic and C_2 and C_4 also have to be harmonic. If a word conforms to SL3, its C_1 and C_2 have to be harmonic, and C_3 and C_4 also have to be harmonic. Finally, if a word conforms to the Full type, all four sibilants within the word must be harmonic. See Table 4 for examples of different types of stimuli:

Table 4 Examples of phonotactic stimuli.

Types	Examples
NE	(5) s a ʤ u ʃ e z [-ant] [+ant]
Full	(6) s e s i z u z [+ant] [+ant]
CD	(7) ʃ e s e ʃ e z [-ant] [+ant]
SL3	(8) z i s o ʃ u ʤ [+ant] [-ant]

As pointed out earlier, the actual realization of the syntactic and phonological stimuli have to be different due to their respective nature. However, the abstract representations of the 4 types of dependencies (NE, Full, CD, and SL3) are identical in both the phonotactic and syntactic contexts.

3.2 Hypothesis

The hypothesis of this study is the absence of context free patterns attested in syntax but not in phonology is due to its unlearnability. Two experiments were designed to investigate whether a context free pattern which is only found in natural language syntax is learnable in the context of phonology. The first experiment (1a) tested the learnability of a context free pattern realized over a sentence (syntax). This was done to ensure this pattern is learnable in experimental settings with ALL. The second experiment (1b) was carried out to test the learnability of the same pattern but realized over a word (phonology) using the same experimental paradigm.

3.2.1 Possible Outcomes

Since in both the syntactic and phonotactic conditions, subjects were trained on only Full and NE stimuli (which both conform NE grammar), and tested with NE-conforming stimuli (Full and NE) and NE-non-conforming stimuli (CD and SL3). If NE grammar was successfully learned during training, subjects's likeliness to consider Full and NE types of stimuli to be belong to the language should be higher than the preference of subjects in the control condition. The likeliness to consider CD and SL3 stimuli to be part of the NE grammar should be similar to the control group subjects' preference. There are 4 possible outcomes regarding whether the NE grammar is learned or not in the two conditions:

1. NE is learned in both Syntax and Phonotactic conditions.

2. NE is learned in neither Syntax and Phonotactic conditions.
3. NE is learned in only Syntax but not Phonotactic condition.
4. NE is learned in only Phonotactic but not Syntax condition.

Outcome (2) can neither reject nor support the hypothesis that context-free grammar is learned only in syntactic context but not phonotactic context as they are uninformative. Outcome (3) supports this hypothesis, whereas outcomes (1) and (4) would rule out the hypothesis.

Table 5 Some possible outcomes of Experiment 1a and 1b.

	Hypothesis supported		**Hypothesis rejected**			
	Outcome (3)		**Outcome (1)**		**Outcome (4)**	
Types	**Syntax condition**	**Syntax condition**	**Syntax condition**	**Syntax condition**	**Syntax condition**	**Phonotactic condition**
Full	> control	> control	> control	> control	~ control	> control
NE	> control	> control	> control	> control	~ control	> control
CD	~ control	~ control	~ control	~ control	~ control	~ control
SL3	~ control	~ control	~ control	~ control	~ control	~ control

Table 5 above shows outcomes that can reject or support the hypothesis that context-free grammars can be learned in syntactic context but not phonological context.

3.3 Experiment 1a: Syntax Experiment

3.3.1 Subjects

Fifty-eight native American-English monolingual speaking, normal hearing adults were recruited for the experiments. Subjects were students from the University of Delaware, aged between 18-27 years old, and their participation was compensated for either with course credit or $10.

3.3.2 Procedure

The experiment took place in a soundproof booth in the Phonetics and Phonology laboratory at the University of Delaware. The experiment consisted of 2 conditions (syntax and control). The procedure for the syntax condition consisted of 2 phases: a training phase and a testing phase. The total duration for both training and testing was about 45 minutes.

During the training phase, subjects listened to sentences that conformed to NE pattern and were instructed to repeat each sentence orally after it was presented. The training contained 200 tokens (40 sentences x 5 repetitions) and the duration was approximately 30 minutes. In the control condition, no training was given—subjects were only given the test.

Training was followed by a testing phase in which the subjects were presented with 100 novel sentences and were asked to judge whether each of them was likely to belong to the artificial language they heard during the training. The test took about 15 minutes to complete. All subjects, regardless of which condition they were in, were given the same test with the exact same 100 test items. However, in the control condition, subjects were asked to judge whether they liked the test sentences.

3.3.3 Stimuli

All training and testing items consisted of 4 bisyllabic words. Each word was constructed as root-suffix. All the roots were with the structure of CV.CV,CV, and contained only the consonants [p, b, t, d, g, h] and the vowels [a, e, i, o, u]. Three different suffixes were used: [-s], [-m], and [-k].

In the syntax condition, training items only included sentences that conformed to NE. Table 6 below summarizes the types of training and test stimuli used.

Table 6 Types of Training and Test Stimuli.

	NE	FULL	SL3	CD
Training	$A_1A_2B_2B_1$	$A_1A_1B_1B_1$	----------------	----------------
Test	$A_1A_2B_2B_1$	$A_1A_1B_1B_1$	$A_1A_1B_2B_2$	$A_1A_2B_1B_2$

The roots in category A were made up of [be, bi, da, do, du, ge, gi, ha, ho, hu, pa, po, pu, te, ti], and those in category B were made up of [ba, bo, bu, de, di, ga, go, gu, he, hi, pe, pi, ta, to, tu].

All stimuli were recorded by a native speaker of English with phonetic training who was unaware of the purpose of this study. She was instructed to produce all sentences with the same sentential intonation as if she would say the English well-formed NE sentence [the rat [the cat chased] escaped]. All the individual bisyllabic words were stressed on the first syllable. Moreover, the [e] and [o] vowels were pronounced diphthongized.

The training set consists of NE (50%) and Full (50%) types of sentences. The test sentences consist of two types of grammatical sentences (NE and Full), and two types of ungrammatical sentences (SL3 and CD). Each type of the sentences makes up of 25% of the entire set of test items.

3.3.4 Predictions

The subjects in the experimental condition should be more likely to consider NE and Full sentences to be part of the grammar during the training if they had successfully learned the NE grammar. The baselines were drawn from the responses of the subjects in the control condition, who did not receive any training prior to the test.

3.3.5 Results

Subjects' responses were collected with the E-Prime 2.0 software (Psychology Software Tools, Pittsburgh, PA) and were modeled using a linear mixed-effects model with a binomial function. The analysis with the binomial function was used because the distribution of the test results was binomial due to the nature of a yes/no task, therefore, the more traditional analyses using t-test or ANOVA which assume normally distributed data are inappropriate. The model was fitted in R (v.2.13.1) (R

Development Core Team 2009), using the *lmer*( ) function from the *lme4* package (Bates, Maechler, and Bolker 2011) for mixed-effects models. The model contained a fixed effect CONDITION with 2 levels (Control and Syntax), and two random effects: SUBJECT and TRIAL. The command used to run this analysis is shown below:

lmer (~YesRate + CONDITION + (1|Subject) + (1|Trial), family="binomial")

For each analysis, the Control condition was coded as the reference level, which was shown as the intercept in the output. With this set-up, the subjects' responses in each condition could be compared directly with those in the Control condition which served as the baseline.

The mean rates of 'yes' responses and their standard errors were tabulated in the Table 7 below.

Table 7 Descriptive statistics of the syntax experiment.

	Mean			
	NE	Full	CD	SL3
Control	0.6014	0.5289	0.6000	0.5862
Exp	0.5214	0.5338	0.5476	0.5462
	Standard Error			
	NE	Full	CD	NE
Control	0.0182	0.0185	0.0182	0.0182
Exp	0.0186	0.0185	0.0185	0.0186

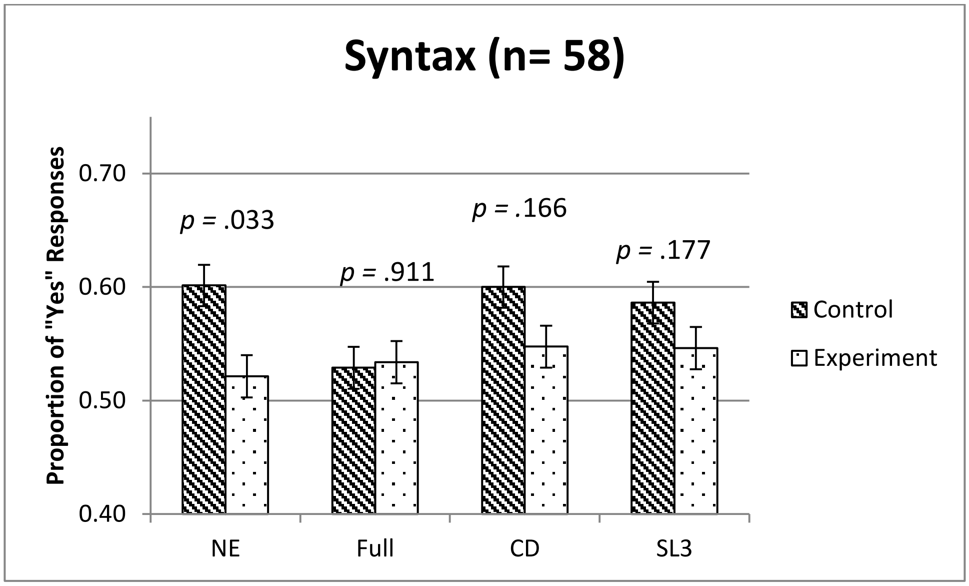

Figure 8 The mean rates of ‘yes’ responses in the syntax experiment.

First of all, the mean rates of experimental subjects’ ‘yes’ responses were lower than the control subjects’. This could mean that subjects had a novelty preference when they performed the test. Linear regression analyses reveal that subjects who were trained on NE grammar performed significantly different than the control subjects only in the NE type (see Table 8 below).

Table 8 Estimates of the control and experimental conditions of 4 types of test stimuli in the syntax experiment.

Types		Estimate	Std. Error	*z*	*p*
NE	(Intercept)	0.4406	0.1211	3.637	0.000276 ***
	Exp	-0.3513	0.1647	-2.133	0.032923*
Full	(Intercept)	0.12673	0.14138	0.896	0.370
	Exp	0.01892	0.16939	0.112	0.911
CD	(Intercept)	0.4364	0.1267	3.444	0.000573 ***
	Exp	-0.2340	0.1690	-1.385	0.166108
SL3	(Intercept)	0.3640	0.1127	3.229	0.00124 **
	Exp	-0.1700	0.1260	-1.349	0.17736

The analyses in Table 8 above were generated from 4 different models, in each of which, only the control and experimental conditions subjects' responses to one type of sentences were included. 'Intercept' corresponds to the control condition, as it was set as the reference level, and the p values for the 'Intercept' groups indicate whether it is significantly different than chance level (i.e. control vs. chance). The *p* values of the 'Exp' row represent whether or not the experimental data is significantly different from the control group (i.e. control vs. experimental). The p values in the (intercept) row are not particularly informative as control subjects preferred all types of sentences except for Full above chance level significantly. What is important is that whether experimental subjects, those who received NE training, behaved significantly different than the control subjects. As shown in Table 8, experimental subjects only reacted significantly differently than the control group when they were presented with NE sentences (p =.033). This suggests that if NE training is given, subjects would react psychological differently only when they were presented with NE type of sentences. The data was further analyzed by converting the binary responses to the rate response. There were two groups in the Experiment 1a: Syntax experimental versus Syntax control conditions. The dependent variable ("yes" rate) was on the interval scale of measurement. Therefore, data were analyzed using an independent samples *t*-test. Data were analyzed using SPSS. Results showed a statistically-significant difference with participants in the Syntax experimental condition obtaining lower "yes" rate when presented with NE type of stimuli than those in the Syntax control condition (t (56)= -2.059, p = .046). The differences between the two groups when presented with Full, CD and SL3 types of stimuli are statistically insignificant.

3.4 Experiment 1b: Phonotactic Experiment

3.4.1 Subjects

Fifty-two native American-English monolingual speaking, normal hearing adults were recruited for the experiments. Subjects were students from the University of Delaware, aged between 18-27 years old, and their participation was compensated for either with course credit or $10.

3.4.2 Procedure

The same procedure was employed as the syntax experiment.

3.4.3 Stimuli

All training and testing items consisted of trisyllabic words with the structure of CV.CV.CVC. The phonological stimuli were constructed parallel to the syntactic ones- each consonant corresponds to a nonsense word, and the [anterior] feature value of each corresponds to the suffix of each word in a sentence. All the roots were with the structure of CV.CV,CV, and contained only the consonants were all sibilants drawn from the set of [s, ʃ, z, ʤ] and the vowels [a, e, i, o, u]. Although [ʒ] would be more ideal instead of [ʤ] as it only differs in one feature compared to [z] in theory, its occurrence is relatively more restricted than [ʤ] in English. Therefore, [ʤ] was used instead of [ʒ].

In the phonotactic condition, training items only included words that conformed to NE grammar, and the dependencies between the sound segments were instantiated by the agreement of [anterior] feature of the sibilants within each word. Since the value for distinctive features are binary, '+' and '-' signs are used for the abstract representations of the stimuli (excluding vowels). Table 9 below summarizes the types of training and test stimuli used.

Table 9 Types of training and test stimuli used in the phonotactic experiment.

	NE	FULL	SL3	CD
Training	+ - - +/ - + + -	+ + + +/ - - - -	----------------	----------------
Test	+ - - +/ - + + -	+ + + +/ - - - -	+ + - -/ - - + +	+ - + -/ - + - +

All stimuli were recorded by a native speaker of English with phonetic training who was unaware of the purpose of this study. She was instructed to produce all words with a penultimate word stress. Moreover, the [e] and [o] vowels were pronounced diphthongized.

The training set consists of NE (50%) and Full (50%) types of sentences. The test sentences consist of two types of grammatical sentences (NE and Full), and two types of ungrammatical sentences (SL3 and CD). Each type of the sentences makes up of 25% of the entire set of test items. The distribution of types of stimuli is exactly the same as the syntax experiment.

3.5 Predictions

The subjects in the experimental condition should be more likely to consider NE and Full words to be part of the grammar they just during the training if they had successfully learned the NE grammar. The baselines were drawn from the responses of the subjects in the control condition, who did not receive any training prior to the test.

3.6 Results

Subjects' responses were collected and analyzed in the same way as the syntax experiment.

The mean rates of 'yes' responses and their standard errors were tabulated in Table 10 below.

Table 10 Descriptive statistics of the phonotactic experiment.

	Mean			
	NE	Full	CD	SL3
Control	0.4969	0.5477	0.5323	0.5215
Exp	0.5308	0.7031	0.5292	0.5785
	Standard Error			
	NE	Full	CD	NE
Control	0.0196	0.0195	0.0196	0.0196
Exp	0.0196	0.0179	0.0196	0.0196

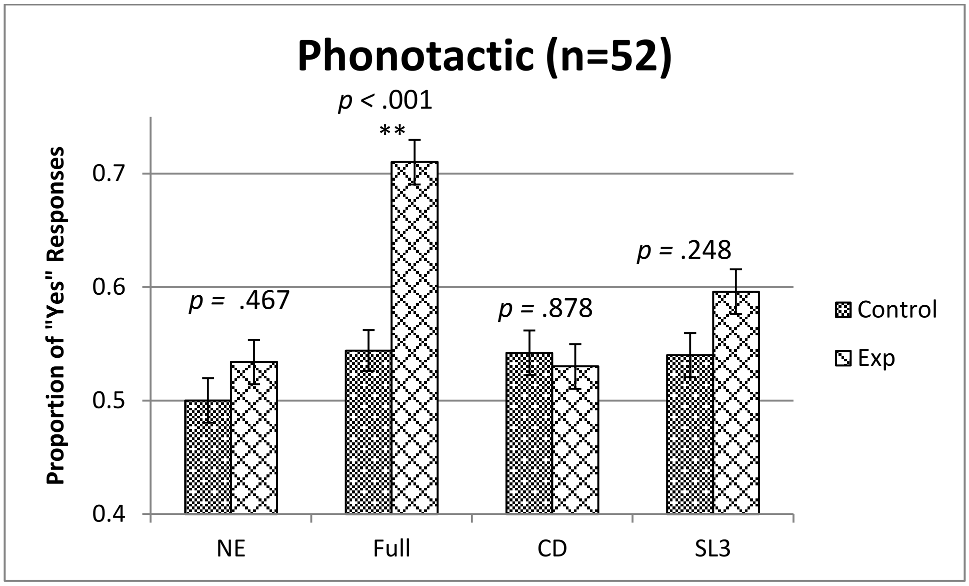

Figure 9 The mean rates of ‘yes’ responses across 4 types of test stimuli

The type of reaction observed in the phonotactic experiment differs from the syntax experiment. Subjects were more likely to say ‘yes’ to grammatical items rather than ungrammatical items during the testing phase. The rates of ‘yes’ responses of experimental subjects were higher than the control subjects. Notably, experimental subjects only reacted significantly different from the control group when presented with the Full type of stimuli (see Table 11 below).

Table 11 Estimates of control and experimental conditions of 4 types of test stimuli in the phonotactic experiment.

Types		Estimate	Std. Error	*z*	*p*
NE	(Intercept)	-0.009695	0.168960	-0.057	0.954
	Exp	0.150582	0.207130	0.727	0.467

Full	(Intercept)	0.2050	0.1462	1.402	0.161
	Exp	0.7274	0.1687	4.312	1.62e-05 ***
CD	(Intercept)	0.16529	0.17467	0.946	0.344
	Exp	-0.03314	0.21676	-0.153	0.878
SL3	(Intercept)	0.1069	0.1713	0.624	0.532
	Exp	0.2453	0.2123	1.155	0.248

The analyses reveal that experimental subjects only behaved significantly different from the control subjects when they were presented with Full type of words ($p < .001$).

The data was further analyzed by converting the binary responses to the rate response. There were two groups in the Experiment 1b: Phonotactic experimental versus Phonotactic control conditions. The dependent variable ("yes" rate) was on the interval scale of measurement. Therefore, data were analyzed using an independent samples *t*-test. Data were analyzed using SPSS. Results showed a statistically-significant difference with participants in the Phonotactic experimental condition obtaining higher "yes" rate when presented with Full type of stimuli than those in the Phonotactic control condition (t (38)= 3.753, p <.001). The differences between the two groups when presented with NE, CD and SL3 types of stimuli are statistically insignificant.

3.7 Discussion

The results obtained from the syntax experiment show that subjects who were exposed to NE grammar reacted significantly different from the control subjects, who had no training, only when they were presented with NE type of sentences. This suggests that the experimental subjects treated NE sentences differently from the Full, CD and SL3 sentences. However, if the NE grammar is successfully learned, the subjects should treat both NE and Full sentences differently from CD and SL3 sentences as Full sentences also conform to the NE grammar. A possible explanation for this is that the experimental subjects were only paying attention to sentences which actively displayed the NE type of dependencies, and ignored the Full sentences as they were ambiguous.

Remarkably, the results obtained from the phonotactic experiment show that experimental subjects reacted significantly differently from the control subjects only when they were presented with the Full type of words even though the training consisted of words which conform to NE and Full. This difference between the syntax and phonotactic learners is noteworthy because given the same type of training in terms of stimuli types, these two groups of experimental subjects appear to have chosen to internalize them differently.

Although both NE and Full strings can be generated from the NE grammar, a context free language, the Full strings themselves also conform to a regular grammar. Therefore, the Full strings are a proper subset of the NE strings. The reverse cannot be said of the NE strings as they can only be generated by NE grammar but not by a regular grammar, hence the NE strings are not a proper subset of the Full grammar. This fact is important for the interpretation of the results from both experiments, because this is precisely how the syntax and phonotactic learners' results diverge. Subjects in the syntax condition reacted significantly differently than the control subjects only when they were given NE sentences, and this suggests that they had internalized the NE grammar. Phontactic subjects reacted significantly than the control subjects only when they were presented with Full words, and this suggests that they had internalized a regular grammar, and not a context-free grammar. Furthermore, although both groups of experimental subjects seemed to have ignored part of the training stimuli, the Full sentences which were overlooked by the subjects in the syntax consition do not contradict the NE rule which was internalized. On the contrary, the NE words which were ignored by the phonotactic subjects are counter-examples of the regular grammar, which they had internalized. Due to this difference, the driving forces for disregarding a particular type of training input must be qualitatively different for the syntax and phonotactic subjects- syntax learners disregarded the Full sentences because they were not informative enough w.r.t. the dependencies they had to learn; phonotactic learners disregarded the NE words because they failed to process the more informative stimuli.

If in fact, the syntax and phonotactic subjects internalized different grammars when both groups were given the same NE grammar in training supports the domain specific hypothesis- humans utilize two different learning mechanisms when they are learning with the same type of language pattern embedded in two different contexts.

This is endorsed by the assumption that the phonological learning mechanism cannot compute grammars that are more complex than regular language, while the syntax mechanism should be able to process context free grammars. If the domain general hypothesis was correct, the NE pattern would have been learned in the phonotactic experiment at least in the same way as it was learned in the syntax experiment.

As mentioned before, the regular Full grammar instantiated by sibilant harmony is an attested pattern in natural languages' phonology, but the NE grammar is unattested. The results from the phonotactic experiment align with the typology, and support the claim that the absence of NE phonotactic patterns in natural languages is not incidental. If the sibilant NE pattern examined in this experiment could in fact be generalized to all types of NE phonotactic patterns, it could be the case that the absence of them is due to the inability of the phonological mechanism to process such grammars.

3.8 Conclusions

This set of experiments is the first experimental study to test the hypothesis of whether there is a distinct learning module for phonology that is highly constrained with computationally well-defined boundaries. The domain specific hypothesis speculates that there are two distinct learning mechanisms for syntax and phonology, and they differ in at least their computational complexities. The syntax learner is hypothesized to be able to learn at least context free grammars, whereas the phonological learner could not. This hypothesis is supported by the results of this study, as they have shown that the NE pattern was learned more readily in the syntax context, than the phonotactic context. Additionally, the phonotactic subjects showed a strong bias towards learning a regular pattern. This regular bias aligns with the phonotactic typology as virtually all attested phonotactic patterns are regular, and suggests that the lack of NE patterns in natural language is due to the limitations of the phonological learner.

If there are two distinct learning mechanisms for two different aspects within language, the domain general theory which supposes the presence of one general learner for different cognitive aspects should be refuted. At the very least, the resources that enable us to learn a context free pattern in syntax are not shared by the phonological domain. This relates to the fMRI findings by Bahlmann et al. (2009) which suggest that the Broca's area (BA 44) are more activated when context-free patterns are being processed in both language and non-language contexts. If the BA 44 region is a domain general resource, it is predicted to be equally activated when syntactic context free and phonotactic context free stimuli are being processed. But based on the present findings, it can be conjectured that BA 44 may not be a domain general resource, or even if it is, the network which is responsible for processing phonotactic patterns excludes it. Obviously, these are just speculations, and should be verified by future work using neuroimaging technology.

To conclude, the present study supports the H2 and H4, both of which assume a distinct learning mechanism for phonology, and H1 and H3 are rejected. However, it is still unclear whether the phonological learner is restricted by the subregular boundary.

Chapter 4

EXPERIMENT 2: SUBREGULAR EXPERIMENTS

4.1 Is the Subregular boundary psychologically real?

Heinz's (2010) *Subregular Hypothesis* states that only phonotactic patterns belonging to certain subregular classes are learnable. In order to determine whether our phonological learner is subjected to computational restrictions, specifically those identified by Heinz (2010), the learnability of regular, non subregular pattern is compared with that of a subregular pattern. If the subregular boundary is not psychologically real, there should be no reason why the learnability of these two patterns would be different.

To date, there are no studies that directly compare the learnability of patterns that belong to the subregular classes and one that does not. Such a comparison can generate 4 logically possible outcomes: 1) both patterns are learnable; 2) SH is learnable, while FL is not; 3) neither pattern is learnable; 4) SH is unlearnable, while FL is. The possible outcomes are summarized in Table 12.

Table 12 Logically possible experimental outcomes that could be obtained from comparing the learnability of a subregular pattern and a non-subregular pattern.

Paradigms	**Subregular (SL/SP/TSL)**	**Non-subregular (Non-SL/-SP/-TSL)**
1	Learnable	Learnable
2	Learnable	Unlearnable
3	Unlearnable	Unlearnable
4	Unlearnable	Learnable

All of these scenarios except for the 4th one are compatible with the Subregular Hypothesis. Therefore, just demonstrating that a non-subregular pattern is learnable under some artificial conditions is not sufficient to reject this hypothesis. However, the creation of an experimental paradigm which produces outcome 4 would falsify the Subregular Hypothesis. Additionally, showing that both subregular and non-subregular patterns are learnable or unlearnable is not particularly informative. A result of the 2nd outcome can be interpreted as evidence in favor of the Subregular Hypothesis.

In order to compare the learnability of two patterns, the patterns must be as minimally different as possible, and the paradigm must give equal training in word kind to each pattern. As explained below, FL and SH are well-matched in many respects. The decision to test the learnability of FL is not arbitrary. This pattern was chosen not only because it is a regular, but non-subregular pattern, but also because it is very similar to SH, an attested pattern. Computationally, the required memory required is the same; only the pattern template is different (see Table 14). This allows for a fair comparison of the learnability of these two patterns. Additionally, the first and last positions of a word are both privileged in terms of saliency and are relevant in phonology. Finally, in C'Lela, a Niger-Congo language, there is an attested pattern that resembles FL, which could plausibly be interpreted as evidence for its learnability. All these properties of FL make it a good candidate for evaluating the Subregular Hypothesis.

4.1.1 FL and SH patterns

One example of a regular sound pattern that is not found in any natural language is long-distance assimilation between only the first and last sounds of a word. Unlike the well-documented long-distance harmony patterns (Hansson 2001, Rose and Walker 2004), FL allows disharmonic intervening segments so long as the first and last sound are harmonic.

The comparison with sibilant harmony, which is documented in Navajo (Sapir and Hoijer 1967), is instructive. Navajo requires sibilants in well-formed words to agree in anteriority. Hypothetical words such as [sototos] and [ʃototoʃ] are both grammatical as the two sibilants in each word agree in anteriority, but [ʃototos] and [sototoʃ] are ill-formed because the two sibilants disagree in anteriority. On the other hand, FL permits both [sototos] and [ʃototoʃ], because the sibilants in the initial and last positions agree in anteriority. Since [ʃototos] and [sototoʃ] do not meet this requirement, they are ill-formed according to FL.

The difference between SH and FL becomes more apparent when examples with sibilants in word-medial positions are examined. FL predicts that [soʃotos] is well-formed because the first and the last sibilants are harmonic. According to SH, on the other hand, [soʃotos] is ill-formed because the word-medial sibilant disagrees with the others. Table 13 summarizes these examples. Note that all words that are well-formed according to SH are also well-formed according to FL (i.e. SH-acceptable words are a proper subset of FL-acceptable words).

Table 13 Examples of legal and illegal strings according to FL and SH grammars. Ellipsis is used to show that these sound segments are not necessarily adjacent to each other.

	Strings
Well-formed according to both SH and FL	[s...s...s]*, [ʃ...ʃ...ʃ]
Well-formed according to SH but not FL	[s...ʃ...s], [ʃ...s...ʃ]
Ill-formed according to both SH and FL	[ʃ...ʃ...s], [s...ʃ...ʃ]
Well-formed according to SH but not FL	None

Computational analysis of these patterns reveals that SH is SP (Heinz 2010), but FL, on the other hand, is neither SL, TSL nor SP. First-last belongs to the Locally Testable class of the Subregular Hierarchy (Rogers and Pullum 2010). This class is a superset of SL, and is a proper subset of regular language.

The learnability of FL can be assessed by comparing its learnability to that of SH. SH is only minimally different from FL as both rules state that [s] can be followed by [s] but not [ʃ], and [ʃ] can be followed by [ʃ] but not [s]. The only difference is the environments of these restrictions, as shown in Table 14:

Table 14 The cooccurrence of sibilants in SH and FL.

x…y	[s]	[ʃ]
[s]	✓	×
[ʃ]	×	✓

x…y	[s]	[ʃ]
[s]	✓	×
[ʃ]	×	✓

From a linguistic and cognitive perspective, FL seems plausible not only because long-distance dependencies between sounds are attested in natural language, but also because word edges have special status in phonology (Beckman 1998, Endress, Nespor and Mehler 2009). Sounds at these positions are usually more perceptually salient, and some phonological rules are edge-sensitive. In this light, FL is not that strange of a pattern.

4.1.1.1 C'Lela

Another reason to think that FL is not a strange pattern is that it is very similar to an attested pattern: a vowel harmony pattern in C'Lela (Dettweiler 2000, Pulleyblank 2002, Archangeli and Pulleyblank 2007). C'Lela is a Niger-Congo language, spoken in Nigeria. The direct object 1st person pronoun [-mi]/[-me] alternates depending on the vowel height of the root. If the vowel in the root is high, the suffix [-mi] surfaces, as in (1). If the root contains a non-high vowel, [-me] surfaces, as in (2).

(9) [buzəkə-mi] 'chased me'
(10) [ɛpkə-me] 'bit me'

(Archangeli and Pulleyblank 2007:8)

C'Lela allows suffix stacking, and interestingly, if there is more than one suffix attached to a root, only the final suffix assimilates to the vowel in the root. The word-medial suffix becomes transparent. Consider the following examples in Table 15.

Table 15 Examples of C'lela (Archangeli and Pulleyblank 2008).

High root with single suffix	High root with two suffixes
(11) i-z**i**s-**i** 'CM-long-CM'	(13) i-z**i**s-i-n**i** 'CM-long-CM-ADJM'
(12) u-p**u**s-**u**'CM-white-CM'	(14) u-p**u**s-u-n**i** 'CM-white-CM-ADJM'
Non-high root with single suffix	Non-high root with two suffixes
(15) i-r**e**k-**e** 'CM-small-CM'	(17) i-r**e**k-i-n**e** 'CM-small-CM-ADJM'
(16) u-g^j**ɔ**z-**o** 'CM-red-CM'	(18) u-g^j**ɔ**z-u-n**e** 'CM-red-CM-ADJM'

The vowels in examples (13) and (14) all agree in height. This is also the case for non-high roots, as in examples (15) and (16). However, when an additional suffix is attached to stems (15) and (16), medial suffixes surface as [-i] and [-u], as in (17) and (18), respectively. The newly added final suffixes still surface as non-high vowels, and therefore are harmonic to the vowel in the root.

One interpretation of the above data is that it is an edge-sensitive vowel harmony pattern. However, it should be noted that prefixes do not seem to participate in the vowel harmony process in C'Lela as shown in examples (15)-(18). In addition, examples with multi-syllabic roots are limited, and therefore no example of a root with vowels of different height (with the exception of [-ə], analyzed as a nonphonemic featureless mora (Pulleyblank 2002:260)) was found. Based on these examples, one can only conclude that the trigger of the vowel harmony in C'Lela is morphologically-bound – the vowel is in a root – and the target is position-bound – the final suffix. This is different from FL assimilation, in which both the trigger and the target are position-bound. Therefore, the vowel harmony pattern of C'Lela, while suggestive of FL, is not exactly the same as FL.

4.2 Hypothesis

The hypothesis of this study is that the absence of certain types of phonological patterns in the word's languages is due to the limitations on what can be extrapolated from the speech input by the phonological learner. This hypothesis was tested empirically in two ALL experiments.

4.3 Methodology

4.3.1 Subjects

Sixty-six native American-English monolingual speaking adults were recruited for the experiments. Subjects were students from the University of Delaware, aged between 18-27 years old, and their participation was compensated for either with course credit or $10.

4.3.2 Procedure

The experiment took place in a soundproof booth in the Phonetics and Phonology laboratory at the University of Delaware. The experiment consisted of 2 experimental conditions (SH and FL) and a control condition. The procedure for both conditions consisted of 2 phases: a training phase and a testing phase. The total duration for both training and testing was about 30 minutes.

During the training phase, subjects listened to words that conformed either to an SH or FL grammar (depending on the experimental condition) and were instructed to repeat each word orally after it was presented. The training contained 200 tokens (40 words x 5 repetitions) and the duration was approximately 15 minutes. In the control condition, no training was given—subjects were only given the test.

Training was followed by a testing phase in which the subjects were presented with pairs of words and were asked to judge whether the first word or the second word of the pair was more likely to belong to the artificial language they just heard during the training. There were 48 pairs of test items in total, and the test took about 10 minutes to complete. All subjects, regardless of which condition they were in, were given the same test with the exact same pairs of 48 pairs of test items.

4.3.3 Stimuli

All training and testing items were trisyllabic, with the structure of CV.CV.CVC, and contained only the consonants [k, s, ʃ] and the vowels [a, ɛ, i, ɔ, u]. Half of the training items consisted of a stop as the second consonant, and the remaining half consisted of a stop as the third consonant. The first and the last consonant were always sibilants.

In the SH condition, training items only included words that conformed to SH and similarly, in the FL condition, training items only included words that conformed to FL were included. Table 16 below summarizes the types of training stimuli used in the three conditions. A complete list of stimuli can be found in the Appendix.

Table 16 Types of training items used in SH, FL and control conditions. Vowels are omitted.

Conditions / Sibilant Tier	SH	FL	Control
[s…s…s]	[s…k…s…s] [s…s…k…s]	[s…k…s…s] [s…s…k…s]	No Training
[ʃ…ʃ…ʃ]	[ʃ…k…ʃ…ʃ] [ʃ…ʃ…k…ʃ]	[ʃ…k…ʃ…ʃ] [ʃ…ʃ…k…ʃ]	
[s…ʃ…s]	None	[s…k…ʃ…s] [s…ʃ…k…s]	

[ʃ…s…ʃ]	None	[ʃ…k…s…ʃ] [ʃ… s…k…ʃ]

One third of the test items contained disagreeing sibilants as the first and the last consonants (e.g. [s…s…ʃ]), and these words conformed to neither the SH nor the FL grammar (i.e. they are examples of *FL/*SH). Another third of the test items contained agreeing sibilants throughout the word (e.g. [s…s…s]), and these words conformed to both SH and FL (FL/SH). Lastly, the remaining one third contained agreeing sibilants only as the first and last consonants (e.g. [s…ʃ…s]), and these words only conformed to FL (FL/*SH). The fourth logically possible type, words that conformed to SH but not FL (*FL/SH) was not used, because all stimuli that conform to SH must also conform to FL.

A two-alternative forced choice design was used; the three types of test stimuli pitted against each other and generated three types of pairings:

a) FL/*SH vs. *FL/*SH (also includes *FL/*SH vs. FL/*SH)

b) FL/SH vs. *FL/*SH (also includes *FL/*SH vs. FL/SH)

c) FL/*SH vs. FL/SH (also includes FL/SH vs. FL/*SH)

4.1.4 Stimuli Recording. Natural stimuli were used for the experiments. A Mandarin Chinese-speaking graduate student with phonetic training who was unaware of the experiments' purpose was recruited to record the stimuli. Explicit training was given to the recorder to ensure that all stimuli were produced consistently. All vowels were pronounced as full vowels. Word stress (with the acoustic correlates of increased pitch and loudness) was placed on the penultimate syllable of all words, and the sibilant [ʃ] was pronounced with rounded lips.

4.4 Predictions

The experiment was designed to investigate whether the choice made by subjects was influenced by the type of grammar they were given in training. Table 17 summarizes the predicted responses if SH and FL are successfully learned in the respective conditions.

Table 17 Predicted preferences for each test pairing if SH and FL grammars were internalized.

Conditions \ Pairs	**FL/*SH vs. *FL/*SH**	**FL/SH vs. *FL/*SH**	**FL/SH vs. FL/*SH**
SH	No preference	FL/SH	FL/SH
FL	FL/*SH	FL/SH	No preference
Control	No preference	No preference	No preference

The results from both the SH and FL groups were compared to those of the Control group – assuming the control group should have no preference for either item in each pairing (since no training was given), the predicted results for each experimental group if they successfully internalized the grammar they were exposed to are shown in Table 18.

Table 18 Predicted results w.r.t.Control group for each test pairing if SH and FL grammars were internalized.

Pairs / **Conditions**	FL/*SH vs. *FL/*SH Rate of FL/*SH	FL/SH vs. *FL/*SH Rate of FL/SH	FL/SH vs. FL/*SH Rate of FL/SH
SH	~ Control	~ Control	> Control
FL	> Control	> Control	~ Control

4.5 Results

The descriptive statistics for the rates of choosing FL/*SH and FL/SH in all three types of test pairings are summarized in Table 19.

Table 19 Descriptive statistics of Control, SH and FL Conditions.

	Conditions		
	Control	**SH**	**FL**
FL/*SH vs. *FL/*SH Mean Rate of FL/*SH *(SE)*	0.49 (*0.03*)	0.49 (*0.03*)	0.52 (*0.03*)
FL/SH vs. *FL/*SH Mean Rate of FL/SH *(SE)*	0.48 (*0.03*)	0.62 (*0.03*)	0.63 (*0.03*)
FL/SH vs. FL/*SH Mean Rate of FL/*SH *(SE)*	0.45 (*0.03*)	0.56 (*0.03*)	0.58 (*0.03*)

Subjects' responses were collected with the E-Prime 2.0 software (Psychology Software Tools, Pittsburgh, PA) and were modeled using a linear mixed-effects model with a binomial function. The analysis with the binomial function was used because the distribution of the test results was binomial due to the nature of a two-alternative forced choice task, therefore, the more traditional analyses using t-test or ANOVA which assume normally distributed data are inappropriate. The model was fitted in R (v.2.13.1) (R Development Core Team 2009), using the lmer( ) function from the lme4 package (Bates, Maechler, and Bolker 2011) for mixed-effects models. The model contained a fixed effect CONDITION with 3 levels (Control, SH, and FL), and two random effects: SUBJECT and TRIAL. The commands used for this analysis is shown below:

lmer (~FLresponse + CONDITION + (1|Subject) + (1|Trial), family= "binomial")

lmer (~FL/SHresponse + CONDITION + (1|Subject) + (1|Trial), family= "binomial")

For each analysis, the Control condition was coded as the reference level, which was shown as the intercept in the output. With this set-up, the subjects' responses in each condition could be compared directly with those in the Control condition. Each model was compared with the empty model, where the fixed effect was replaced by "1". The function anova( ) was used to perform a likelihood ratio test between the empty model and the respective individual model to check if CONDITION was an important factor in its own right in each model.

The results for each type of pairing were analyzed separately because each pairing had a different dependent variable. The results were analyzed by examining the rate of choosing one type of stimuli over the other within a pairing. For example, in pairing (a), where FL/*SH was pitted against *FL/*SH, the rate of choosing FL/*SH was analyzed: the subjects' response was coded as "1" if they chose the FL/*SH item and "0" otherwise. For pairings (b) and (c), the subjects' response was coded as "1" if they chose the item that conformed to FL/SH and "0" otherwise.

In the analysis, the one-tailed test for cases in which the results were expected to be "Higher than Control" was used. For cases which "Same as Control" was predicted, the two-tailed test was used.

The rates of choosing FL/*SH when subjects were presented with FL/*SH vs. *FL/*SH in all three conditions are shown in Figure 10.

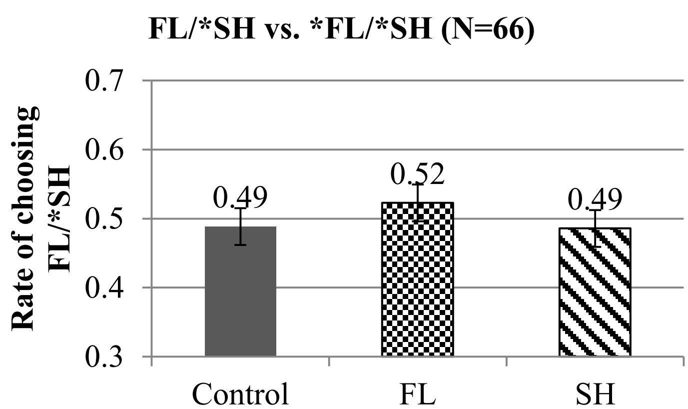

Figure 10 Mean rates of choosing FL/*SH when presented with FL/*SH vs. *FL/*SH.

The likelihood ratio tests showed that only 2 out of 3 models with the fixed factor CONDITION were significantly different from their respective empty models. The first model FL/*SH vs. *FL/*SH was not significantly different from its empty model ($\chi 2 = 1.05$, $p = .59$), which means CONDITION is not an important predictor in this model. The second model, FL/SH vs. *FL/*SH and third model FL/SH vs. FL/*SH were both significantly different from their empty models ($\chi 2 = 14.22$, $p < .001$ and $\chi 2 = 10.71$, $p =.005$) respectively. This means CONDITION is an important factor in its own right in both models.

The model for the FL/*SH vs. *FL/*SH pairings showed that neither the SH nor the FL groups' responses were significantly different from the Control group's (shown as Intercept in Table 20). The log-odds of the SH subjects choosing FL/*SH was not significantly higher than the Control subjects' ($p_{1\text{-tailed}} = .47$), nor was the log-odds of the FL subjects choosing FL/*SH significantly different from the Control subjects' ($p_{2\text{-tailed}} = .20$).

Table 20 Estimates of the conditions in the analysis of subjects' response in pairing FL/*SH vs. *FL/*SH.

*FL/*SH vs. *FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**
(Intercept)	-0.04638	0.14192	-0.327	0.744	0.372
Condition: SH	-0.0119	0.16435	-0.072	0.942	0.471

Condition: FL	0.14139	0.16439	0.860	0.390	0.195

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

The rates of choosing FL/SH when subjects were presented with FL/SH vs. *FL/*SH pairings in all three conditions are shown in Figure 3.

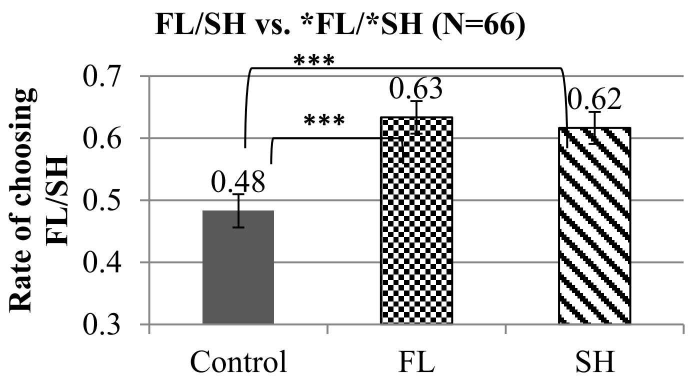

Figure 11 Mean rates of choosing FL/SH when presented with78 FL/SH vs. *FL/*SH.

The model for the FL/SH vs. *FL/*SH pairings suggests that the log-odds of the SH subjects choosing FL/SH was significantly higher than the Control subjects’ (p1-tailed < .001), and the log-odds of the FL subjects choosing FL/SH was also significantly higher than the Control subjects’ (p1-tailed <.001).

Table 21 Estimates of the conditions in the analysis of subjects’ response in pairing FL/SH vs. *FL/*SH.

*FL/SH vs. *FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**

(Intercept)	-0.07203	0.16927	-0.426	0.670436	0.335218
Condition: SH	0.58131	0.18073	3.216	0.001298**	0.000649***
Condition: FL	0.6581	0.18134	3.629	0.000284***	0.000142***

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

The rates of choosing FL/SH when subjects were presented with the FL/SH vs. FL/*SH pairings in all three conditions are shown in Figure 4.

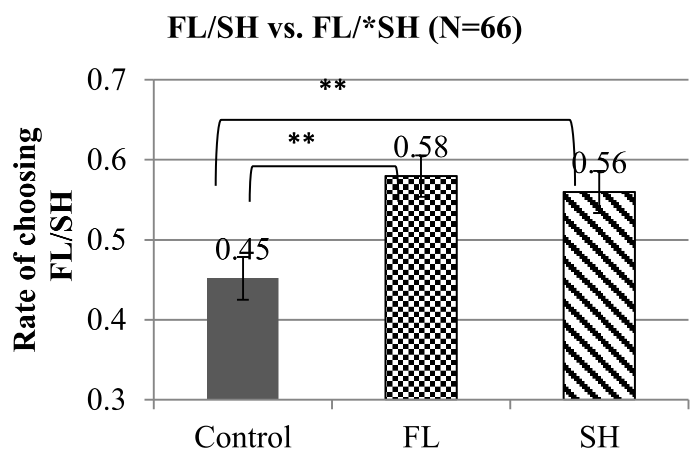

Figure 12 Mean rates of choosing FL/SH when presented with FL/SH vs. FL/*SH.

The model for the third test pairing, FL/SH vs. FL/*SH, suggests that the log-odds of the SH subjects choosing FL/SH was significantly higher than the Control subjects’ (p1-tailed = .004), and the log-odds of the FL subjects choosing FL/SH was significantly higher than the Control subjects’ (p2-tailed =.002).

Table 22 Estimates of the conditions in the analysis of subjects' response in pairing FL/SH vs. FL/*SH.

*FL/SH vs. FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**
(Intercept)	-0.2037	0.1526	-1.335	0.18196	0.09098.
Condition: SH	0.4544	0.1697	2.678	0.0074**	0.0037**
Condition: FL	0.5394	0.1700	3.172	0.00151**	0.000755***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

A separate analysis was run to test whether the FL group performed significantly different from the SH group. The SH group was coded as the reference group (Intercept) in this analysis, and since no specific direction was predicted for the results, 2-tailed tests were used.

Table 23 Estimates of the conditions in 3 types of test pairings with SH group as the reference group.

*FL/SH vs. FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**
(Intercept)	-0.05818	0.14195	-0.41	0.682
Condition:Control	0.01182	0.16435	0.072	0.943
Condition: FL	0.15308	0.16441	0.931	0.352
*FL/SH vs. FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**
(Intercept)	0.50927	0.17142	2.971	0.00297**
Condition:Control	-0.58127	0.18073	-3.216	0.0013**
Condition: FL	0.07667	0.18317	0.419	0.67552
*FL/SH vs. FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**
(Intercept)	0.25095	0.15285	1.642	0.10063
Condition:Control	-0.45467	0.16968	-2.68	0.00737**
Condition: FL	0.08469	0.17022	0.498	0.61881

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

As shown in Table 23, the responses of the FL group in all three types of test pairings are not significantly different from the SH group: FL/*SH vs. *FL/*SH (p = .352), FL/SH vs. *FL/*SH (p = .676) and FL/SH vs. FL/*SH (p = .619).

The results obtained match the predictions made by the SH but not the FL grammar. Therefore, it can be concluded that the SH subjects were able to internalize the SH grammar.

The data was further analyzed by converting the binary responses to the rate of response. There were three groups in the experiment, each representing a training conditioning: Control, SH, and FL. There are three dependent variables- 1) the rate of choosing FL/*SH item when presented with FL/*SH and *FL/*SH , 2) the rate of choosing FL/SH item when presented with FL/SH and *FL/*SH, and 3) the rate of choosing FL/SH when presented with FL/*SH and FL/SH. Data were analyzed using a one-way analysis of variance (ANOVA).

Preliminary comparisons revealed that the homogeneity assumption underlying an ANOVA was met. Therefore, post hoc comparisons were apportioned using the Tukey adjustment. The ANOVA with DV(1) the rate of choosing FL/*SH item when presented with FL/*SH and *FL/*SH did not show a statistically significant difference between groups ($F(2, 60) = .296$, $p = .745$). The ANOVA with DV (2) the rate of choosing FL/SH item when presented with FL/SH and *FL/*SH showed a statistically significant difference between groups ($F(2, 60) = 7.994$, $p < .001$). The ANOVA with DV (3) the rate of choosing FL/SH item when presented with FL/SH and FL/*SH also showed a statistically significant difference between groups ($F(2, 60) = 4.718$, $p = .013$).

Post hoc analyses demonstrated that both SH condition and FL condition produced significantly higher rates of choosing FL/SH when presented with FL/SH vs. *FL/*SH than the Control condition ($p = .006$) and ($p < .001$) respectively. Furthermore, both SH condition and FL condition produced significantly higher rates of choosing FL/SH when presented with FL/SH vs. FL/*SH than the Control condition (p = .047) and (p = .017) respectively.

The results for the FL condition were unexpected. Not only was FL not learned by the FL subjects, but their performance was not significantly different from that of the SH subjects in all three test pairings. When they were given the pairing of *FL/*SH vs. FL/*SH, they did not perform significantly different from Control. If FL had been learned successfully, the rate of choosing FL/*SH should be higher than Control. For the second type of pairing, the FL subjects' rate of choosing FL/SH when they were given the choice of *FL/*SH vs. FL/SH was significantly higher than Control, a choice consistent with both the FL and the SH grammars. Lastly, the rate of choosing FL/SH when the subjects were given the choice of FL/*SH vs. FL/SH was also significantly higher than Control – here both items conform to FL, but the subjects showed a preference for the item that also conforms to SH. In sum, the FL subjects chose the items that conform to SH significantly more than the items that do not, but failed to choose items that only conform to FL. Combining the results from all three pairings, it was concluded that the FL subjects were unable to internalize the FL grammar.

In addition, the FL subjects were not expected to internalize the SH grammar, because the FL training included items that do not conform to SH (e.g. [s...ʃ…s] and [ʃ…s…ʃ]). Yet the FL subjects seemed to ignore these words, and chose to accept the SH grammar anyway. It could be the case that the subjects were heavily biased towards learning SH, and the presence of stimuli that conform to both SH and FL led them to falsely assume the SH grammar. Thus, as a follow-up, an additional experiment was conducted to alleviate this potential SH bias by replacing the ambiguous FL/SH words with words that conform to only FL and not to SH (i.e. FL/*SH).

4.6 Additional Condition- Intensive FL (IFL)

Another 22 monolingual English speakers were recruited for this condition. The Intensive FL training stimuli were constructed similarly to the FL stimuli in terms of length, syllable structure, and the phoneme inventory used. Words that conform to both SH and FL, e.g. [s…s…s] and [ʃ…ʃ…ʃ] were replaced by words that only conform to FL e.g. [s…ʃ…s] and [ʃ…s…ʃ]. Instead of 4 types of training stimuli, only 2 were used. The test used in Experiment 1 was used in this condition. The procedure was the same as SH and FL conditions.

The results of the IFL condition were significantly different from the FL condition.

Table 24 Descriptive statistics of IFL condition.

IFL	FL/*SH vs. *FL/*SH Rate of choosing FL/*SH	FL/SH vs. *FL/*SH Rate of choosing FL/SH	FL/SH vs. FL/*SH Rate of choosing FL/SH
Mean	0.553977	0.414773	0.357955
SE	0.026532	0.026297	0.025588

The rates of choosing FL/*SH when subjects were presented with the pair of FL/*SH vs. *FL/*SH in the IFL condition are shown in Figure 13, the rates of choosing FL/SH when subjects were presented with the pair of FL/SH vs. *FL/*SH are shown in Figure 14, and the rates of choosing FL/SH when subjects were presented with the pair of FL/SH vs. FL/*SH are shown in Figure 15.

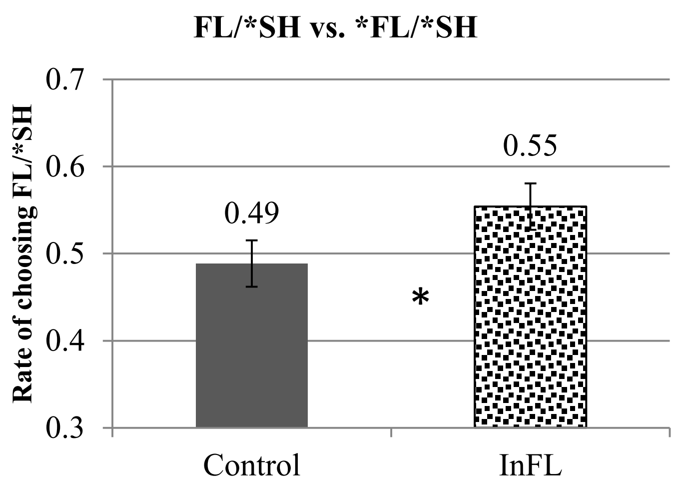

Figure 13 The rates of choosing FL/*SH when subjects were presented with the pair of FL/*SH vs. *FL/*SH in the IFL condition.

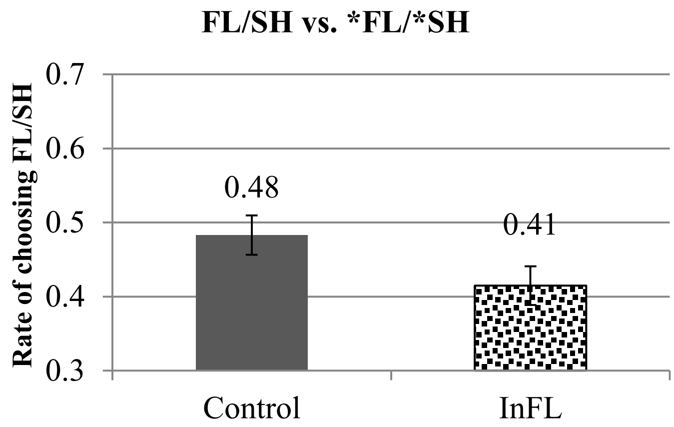

Figure 14 The rates of choosing FL/SH when subjects were presented with the pair of FL/SH vs. *FL/*SH in the IFL condition.

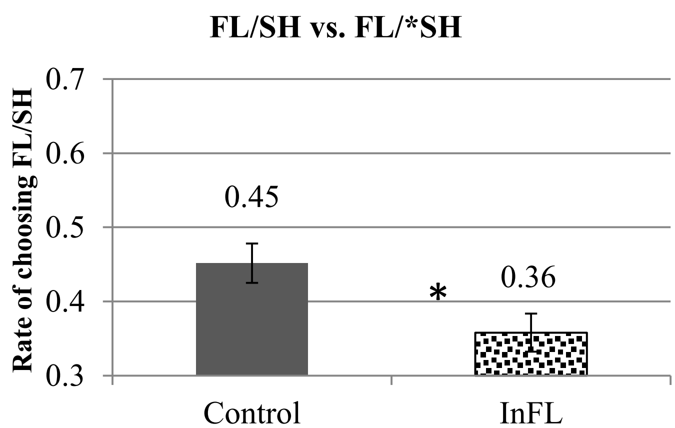

Figure 15 The rates of choosing FL/SH when subjects were presented with the pair of FL/SH vs. FL/*SH in the IFL condition.

The data obtained from the IFL condition was added to the original dataset, and the whole dataset was regressed again with the same model. The fixed effect now consists of 4 levels: Control, SH, FL and IFL. The estimates, *z* values and *p* values of the three models (one for each type of test pairing) are included in Table 25.

Table 25 Estimated of the Control, SH, FL and IFL conditions.

*FL/*SH vs. *FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**
(Intercept)	-0.04709	0.13903	-0.339	0.7348	0.3674
Condition: SH	-0.01173	0.15632	-0.075	0.9402	0.4701
Condition: FL	0.14076	0.15636	0.9	0.368	0.184
Condition: IFL	0.27044	0.15673	1.726	0.0844.	0.0422*

*FL/SH vs. *FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**
(Intercept)	-0.07222	0.16698	-0.433	0.665367	0.332684
Condition: SH	0.57972	0.18039	3.214	0.00131**	0.000655***
Condition: FL	0.6564	0.181	3.627	0.000287***	0.000144***
Condition: IFL	-0.29479	0.17941	-1.643	0.100352	0.050176.

*FL/SH vs. FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**
(Intercept)	-0.2045	0.1563	-1.309	0.19063	0.095315.
Condition: SH	0.457	0.1763	2.592	0.00955**	0.004775**
Condition: FL	0.5418	0.1767	3.067	0.00216**	0.00108**
Condition: IFL	-0.4116	0.1787	-2.303	0.02128*	0.01064*

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

When the IFL group was given the pairing of *FL/*SH vs. FL/*SH, the log-odds of choosing FL/*SH was significantly higher than that of the Control group (p1-tailed = .04). IFL group's log-odds of choosing FL/SH when they were given the choice of *FL/*SH vs. FL/SH went the opposite direction of the prediction made by FL grammar. Due to this reason, we ran a post-hoc 2-tailed test in order to interpret the results. The analysis indicated that the log-odds of the IFL group choosing FL/SH is not significantly different from the Control group (p2-tailed = .100). Lastly, the log-odds of choosing FL/SH when subjects heard FL/SH vs. FL/*SH again went the opposite direction of FL grammar prediction. A 2-tailed post-hoc test was run, and confirms that IFL subjects were less likely to choose FL/SH when they heard FL/SH vs. FL/*SH was significantly lower than Control group (p2-tailed = .021). Combining the results from all three different pairings, we can conclude that subjects in the IFL group only preferred stimuli that conformed to FL, but not to SH (i.e. FL/*SH).

These results indicated that subjects who were given only FL/*SH stimuli during training internalized a different rule than FL. Since all FL/*SH stimuli were either instantiated as [s… ʃ…s] or [ʃ…s…ʃ] (on sibilant level), it is likely that these subjects internalised a sibilant disharmony rule which requires neighbouring sibilants to be disharmonic to each other. This could explain why subjects did not prefer FL/SH (e.g. [s...s...s]) type of words. Nonetheless, Intensive FL subjects definitely failed to internalise the FL grammar that was intended in this study, and together with the SH and FL condition results obtained from a carefully controlled experimental setting, FL is harder to learn than SH.

4.7 Discussion

The experiments in this study were designed to test the learnability of two phonotactic patterns in the fairest possible way. The learnability of two patterns which are minimally different in phonological terms, but differ in their computational characterizations, were compared. The results have shown that SH was readily learned by humans in this paradigm. A mere 20 minutes of exposure to the grammar was sufficient to significantly affect the subjects' behavior. The performance of the SH subjects matched the predictions in all three types of test-pairings, and therefore provides strong evidence that the SH grammar was internalized.

These results were expected, as SH is both attested and belongs to SP. On the other hand, subjects who were exposed to the FL grammar did not perform according to the predictions. The only way to establish whether FL was learned is to examine the subjects' overall performance in all three pairings. Since only the results for one pairing concurs with the FL grammar's prediction, there was insufficient evidence to claim that FL was successfully learned in this experiment.

It could be true that FL would be learnable if the amount of training was increased or if the stimuli were presented in a different format/method. The crucial argument drawn from these results is that given the same experimental setting, and the same amount of training, the SH grammar was learned but the FL grammar was not (see Table 12). FL was at least more challenging for the subjects to internalize than SH was.

Furthermore, the subjects who were exposed to FL performed very similarly to the subjects who were exposed to SH. It must to be noted that there is an overlap of words that conform to both SH and FL (i.e. FL/SH), and the proportion of such words was limited to 50% of the entire set of training items. The remaining 50% consisted of words that did not conform to SH. That means that for every one word that could be construed as evidence for SH, there was another word that was not consistent with SH. Their performance could be explained by a heavy SH bias, which is influential enough to suppress the counterevidence. Pearl (2008) suggests that children implement a filter on the data when they are faced with ambiguous linguistic input that causes them to ignore information in the ambiguous data. If this theory can be extended to adults learning a new language in an experimental paradigm, the implementation of a filter would be a plausible explanation for why the FL subjects ignored part of the training data.

Because of the apparent SH bias exhibited by the subjects in the FL condition, an additional condition, Intensive First-Last (IFL), was run. All the training items in the original FL condition which conformed to both SH and FL were replaced with words that conformed to only FL in IFL. The purpose of testing this condition was to verify whether FL could be learned if the SH bias was alleviated by removing potentially distracting or ambiguous stimuli. The results for this condition were significantly different from the results obtained in the FL condition. However, these results were still inconsistent with the predictions made for the FL grammar. The IFL subjects consistently picked the words that only conformed to FL and not SH (recall that SH words necessarily also conform to FL). This problem could be due to the distribution of the training stimuli – the IFL subjects were given only a subset of the FL grammar; more precisely, they were only given words that conformed to FL but not SH. As a result, the IFL training items are also consistent with a sibilant disharmony rule which requires neighboring sibilants to be disharmonic. The lack of FL/SH words in the training may have prevented the subjects from generalizing to FL/SH as a possible word. Thus the IFL subjects may have been reluctant to generalize to the full-fledged FL grammar and instead assumed the disharmony grammar.

A possible follow-up could be designed to correct this problem. Instead of replacing all the training items that conform to both FL and SH, just the proportion of them could be lowered. This would differentiate the FL grammar from the disharmony grammar (as FL/SH is inconsistent with the disharmony rule), but the SH bias could still be alleviated due to the fewer occurrences of such words.

Nonetheless, even if FL was learned in this paradigm, the results would be meaningless unless SH was not learned under the same paradigm (see Table 12). In other words, the experiments run to date match the predictions of the Subregular Hypothesis.

4.8 Conclusions

The experimental results of this study have provided empirical evidence for the difference in learnability of two carefully matched phonotactic patterns. The SH pattern is an attested long-distance dependency pattern that belongs to the Strictly-Piecewise class. FL, on the other hand, is an unattested, non-SL, non-SP, non-TSL but regular pattern. The learnability of these two patterns was compared, and the results suggest that SH was learnable in the experimental paradigm used in this study, while FL was not. The results concur with the H4 which predicts the domain-specific phonological learning mechanism is subjected to computational constraints and also paradigm 2 as indicated in Table 12 that the SL/SP/TSL pattern was and while the non-SL/SP/TSL pattern was not. This suggests that phonotactic patterns that reside outside of the SL, SP, and TSL classes are less easily learned than those that reside within them. These findings imply that the computational boundaries proposed by the Subregular Hypothesis are psychologically real. In sum, the experimental results support that the phonological learning mechanism is constrained by specific computational properties, such that subregular patterns are more easily learned than regular but non-subregular patterns.

Chapter 5

EXPERIMENT 3: NON-LINGUISTIC EXPERIMENTS

5.1 Experiment 3a: Visual Shape Experiment

A series non-linguistic visual experiments was carried out to determine whether the subregular restrictions on the phonological learning mechanism also apply to pattern learning in different cognitive domain.

5.1.1 Subjects

Fifty-nine native American-English monolingual speaking adults with normal or corrected vision were recruited for the experiments. Twenty subjects were assigned to the control condition, another 20 were assigned to FL condition, and 19 were assigned to the SH condition. Subjects were students from the University of Delaware, aged between 18-27 years old, and their participation was compensated for either with course credit or $10.

5.1.2 Procedure

The experiment took place in a soundproof booth in the Phonetics and Phonology laboratory at the University of Delaware. The experiment consisted of 2 experimental conditions (SH and FL) and a control condition. The procedure for both conditions consisted of 2 phases: a training phase and a testing phase. The total duration for both training and testing was about 20 minutes.

During the training phase, subjects were instructed to watch the computer screen placed in front of them, with shape sequences that conformed either to an SH or FL grammar (depending on the experimental condition). They were informed that all the sequences belong to a particular set. The training contained 200 tokens (40 words x 5 repetitions) and the duration was approximately 15 minutes. In the control condition, no training was given—subjects were only given the test.

Training was followed by a testing phase in which the subjects were presented with pairs of sequences and were asked to judge whether the first sequence or the second sequence of the pair was more likely to belong to the set of sequences they just saw during the training. There were 48 pairs of test items in total, and the test took about 10 minutes to complete. All subjects, regardless of which condition they were in, were given the same test with the exact same pairs of 48 pairs of test items.

5.1.3 Stimuli

The FL and SH shape training stimuli were constructed based on the linguistic stimuli. Each sound segment of the linguistic stimuli was converted to a certain form with a certain color. All vowels were converted to circles of different colors, sibilants were converted to triangles, and the stop [k] was converted to a square. See Table 26 for the complete table of conversion.

Table 26 Segment-to-shape conversion.

Shapes	Sound	Shapes	Sound
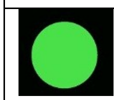	[a]	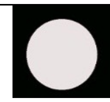	[u]
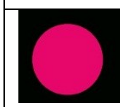	[ɛ]	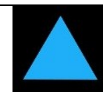	[s]
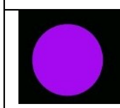	[i]	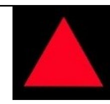	[ʃ]
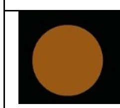	[ɔ]	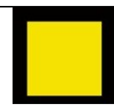	[k]

Each sequence consisted of 7 shapes, and each shape was presented one by one on the center of the computer screen for 400ms. The total duration of each shape sequences is 2800ms (400ms/shape x 7 shapes = 2800ms/sequence).

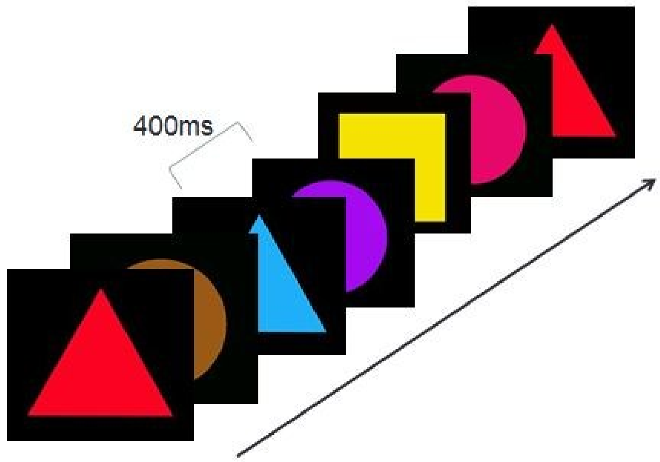

Figure 16 The presentation of a shape sequence.

Each sequence was preceded by a 500ms fixation screen marked with "+", followed by a blank screen of another 500ms. During the test phase, sequences were paired just like the linguistic condition. Each test stimuli was preceded by either "1" or "2" followed by a fixation slide to indicate the order of that sequence within a test pair.

5.1.4 Predictions

The predictions are exactly the same as the predictions for Experiment 2.

5.1.5 Results

The descriptive statistics for the rates of choosing FL/*SH and FL/SH in all three types of test pairings are summarized in Table 27.

Table 27 Descriptive statistics of Control, SH and FL Conditions in Experiment 3a.

	Conditions		
	Control	**SH**	**FL**
FL/*SH vs. *FL/*SH			
Mean Rate of FL/*SH *(SE)*	0.513 (0.028)	0.582 (0.028)	0.694 (0.026)
FL/SH vs. *FL/*SH			
Mean Rate of FL/SH *(SE)*	0.484 (0.028)	0.766 (0.024)	0.631 (0.027)
FL/SH vs. FL/*SH			
Mean Rate of FL/*SH *(SE)*	0.450 (0.028)	0.648 (0.027)	0.419 (0.028)

The results obtained from the visual experiment were analyzed the same way as the results obtained from Experiment 2.

The rates of choosing FL/*SH when subjects were presented with FL/*SH vs. *FL/*SH in all three conditions are shown in Figure 17.

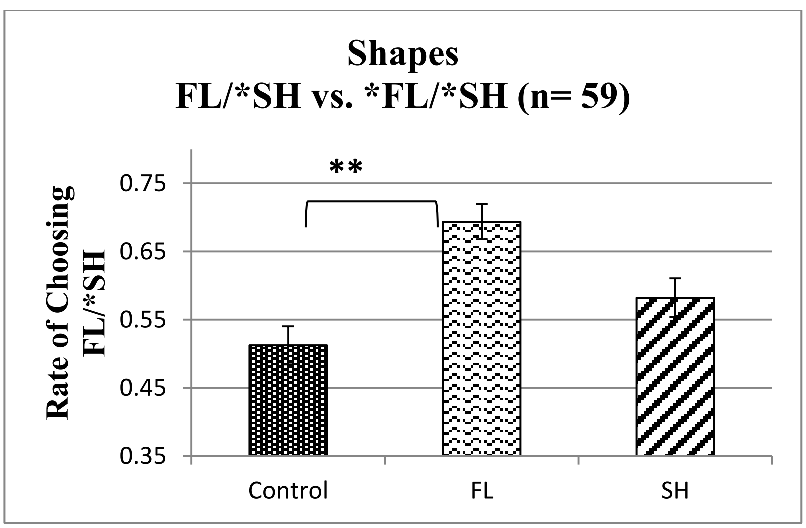

Figure 17 Mean rates of choosing FL/*SH when presented with FL/*SH vs. *FL/*SH in Exp 3a.

The likelihood ratio tests showed that all 3 models with the fixed factor CONDITION were significantly different from their respective empty models. The first model FL/*SH vs. *FL/*SH, the second model, FL/SH vs. *FL/*SH and third model FL/SH vs. FL/*SH were all significantly different from their empty models ($\chi2 = 9.79$, $p = .007$, $\chi2 = 16.45$, $p < .001$ and $\chi2 = 14.91$, $p <.001$) respectively. This means CONDITION is an important factor in its own right in all three models.

The model for the FL/*SH vs. *FL/*SH pairings showed that the FL groups' responses were significantly different from the Control group's (shown as Intercept in Table 28). The log-odds of the FL subjects choosing FL/*SH significantly different from the Control subjects' (p1-tailed =.001), while the log-odds of the SH subjects choosing FL/*SH was not significantly higher than the Control subjects' (p2-tailed = .24).

Table 28 Estimates of the conditions in the analysis of subjects' response in pairing FL/*SH vs. *FL/*SH in Exp 3a.

*FL/*SH vs. *FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**
(Intercept)	0.054	0.190	0.282	0.778	0.389
Condition: SH	0.321	0.274	1.170	0.242	0.121
Condition: FL	0.886	0.277	3.205	0.00135**	0.000675***

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

The rates of choosing FL/SH when subjects were presented with FL/SH vs. *FL/*SH pairings in all three conditions are shown in Figure 18.

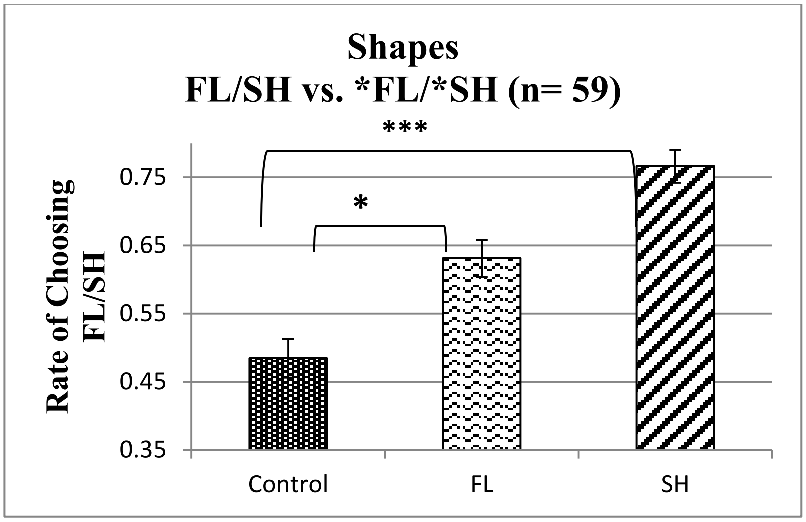

Figure 18 Mean rates of choosing FL/SH when presented with78 FL/SH vs. *FL/*SH.in Exp 3a.

The model for the FL/SH vs. *FL/*SH pairings suggests that the log-odds of the SH subjects choosing FL/SH was significantly higher than the Control subjects’ (p1-tailed < .001), and the log-odds of the FL subjects choosing FL/SH was also significantly higher than the Control subjects’ (p1-tailed =.014).

Table 29 Estimates of the conditions in the analysis of subjects’ response in pairing FL/SH vs. *FL/*SH in Exp 3a.

*FL/SH vs. *FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**
(Intercept)	-0.066	0.237	-0.278	0.781	0.391
Condition: SH	1.513	0.353	4.282	0.0000186***	0.0000093***
Condition: FL	0.742	0.339	2.185	0.0289*	0.01445*

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

The rates of choosing FL/SH when subjects were presented with the FL/SH vs. FL/*SH pairings in all three conditions are shown in Figure 19.

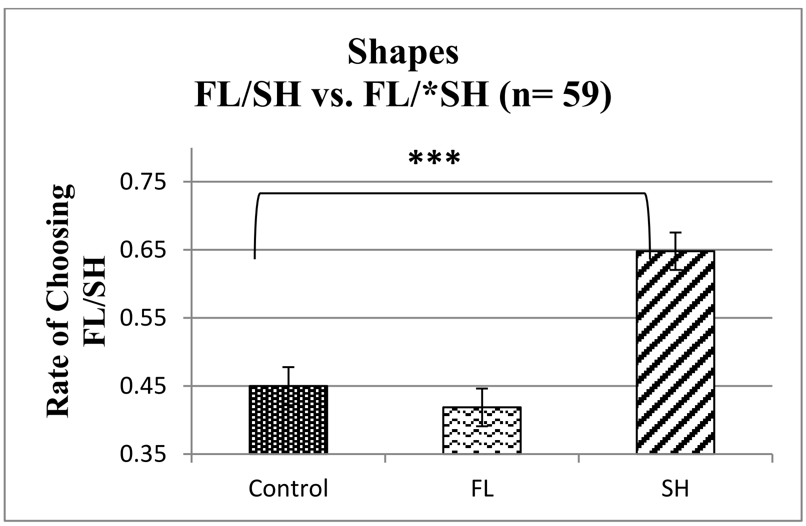

Figure 19 Mean rates of choosing FL/SH when presented with FL/SH vs. FL/*SH in Exp 3a.

The model for the third test pairing, FL/SH vs. FL/*SH, suggests that the log-odds of the SH subjects choosing FL/SH was significantly higher than the Control subjects' (p1-tailed = .001), and the log-odds of the FL subjects choosing FL/SH was not significantly higher than the Control subjects' (p2-tailed =.582).

Table 30 Estimates of the conditions in the analysis of subjects' response in pairing FL/SH vs. FL/*SH in Exp 3a.

*FL/SH vs. FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**
(Intercept)	-0.220	0.197	-1.118	0.264	0.132
Condition: SH	0.924	0.284	3.252	0.00115**	0.000575***
Condition: FL	-0.152	0.277	-0.550	0.582	0.291

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

The data was further analyzed by converting the binary responses to the rate of response. Experiment 3a was analyzed in the same way as Experiment 2. There were three groups in the experiment, each representing a training conditioning: Control, SH, and FL.

Preliminary comparisons revealed that the homogeneity assumption underlying an ANOVA was met for the last two DV but not the first DV. The F statistics is consequently replaced by the Welch statistics for the first DV. Post hoc comparisons were apportioned using the Tukey adjustment is used for the last two DV, and Games-Howell is used for the first DV. The ANOVA with DV(1) the rate of choosing FL/*SH item when presented with FL/*SH and *FL/*SH showed a statistically significant difference between groups (Welch Statistics (2, 33.347) = 4.964, p =.745).

The ANOVA with DV (2) the rate of choosing FL/SH item when presented with FL/SH and *FL/*SH showed a statistically significant difference between groups ($F (2, 56) = 8.873$, $p < .001$). The ANOVA with DV (3) the rate of choosing FL/SH item when presented with FL/SH and FL/*SH also showed a statistically significant difference between groups ($F (2, 56) = 7.765$, $p < .001$).

Post hoc analyses demonstrated that FL condition produced significantly higher rates of choosing FL/*SH when presented with FL/SH vs. *FL/*SH than the Control condition ($p = .012$). SH condition produced significantly higher rates of choosing FL/SH when presented with FL/SH vs. *FL/*SH ($p < .001$), and FL condition's rate of choosing FL/SH is only marginally significantly higher than control condition ($p = .076$). Finally, only SH condition produced significantly higher rates of choosing FL/SH when presented with FL/SH vs. FL/*SH than the Control condition ($p = .007$).

A separate analysis was run to test whether the FL group performed significantly different from the SH group. The SH group was coded as the reference group (Intercept) in this analysis, and since no specific direction was predicted for the results, 2-tailed tests were used.

Table 31 Estimates of the conditions in 3 types of test pairings with SH group as the reference group in Exp 3a.

*FL/SH vs. FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**
(Intercept)	0.375	0.198	1.894	0.0582.
Condition:Control	-0.321	0.274	-1.170	0.242
Condition: FL	0.565	0.282	2.005	0.045*
*FL/SH vs. FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**
(Intercept)	1.447	0.263	5.496	0.0000000388***
Condition:Control	-1.513	0.353	-4.282	0.0000185***
Condition: FL	-0.772	0.358	-2.156	0.0311*
*FL/SH vs. FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**
(Intercept)	0.704	0.208	3.391	0.000697***
Condition:Control	-0.924	0.284	-3.252	0.001146**
Condition: FL	-1.077	0.285	-3.775	0.00016***

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

As shown in Table 31, the responses of the FL group in all three types of test pairings are significantly different from the SH group: FL/*SH vs. *FL/*SH ($p = .045$), FL/SH vs. *FL/*SH ($p = .031$) and FL/SH vs. FL/*SH ($p < .001$). These results are different from the ones obtained from Experiment 2, where the FL and the SH subjects’ responses were not significantly different from each other in all three models.

The results obtained match the predictions made by the SH *and* the FL grammar. Therefore, it can be concluded that the SH subjects were able to internalize the SH grammar, and the FL subjects were able to internalize the FL grammar. These results are clearly different from the ones obtained in Experiment 2, and they suggest that there is no SH bias when stimuli were presented visually in a non-linguistic context. Another possible explanation for these results is that the task was too easy due to the canonical shapes used to construct the stimuli in this experiment (see Section 5.1.6 below for a more detailed discussion).

5.1.6 Discussion

The experiments in this study were designed to test the learnability of SH and FL pattern within a visual non-linguistic context. The idea is to compare whether these results would be similar to the ones obtained from the linguistic context in order to investigate whether the limitations of the phonological learner was general restrictions for learning patterns in another cognitive domain. The patterns which were tested in this experiment were embedded in canonical shapes with various colors rather than in words as in Experiment 2. The results have shown that both SH and FL were readily learned by humans in this context. The performance of the SH and FL subjects matched the predictions in all three types of test-pairings, and therefore provides strong evidence that both the SH and FL grammar was internalized. Another difference was noted in the informal interviews after the experiment was completed: most of the subjects in this visual experiment could verbalize the actual pattern they picked up, whereas, the subjects in the linguistic experiment could not. While it could be argued that the SH and FL grammar were learned more successfully in the visual context, it is also possible that the canonical shapes and colors are more familiar to the subjects than are sound segments.

A potential problem that could be caused by the visual stimuli is that complexity of the shape/color combinations cannot match the complexity of the sound segments. This could make the task of learning SH/FL grammar in the visual context easier in general than in the linguistic context. A possible follow-up could be designed to correct this problem. Instead of using canonical shapes, non-canonical shapes could be used.

Another possible problem is that the duration of each shape sequence was longer than linguistic stimuli (2800ms vs. 1400ms). Therefore, the subjects in the visual experiment had a longer exposure time of each sequence than the subjects in the linguistic experiment had. However, it is also true that the subjects in the linguistic experiment were asked to repeat each word after it was presented, but the visual subjects were not. However, it is still unclear whether or not the act of repetition can enhance the learning effect, therefore, it cannot be judged whether verbal repetition could be compensated by longer duration of exposure of each shape sequence. One way to solve this problem is to increase the duration of the linguistic stimuli to match the visual stimuli, but this would make each word sound unnatural. Another way is to shorten the presentation of the visual stimuli, but as some subjects in the pilot study complained, this would make the screen difficult to watch due to the very rapid flashing.

Nonetheless, under the current paradigm, the learnability of SH and FL grammars was shown to be different in the linguistic and the visual domains. Subjects in visual experiment could learn the FL pattern as well as the SH pattern, while subjects in the linguistic experiment learned the SH pattern better than the FL pattern. The difference in the results obtained from these two experiments suggest that visual learning mechanism is less restrictive than the phonological learning mechanism, and this suggests that the phonological learner is distinct from the visual learner.

5.1.7 Conclusions

The experimental results of this study have provided empirical evidence for the difference in restrictiveness in two cognitive domains (phonological vs. non-linguistic visual). This means the learning mechanism for phonology is not the same as the one for non-linguistic visual patterns. These findings are consistent with the domain-specific hypothesis. However, the phonological domain and the non-linguistic visual domain differ in at least two dimensions: 1) linguistic vs. non-linguistic and 2) auditory vs. visual, therefore, the difference observed in the learnability of FL between these two experiments could be due to either dimensions. In order to be able to tease two dimensions apart, another experiment was carried out (non-linguistic auditory experiment). This experiment eliminates the difference visual vs. auditory differences in the stimuli, and if the difference in the learnability of FL maintains in the non-linguistic auditory experiment, one can be sure that the subregular restrictions observed in the phonological learner is only exclusive to the phonological learner.

5.2 Experiment 3b: Auditory Experiment

5.2.1 Subjects

Sixty native American-English monolingual speaking adults with normal hearing were recruited for the experiments. Twenty subjects were assigned to each of the three conditions (Control, SH, and FL). Subjects were students from the University of Delaware, aged between 18-27 years old, and their participation was compensated for either with course credit or $10.

5.2.2 Procedures

The experiment took place in a soundproof booth in the Phonetics and Phonology laboratory at the University of Delaware. The experiment consisted of 2 experimental conditions (SH and FL) and a control condition. The procedure for both conditions consisted of 2 phases: a training phase and a testing phase. The total duration for both training and testing was about 20 minutes.

During the training phase, subjects were instructed to listen to sequences of drumbeats that conformed either to an SH or FL grammar (depending on the experimental condition). They were informed that all the sequences belong to a particular set. The training contained 200 tokens (40 sequences x 5 repetitions) and the duration was approximately 10 minutes. In the control condition, no training was given—subjects were only given the test.

Training was followed by a testing phase in which the subjects were presented with pairs of sequences and were asked to judge whether the first sequence or the second sequence of the pair was more likely to belong to the set of sequences they just heard during the training. There were 48 pairs of test items in total, and the test took about 10 minutes to complete. All subjects, regardless of which condition they were in, were given the same test with the exact same pairs of 48 pairs of test items.

5.2.3 Stimuli

The auditory stimuli were constructed based on the linguistic stimuli in Experiment 2. Each sound segment was converted to a certain form with a certain color. All vowels were converted to cymbals, sibilants were converted to tom drum, and the stop [k] was converted to a snare drum. See Table 32 for the complete table of conversion.

Table 32 Segment-to-drumbeat conversion.

Drumbeat	**Sound**	**Drumbeat**	**Sound**
closed high hat	[a]	open high hat	[u]
ride	[ɛ]	low tom	[s]

“washy” high hat	[i]	high tom	[ʃ]
crash	[ɔ]	snare drum	[k]

A control condition was conducted to ensure subjects did not have any inherent preference for any particular type of drumbeat sequences. The mean length of the linguistic stimuli (~1400ms), and in order to match this duration, each drumbeat was constructed to last for 200ms, and since there were 7 beats in each sequence, the total duration of each sequence was also 1400ms (200ms/beat x 7 beats = 1400ms/sequence).

5.2.4 Predictions

The predictions for this experiment are the same as the ones in Experiment 2.

5.2.5 Results

The descriptive statistics for the rates of choosing FL/*SH and FL/SH in all three types of test pairings are summarized in Table 33.

Table 33 Descriptive statistics of Exp. 3b.

	Conditions		
	Control	**SH**	**FL**
FL/*SH vs. *FL/*SH			
Mean Rate of FL/*SH *(SE)*	0.494 (0.028)	0.481 (0.028)	0.488 (0.028)

FL/SH vs. *FL/*SH			
Mean Rate of FL/SH *(SE)*	0.534 (0.028)	0.406 (0.027)	0.531 (0.028)
FL/SH vs. FL/*SH			
Mean Rate of FL/*SH *(SE)*	0.497(0.028)	0.397 (0.027)	0.559 (0.028)

The results obtained from this auditory non-linguistic experiment were analyzed the same way as the results obtained from Experiments 2 and 3a.

The rates of choosing FL/*SH when subjects were presented with FL/*SH vs. *FL/*SH in all three conditions are shown in Figure 20.

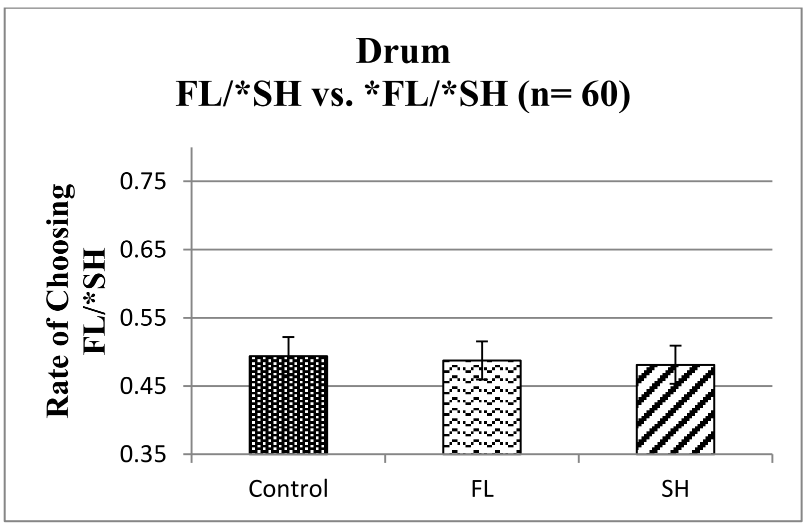

Figure 20 Mean rates of choosing FL/*SH when presented with FL/*SH vs. *FL/*SH in Exp 3b.

The likelihood ratio tests showed that 2 out of 3 models with the fixed factor CONDITION were significantly different from their respective empty models. The first model FL/*SH vs. *FL/*SH is not significantly different from its empty model, ($\chi 2$ = 0.10, p = .951). The second model, FL/SH vs. *FL/*SH and third model FL/SH vs. FL/*SH were both significantly different from their empty models ($\chi 2$ = 7.77, p = .021 and $\chi 2$ = 9.56, p = .008) respectively. This means CONDITION is an important factor in its own right only in second and the third models.

The model for the FL/*SH vs. *FL/*SH pairings showed that the FL groups' responses were not significantly different from the Control group's (shown as Intercept in Table 34). The log-odds of the FL subjects choosing FL/*SH is not significantly different from the Control subjects' (p1-tailed =.437), while the log-odds of the SH subjects choosing FL/*SH was not significantly higher than the Control subjects' (p2-tailed = .376).

Table 34 Estimates of the conditions in the analysis of subjects' response in pairing FL/*SH vs. *FL/*SH in Exp. 3b.

*FL/*SH vs. *FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**
(Intercept)	-0.025	0.114	-0.220	0.826	0.413
Condition: SH	-0.050	0.158	-0.317	0.751	0.3755
Condition: FL	-0.025	0.158	-0.158	0.874	0.437

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

The rates of choosing FL/SH when subjects were presented with FL/SH vs. *FL/*SH pairings in all three conditions are shown in Figure 21.

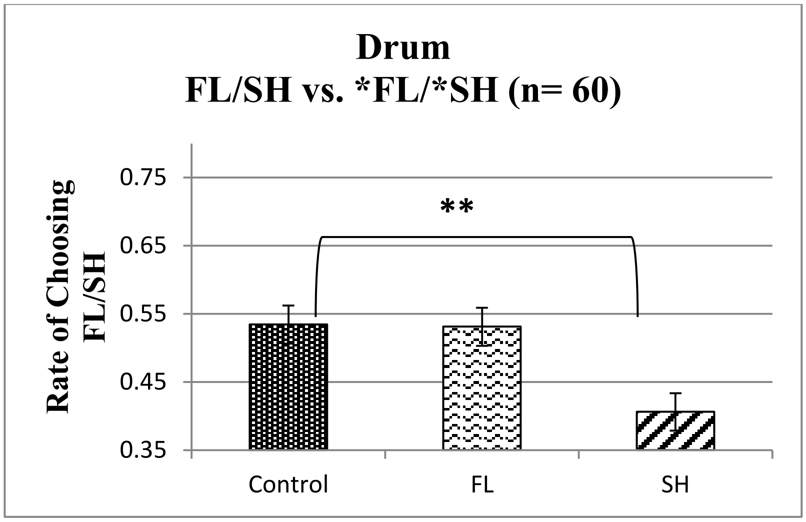

Figure 21 Mean rates of choosing FL/SH when presented with FL/SH vs. *FL/*SH in Exp. 3b.

The mean rate of choosing FL/SH of the SH subjects was lower than that of the Control subjects. This result was unexpected because it indicates the SH subjects preferred the items that did not conform to the SH grammar more, hence a novelty preference, and this is the opposite of what was found in the linguistic condition. Since this was not the predicted direction of the results, the two-tailed test was used for the models in this experiment to detect whether the experimental results were significantly different from the control results. The model for the FL/SH vs. *FL/*SH pairings suggests that the log-odds of the SH subjects choosing FL/SH was significantly different from the Control subjects' (p1-tailed = .006), and the log-odds of the FL subjects choosing FL/SH was not significantly different than the Control subjects' (p2-tailed =.957).

Table 35 Estimates of the conditions in the analysis of subjects' response in pairing FL/SH vs. *FL/*SH in Exp. 3b.

*FL/SH vs. *FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**
(Intercept)	0.149	0.184	0.814	0.4159	0.20795
Condition: SH	-0.570	0.225	-2.531	0.0114*	0.0057**
Condition: FL	-0.012	0.224	-0.054	0.957	0.4785

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

The rates of choosing FL/SH when subjects were presented with the FL/SH vs. FL/*SH pairings in all three conditions are shown in Figure 22.

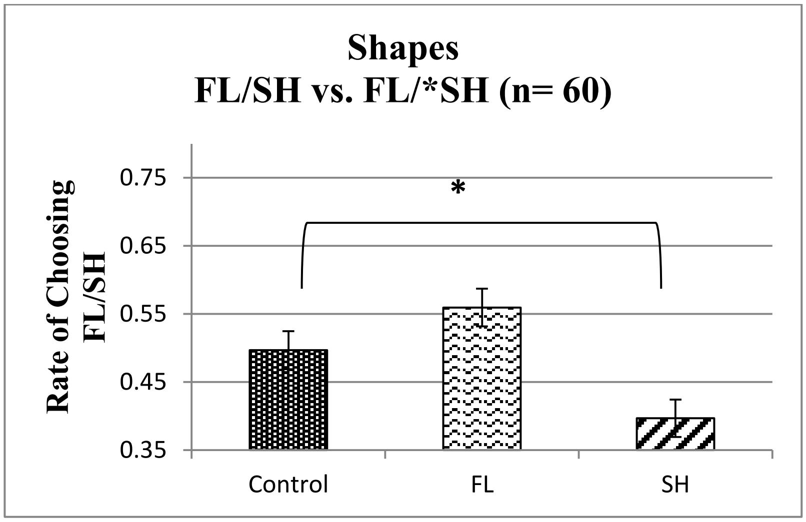

Figure 22 Mean rates of choosing FL/SH when presented with FL/SH vs. FL/*SH in Exp 3b.

The model for the third test pairing, FL/SH vs. FL/*SH, suggests that the log-odds of the SH subjects choosing FL/SH was significantly different from the Control subjects' (p2-tailed = .047), and the log-odds of the FL subjects choosing FL/SH was not significantly different from the Control subjects' (p2-tailed =.227).

Table 36 Estimates of the conditions in the analysis of subjects' response in pairing FL/SH vs. FL/*SH in Exp. 3b.

*FL/SH vs. FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**	***p* (1-tailed)**
(Intercept)	-0.013	0.154	-0.084	0.933	0.4665

Condition: SH	-0.436	0.219	-1.986	0.047*	0.0235*
Condition: FL	0.263	0.218	1.209	0.227	0.1135

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

The data was further analyzed by converting the binary responses to the rate of response. Experiment 3b was analyzed in the same way as Experiment 2 and Experiment 3a. There were three groups in the experiment, each representing a training conditioning: Control, SH, and FL.

Preliminary comparisons revealed that the homogeneity assumption underlying an ANOVA was met. Therefore, post hoc comparisons were apportioned using the Tukey adjustment. The ANOVA with DV(1) the rate of choosing FL/*SH item when presented with FL/*SH and *FL/*SH did not show a statistically significant difference between groups ($F (2, 57) = .066, p =.936$). The ANOVA with DV (2) the rate of choosing FL/SH item when presented with FL/SH and *FL/*SH showed a statistically significant difference between groups ($F (2, 57) = 3.896, p = .027$). The ANOVA with DV (3) the rate of choosing FL/SH item when presented with FL/SH and FL/*SH also showed a statistically significant difference between groups ($F (2, 57) = 4.824, p = .012$).

Post hoc analyses demonstrated that SH condition produced significantly *lower* rates of choosing FL/SH when presented with FL/SH vs. *FL/*SH than the Control condition ($p = .046$). Neither SH condition and FL condition produced significantly different rates of choosing FL/SH when when presented with FL/*SH vs. *FL/*SH.

A separate analysis was run to test whether the FL group performed significantly different from the SH group. The SH group was coded as the reference group (Intercept) in this analysis, and since no specific direction was predicted for the results, 2-tailed tests were used.

Table 37 Estimates of the conditions in 3 types of test pairings with SH group as the reference group in Exp 3b.

*FL/*SH vs. FL/SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**
(Intercept)	-0.075	0.114	-0.659	0.510
Condition:Control	0.050	0.158	0.317	0.751
Condition: FL	0.025	0.158	0.158	0.874
*FL/SH vs. *FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**
(Intercept)	-0.420	0.186	-2.261	0.0237*
Condition:Control	0.569	0.225	2.530	0.0114*
Condition: FL	0.557	0.225	2.473	0.0134*
*FL/SH vs. FL/*SH*	**Estimate**	**Standard Error**	*z*	***p* (2-tailed)**
(Intercept)	-0.449	0.158	-2.849	0.00438**
Condition:Control	0.436	0.219	1.986	0.04704*
Condition: FL	0.699	0.220	3.179	0.00148**

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

As shown in Table 37, the responses of the FL group in 2 out of 2 types of test pairings are significantly different from the SH group: FL/*SH vs. *FL/*SH (p = .874), FL/SH vs. *FL/*SH (p = .013) and FL/SH vs. FL/*SH (p = .001). These results are different from the ones obtained from Experiments 2 and 3a in different way. In Experiment 2, the FL and the SH subjects' responses were not significantly different from each other in all three models, and in Experiment 3a, the FL and the SH subjects' responses were significantly different in all 3 models.

The results obtained match the predictions made by the SH but not the FL grammar. This is similar to what was found in Experiment 2 where phonological stimuli were used. Similarly, it can be concluded that the SH subjects were able to internalize the SH grammar, and the FL subjects were able to internalize the FL grammar in this non-linguistic auditory context. However, the details of the results are clearly different from the ones obtained in Experiments 2 and 3a but in different ways. Firstly, although FL was not as readily learned as SH in this non-linguistic auditory context, the SH bias, unlike the results obtained from the linguistic experiment. Secondly, FL was not as learnable as SH, which is not what was found in the non-linguistic visual experiment. Lastly, the subjects in this experiment showed a preference for sequences which were dissimilar to the training set, despite of the question asked was to choose the sequences which they thought was more likely to belong to the training set. The last point may be trivial because as long as the preference for similar/dissimilar items is consistent and systematic, one can still conclude the subjects reacted psychologically differently for one particular set of items over the other. On the other hand, the first two points were more noteworthy, as those are the differences that would influence how the conclusions should be drawn.

5.2.6 Discussion

The experiments in this study were designed to test the learnability of SH and FL pattern within a auditory non-linguistic context. The idea is to compare whether these results would be similar to the ones obtained from the linguistic context in order to investigate whether the limitations of the phonological learner was general restrictions for learning patterns in another auditory cognitive domain. The patterns which were tested in this experiment were embedded in sequences of drumbeats and were presented aurally. The results have shown that SH was readily learned by humans in this context, but not FL. The performance of the SH subjects matched the predictions in all three types of test-pairings even though they showed a preference of dissimilar items rather than similar items. However, they still showed a systematic dispreference for sequences that conformed to the SH, and therefore the results still constitute as strong evidence that the SH grammar was internalized. Subjects who were exposed to the FL grammar in this experiment showed no significant difference from the Control subjects, who had no training at all. This is interesting because even though FL was not as learnable as SH in the non-linguistic auditory context, the results are different from the phonotactic experiment (Exp 2). The FL subjects in this experiment did not show a bias towards learning SH unlike what was seen in the phonotactic experiment. Based on these results, it seems like the non-linguistic auditory learner has the same subregular restrictions as the phonological learner, but due to the lack of SH bias in the auditory learner, the conclusion that these two domains share one single learner could not be drawn. Other than the linguistic/non-linguistic difference of the stimuli, subjects in the non-linguistic auditory experiment were also not asked to repeat what they heard during the training, as it may seem like an impossible task. This rules out the possibility that the difference between this experiment and the Experiment 2 was due to the length of exposure.

The difference observed in the auditory non-linguistic experiment and the visual non-linguistic experiment is more apparent. The FL patterns could be learned in the visual domain but not in the auditory domain. However, the question of whether the visual canonical shapes were easier to process than the drum beats could be problematic to answer. First, the labels of canonical shapes are familiar to all of the subjects who participated, but the labels of the type of drum beats used were not. Again, this problem could be eradicated by using non-sense shapes in the future work.

Yet, under the current paradigm, the learnability of SH and FL grammars was shown to be similar in the linguistic and the non-linguistic auditory domains despite the lack of SH bias in the non-linguistic domain. Subjects in non-linguistic auditory experiment could not learn the FL pattern as well as the SH pattern, but they also did not ignore the FL training stimuli, and assumed the SH pattern when they were trained on FL. The similary in the results obtained from these linguistic and the non-linguistic auditory experiments suggests that auditory learning mechanism share some similarities with the phonological learning mechanism, but whether they are identical remains questionable. However, the presence of the SH bias shown in the linguistic auditory condition and the lack of it in the non-linguistic condition suggests that the linguistic stimuli were treated differently than the non-linguistic stimuli.

5.2.7 Conclusions

The experimental results of this study have provided empirical evidence for the similarity in restrictiveness in two cognitive domains (phonological vs. non-linguistic auditory). Subjects who were in the phonotactic experiment and the auditory non-linguistic experiment both failed to learn a regular but non-SL/-SP/-TSL pattern.

Chapter 6

GENERAL DISCUSSION AND CONCLUSIONS

This dissertation investigated the domain specificity of phonological learning. It began by asking the question of whether the phonological learning mechanism is the same as the one for syntax learner. This question was raised due to the asymmetry that is observed in the typology of phonological patterns and the syntactic patterns. The attested phonological patterns belong to the SL, SP or TSL classes of language. These are all proper subset of the regular languages which is a proper subset of the context free language where certain syntactic patterns can be found. In other words, the patterns found in syntax are computationally more complex than the patterns found in phonology. The fact that context free patterns are not found in phonology could be due to a number of reasons- it could be purely accidental or due to some other reasons. Heinz and Idsardi (2011) suggest that one of the possible reasons could be because the learner which is responsible for learning phonology and syntax are different, and their ability to compute grammars of different complexity is also different. If a single general learner is responsible for learning language of different aspects, this learner must be able to learn the context free syntactic pattern which is present in natural language's syntax, as well as context free pattern in the phonological domain, which is unattested in natural languages' phonology. However, if the phonological learner is distinct from the syntactic learner, and the absence of context free phonotactic learner could be due to the inability of the phonological learner to learn such pattern. These hypotheses were tested in Experiment 1, where the context free pattern was embedded in both a word (phonology) and a sentence (syntax). The results obtained from Experiment 1 suggest that the context free pattern was learned in the syntactic context, but not in the phonological context. These results in turn support that the claim that the phonological learner is distinct from the syntactic learner.

Experiment 2 was run to investigate whether the fact that only certain subregular classes of patterns (SL, SP and TSL) are found in natural languages' phonology is due to the restrictions of the phonological learner. The learnability of a SP pattern, sibilant harmony was compared with the learnability of a regular, but non-SL/-SP/-TSL pattern, first last assimilation. If the learnability of these two patterns are not different (i.e. both are learnable), then the absence of FL is irrelevant to its learnability. However, if the sibilant harmony pattern is easier to learn than the first last assimilation, this may suggest that the absence of first-last is due to the limitations of the phonological learner. The results from Experiment 2 supports the latter, as subjects who were exposed to the FL grammar did not learn FL, but subjects who were exposed to SH learned SH grammar readily. This implies that the subregular boundary present in the typology can be accounted for by the computational restrictions of the phonological learner. It is also noted that a strong SH bias is present in the phonological domain. Subjects who were exposed to FL grammar internalized the SH rule even though half of the training stimuli that were given were inconsistent with SH.

Experiments 3a and 3b were carried out to investigate whether the subregular constraints shown in the phonological learning mechanism was shared by other non-linguistic domains. Experiment 3a tested the learnability of FL and SH when the stimuli were presented visually as sequences of shapes rather than words. The results obtained suggest that both FL and SH grammars were learnable in this paradigm. If this paradigm were on par with the one in Experiment 2, this would imply that the visual pattern learning mechanism is not restricted by the subregular constraints. Experiment 3b tested the learnability of FL and SH when the stimuli were presented aurally as sequences of drum beats. The reason for using drum beats is to test whether the difference in results obtained in Experiment 2 and 3a is due to a linguistic/non-linguistic difference or a visual/auditory difference.

The results obtained in Experiment 3b show that the FL pattern is more difficult to learn than the SH pattern in a non-linguistic context task. These results are seemingly the same as the ones obtained in Experiment 2, but in Experiment 3b, no SH bias was seen. Subjects who were exposed to FL in the non-linguistic auditory domain simply did not perform significantly different from the control subjects, and they did perform significantly different from the SH subjects. Even though the SH bias is not present in the non-linguistic context, it is still striking that the non-linguistic auditory learner also fail to learn a non-subregular FL pattern.

There is no straight-forward way to interpret these results, but one of the possible interpretations is that the subregular learning restrictions are applied to auditory domain in general instead of just to the phonological domain. In other words, it is possible that the auditory domains (both phonological and non-linguistic) share one general learning mechanism, and the SH bias was just an added filter only applied when phonotactic patterns are learned. It is not that surprising the same general learner is used for learning phonotactic patterns and arbitrary drum beats. Both types of strings are auditory, and are built up by very basic units (a sound segment vs. a basic beat), unlike the basic units of sentences (words).

Another possible interpretation is that the subregular constraints and the SH bias are combined properties of the phonological learner, and even though the same subregular restrictions are seen in the non-linguistic auditory learner, it lacks the SH bias, and this may suggest that even though some of the properties in these two distinct learners overlap, they are still distinct mechanisms. Although Experiment 2 and Experiment 3a and 3b provide different results with regard to whether FL could be learned, SH was always successfully learned in these experiments. It is worth noting that under no circumstance did these experiments show that FL could be learned while SH could not.

To conclude, the two domain specific hypotheses (H2 and H4) are supported by the experimental results, and that implies that the phonological learner is distinct from the syntactic learner. If two close-knitted domains such as syntax and phonology employ different mechanisms during learning, then it becomes more plausible that other cognitive domains have their own specialized learning mechanisms.

Results from Experiment 2 have shown that the phonological learner is restricted by computational boundaries (the subregular boundaries), therefore, H4 is supported. These results support the claim that the hypothesis space that was previously assumed for phonological patterns should be tightened from regular to subregular.

However, in the non-linguistic experiments, the results suggest that there is a partial overlap in the computational properties of the phonological and the non-linguistic auditory learning mechanisms. It should be noted though, the results in

Experiment 1 suggest, a context free syntactic pattern can be learned via aural means. One major difference between the syntactic stimuli and the phonological/drumbeat stimuli is that they were composed of units that are less basic than sound segment/drumbeat. The higher complexity of the syntactic stimuli could be a possible reason why its learner is different from the phonological and the non-linguistic auditory learners.

A possible way to test whether the learning mechanisms are divided according to the complexity of the stimuli's basic units is to embed the context free pattern in a type of non-linguistic stimuli that is of higher complexity than drum beats (such as musical phrases). If the context free pattern could be learned in this context, this hypothesis would be supported.

Future work that addresses questions raised in the above should be continued. With the help of various neuroimaging techniques, finer differences between the various learning mechanisms which are undetected by the ALL paradigm could be discovered, and that may be where the answers to these questions reside.

All in all, the experimental results collected from the three sets of experiments in this dissertation are consistent with the hypothesis that the phonological learning mechanism is distinct from the syntax one, and a non-linguistic auditory learning mechanism is different from a non-linguistic visual learning mechanism. The results are also consistent with the hypothesis that the phonological learning mechanism is restricted by computational constraints- a subregular pattern is more easily learned than a regular pattern.

REFERENCES

Archangeli, Diana and Douglas Pulleyblank. 2007. Harmony. In The Cambridge Handbook of Phonology, ed. by Paul de Lacy, 353-378. Cambridge: Cambridge University Press.

Aslin, Richard N., Jenny Saffran, and Elissa L. Newport. 1998. Computation of conditional probability statistics by human infants. Psychological Science 9:321-324.

Barnes, Jonathan. 2002. Positional neutralization: a phonologization approach to typological patterns. Ph.D. dissertation, University of California, Berkeley.

Bates, Douglas, Martin Maechler and Ben Bolker. 2011. lme4: Linear mixed-effects models using s4 classes [Computer software manual]. Available from http://CRAN.R-project.org/package=lme4

Beckman, Jill. 1998. Positional Faithfulness. Ph.D. dissertation, University of Massachusetts, Amherst.

Berent, Iris, Colin Wilson, Gary F. Marcus, and Douglas K. Bemis. 2012. On the role of variables in phonology: remarks on Hayes and Wilson 2008. Linguistic Inquiry 43(1): 97-119.

Blevins, Juliette. 2004. Evolutionary Phonology. Cambridge: Cambridge University Press.

Chambers, Kyle E., Kristine H. Onishi, and Cynthia Fisher. 2003. Infants learn phonotactic regularities from brief auditory experience. Cognition 87: B69-B7

Dell, Gary. S., Kristopher D. Reed, David R. Adams, and Antje S. Meyer. 2000. Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. Journal of Experimental Psychology: Learning, Memory, and Cognition 26:1355-1367.

Dettweiler, Stephen H. 2000. Vowel harmony and neutral vowels in C'Lela. The Journal of West African Languages 28, 3-18.

Endress, Ansgar.D., Marina Nespor, and Jacques Mehler. 2009. Perceptual and memory constraints on language acquisition. Trends in Cognitive Science 13: 348–353.

Finley, Sara. 2011. The privileged status of locality in consonant harmony. Journal of Memory and Language 65:74-83.

Finley, Sara. To appear. Testing the limits of long-distance learning: Learning beyond the three-segment window. Cognitive Science.

Finley, Sara and William Badecker. 2009. Artificial language learning and feature-based generalization. Journal of Memory and Language 61: 423-437.

Fitch, W. Tecumseh and Angela D. Friederici, 2012. Artificial grammar learning meets formal language theory: an overview. The Philosophical Transactions of the Royal Society B 367: 1933-1955.

Goldrick, Matt. 2004. Phonological features and phonotactic constraints in speech production. Journal of Memory and Language: 51: 586-603.

Goldsmith, John and Jason Riggle. To Appear. Information Theoretic Approaches to Phonological Structure: The Case of Finnish Vowel Harmony. Natural Language and Linguistic Theory.

Goldwater, Sharon and Mark Johnson. 2003. Learning OT Constraint Rankings Using a Maximum Entropy Model. In Proceedings of the Stockholm Workshop on 'Variation within Optimality Theory. April 26-27, 2003 at Stockholm Univ. Sweden, ed. by Jennifer Spenader, Anders Eriksson, and Östen Dahl, 111-120. Stockholm: Stockholm University, Sweden.

Gómez, Rebecca. L. 2002. Variability and detection of invariant structure. Psychological Science, 13, 431-436.

Hale, Mark, and Charles Reiss. 2000. Substance abuse and “dysfunctionalism”: current trends in phonology. Linguistic Inquiry 31(1):157–169.

Hansson, Gunnar. 2001. Theoretical and typological issues in consonant harmony. Ph.D dissertation, University of California, Berkeley.

Hayes, Bruce, and Colin Wilson. 2008. A maximum entropy model of phonotactics and phonotactic learning. Linguistic Inquiry 39(3):379–440.

Heinz, Jeffrey. 2010. Learning long-distance phonotactics. Linguistic Inquiry 41(4): 623-661.

Heinz, Jeffrey. 2011 a. Computational phonology part I: Foundations. Language and Linguistics Compass 5(4):140-152.

Heinz, Jeffrey. 2011 b. Computational phonology part II: Grammars, learning, and the future. Language and Linguistics Compass 5(4):153-168.

Heinz, Jeffrey, Chetan Rawal, and Herbert Tanner. 2011. Tier-based Strictly Local Constraints for Phonology. In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics, Portland, Oregon, USA.

Kaplan, Ronald and Martin Kay. 1994. Regular models of phonological rule systems. Computational Linguistics 20:331-378.

Koo, Hahn and Lydia Callahan. To appear. Tier-adjacency is not a necessary condition for learning phonotactic dependencies. Language and Cognitive Processes.

McNaughton, Robert and Seymour Papert. 1971. Counter-Free Automata. Cambridge, MA: MIT Press.

Newport, Elissa L. and Richard N. Aslin, R. 2004. Learning at a distance I. Statistical learning of nonadjacent dependencies. Cognitive Psychology 48:127-162.

Nowak, Martin, Natalia Komarova, and Partha Niyogi (2002). Computational and evolutionary aspects of language. Nature 417: 611-617.

Ohala, John. 1993. The phonetics of sound change. In Historical linguistics: problems and perspectives, ed. By Charles Jones, 237-278. Harlow: Longman.

Onishi, Kristine H., Kyle. E. Chambers and Cynthia Fisher. 2002. Learning phonotactic constraints from brief auditory experience. Cognition 83:B13-B23

Onnis, Luca, Padraic Monaghan, Korin Richmond, and Nick Chater. 2005. Phonology impacts segmentation in online speech processing. Journal of Memory & Language, 53, 225-237

Pearl, Lisa. 2008. Putting the Emphasis on Unambiguous: The Feasibility of Data Filtering for Learning English Metrical Phonology, BUCLD 32: Proceedings of the 32nd annual Boston University Conference on Child Language Development, Chan, H., Jacob, H., and Kapia, E (eds.), Somerville, MA: Cascadilla Press, 390-401.

Peperkamp, Sharon, Katrin Skoruppa and Emmanuel Dupoux. 2006. The role of phonetic naturalness in phonological rule acquisition. In Proceedings of the 30th Annual Boston University Conference on Language Development, ed. by D. Bamman, T. Magnitskaia & C. Zaller, 464-475. Sommerville: Cascadilla Press.

Prince, Alan and Paul Smolensky. 1993. Optimality Theory: Constraint interaction in generative grammar. Cambridge, MA: Blackwell.

Pulleyblank, Douglas. 2002. Harmony drivers: no disagreement allowed. In Proceedings of the Twenty-eighth Annual Meeting of the Berkeley Linguistics Society, 249-267. Berkeley Linguistics Society, Berkeley, California.

Pycha, Anne, Pawel Nowak, Eurie Shin and Ryan Shosted. 2003. Phonological rule-learning and its implications for a theory of vowel harmony. In Proceedings of the West Coast Conference on Formal Linguistics 22 ed. by G. Garding and M. Tsujimura, 423-435. Somerville, MA: Cascadilla Press.

R Development Core Team. 2009. R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Available from http://www.R-project.org

Rogers, James, Jeffrey Heinz, Gil Bailey, Matt Edlefsen, Molly Visscher, David Wellcome, and Sean Wibel. 2010. On languages piecewise testable in the strict sense. In The Mathematics of Language ed. by Christian Ebert, Gerhard Jäger, and Jens Michaelis,vol. 6149 of Lecture Notes in Artifical Intelligence, 255–265. NY: Springer.

Rogers, James and Geoff Pullum. 2011. Aural pattern recognition experiments and the subregular hierarchy. Journal of Logic, Language and Information 20:329-342.

Rose, Sharon and Rachel Walker. 2004. A typology of consonant agreement as correspondence. Language 80:475–531.

Sapir, Edward, and Harry Hoijer. 1967. The phonology and morphology of the Navaho language. Berkeley: University of California Press.

Seidl, Amanda, Alejandrina Cristià, Amelie Bernard, and Kristine H. Onishi. 2009. Allophonic and phonemic contrasts in infants' learning of sound patterns. Language Learning and Development 5(3):191-202.

Vogel, Irene. 2009. The status of the Clitic Group. In Phonological domains: Universals and deviations, ed. by Janet Grijzenhout and Baris Kabak, 15-46. Berlin / New York: Mouton de Gruyter.

Vogel, Irene. 2010. The phonology of compounds. In Current Issues in Linguistic Theory, 311: Cross-Disciplinary Issues in Compounding, ed. by Sergio Scalise and Irene Vogel, 145-164. Amsterdam / Philadelphia: John Benjamins.

Wilson, Collin. 2003. Experimental investigation of phonological naturalness. In Proceedings of the 22nd West Coast Conference on Formal Linguistics ed. by G. Garding and M. Tsujimura, 533-546. Somerville, MA: Cascadilla Press.

Appendix A

EXP 1: TRAINING ITEMS

Conditions	Items			
	ʃaʃɛkuʃ	sasɔkɔs	ʃakuʃuʃ	sakɛsas
	ʃɛʃukuʃ	sɛsukɔs	ʃɛkɛʃuʃ	sɛkɔsɛs
	ʃiʃɔkɛʃ	sisɛkɔs	ʃikiʃɔʃ	sikisis
	ʃɔʃɔkɛʃ	sɔsakɔs	ʃɔkaʃɛʃ	sɔkusas
SH	ʃuʃakiʃ	susukɔs	ʃukɔʃɔʃ	sukasus
	ʃaʃɔkuʃ	sasakus	ʃakuʃiʃ	sakisɔs
	ʃɛʃɔkiʃ	sɛsukis	ʃɛkɛʃɔʃ	sɛkɛsɛs
	ʃiʃukɔʃ	sisɔkɔs	ʃikɛʃaʃ	sikasɔs
	ʃɔʃɔkuʃ	sɔsikas	ʃɔkɛʃiʃ	sɔkɔsis
	ʃuʃakɛʃ	susɔkis	ʃukɔʃɛʃ	sukɛsas
	ʃasɔkuʃ	saʃakus	ʃaʃɛkuʃ	sasɔkɔs
	ʃɛsɔkiʃ	sɛʃukis	ʃɛʃukuʃ	sɛsukɔs
	ʃisukɔʃ	siʃɔkɔs	ʃiʃɔkɛʃ	sisɛkɔs
	ʃɔsɔkuʃ	sɔʃikas	ʃɔʃɔkɛʃ	sɔsakɔs
FL	ʃusakɛʃ	suʃɔkis	ʃuʃakiʃ	susukɔs
	ʃakusiʃ	sakiʃɔs	ʃakuʃuʃ	sakɛsas
	ʃɛkɛsɔʃ	sɛkɛʃɛs	ʃɛkɛʃuʃ	sɛkɔsɛs
	ʃikɛsaʃ	sikaʃɔs	ʃikiʃɔʃ	sikisis
	ʃɔkɛsiʃ	sɔkɔʃis	ʃɔkaʃɛʃ	sɔkusas
	ʃukɔsɛʃ	sukɛʃas	ʃukɔʃɔʃ	sukasus
	ʃasɔkuʃ	saʃakus	ʃakusiʃ	sakiʃɔs
	ʃɛsɔkiʃ	sɛʃukis	ʃɛkɛsɔʃ	sɛkɛʃɛs
	ʃisukɔʃ	siʃɔkɔs	ʃikɛsaʃ	sikaʃɔs
	ʃɔsɔkuʃ	sɔʃikas	ʃɔkɛsiʃ	sɔkɔʃis
IFL	ʃusakɛʃ	suʃɔkis	ʃukɔsɛʃ	sukɛʃas
	ʃasɛkuʃ	saʃɔkɔs	ʃakusuʃ	sakɛʃas
	ʃɛsukuʃ	sɛʃukɔs	ʃɛkɛsuʃ	sɛkɔʃɛs
	ʃisɔkɛʃ	siʃɛkɔs	ʃikisɔʃ	sikiʃis
	ʃɔsɔkɛʃ	sɔʃakɔs	ʃɔkasɛʃ	sɔkuʃas
	ʃusakiʃ	suʃukɔs	ʃukɔsɔʃ	sukaʃus

Appendix B

EXP 1: TEST ITEMS

FL/*SH vs. *FL/*SH		FL/SH vs. *FL/*SH		FL/SH vs. FL/*SH	
Word 1	Word 2	Word 1	Word 2	Word 1	Word 2
sɛkɔʃɔs	ʃɛkɔʃɔs	sukisas	sukisaʃ	ʃɔʃukɔʃ	ʃɔsukɔʃ
ʃasɔkaʃ	ʃasɔkas	ʃɛʃɛkaʃ	sɛʃɛkaʃ	sukɛsus	sukɛʃus
suʃɛkɔs	suʃɛkɔʃ	ʃɔkuʃiʃ	ʃɔkuʃis	sisakus	siʃakus
ʃikisaʃ	sikisaʃ	sisɔkus	ʃisɔkus	ʃakaʃɔʃ	ʃakasɔʃ
sɛʃɔkɔs	sɛʃɔkɔʃ	ʃɛkɛʃaʃ	ʃɛkɛʃas	ʃɔkuʃɔʃ	ʃɔkusɔʃ
ʃakɔsaʃ	sakɔsaʃ	ʃɔʃukiʃ	sɔʃukiʃ	susɛkus	suʃɛkus
sukɛʃɔs	ʃukɛʃɔs	sikɔsus	sikɔsuʃ	sikasus	sikaʃus
ʃisikaʃ	ʃisikas	susikas	ʃusikas	ʃaʃakɔʃ	ʃasakɔʃ
ʃikuʃis	sikuʃis	sɔkasiʃ	sɔkasis	sɔʃakas	sɔsakas
ʃɛsakis	ʃɛsakiʃ	ʃikɔʃis	ʃikɔʃiʃ	ʃusɛkiʃ	ʃuʃɛkiʃ
sɔkisɔʃ	ʃɔkisɔʃ	ʃɛsikɔs	sɛsikɔs	sakuʃɛs	sakusɛs
saʃɛkɛʃ	saʃɛkɛs	saʃikuʃ	ʃaʃikuʃ	ʃɛkɔsaʃ	ʃɛkɔʃaʃ
ʃɔsikɔs	ʃɔsikɔʃ	ʃɔsakis	sɔsakis	saʃukɛs	sasukɛs
siʃukiʃ	siʃukis	siʃɔkiʃ	ʃiʃɔkiʃ	ʃɛsɔkaʃ	ʃɛʃɔkaʃ
sɛkasiʃ	ʃɛkasiʃ	sɛkisɔʃ	sɛkisɔs	sɔkaʃas	sɔkasas
ʃakɛʃɛs	sakɛʃɛs	ʃakiʃus	ʃakiʃuʃ	ʃukɛsiʃ	ʃukɛʃiʃ

Appendix C

EXPERIMENT 2: Syntax Training Items

dubek pubim gatam petuk	dudos hahus budes gopis
gihok tegik pepek heguk	pahum dobes gohis tabom
biges podas tobus badis	tipus datem degom hetas
gipam hopos bopes tudim	tedum begek pebak bobum
dohum gipas tupes bubom	hahus tites pipis gutos
pugim hopak togok gahem	bepak potem hidim tuguk
hupok gipak guhik detuk	dobis pados bades togus
pobes duhak toguk hihis	gitis dupuk dehek tugos
pobik tihok dihek bogak	hagik pahus tubos gahek
hugik tedus bagas gupek	gepuk pugik gagak babok
tetik dudak hidek gobuk	hadak behom hedim gadek
hudas hodus hebos pipis	tehom gebem gagum dihim
bidom hugim detum petam	tiham potim gubom hipem
hopum bigem tadim togom	patek dotis hetas bugok
dages pubis gotos pibas	gepom bihus tagos pigam
paduk hugek buhik tatak	dupos gehak tubuk hedes
hoham dapom bopim didem	bepas tipom pebum didis
getek hahak hipik detok	betis gepum tatom bubas
dobis hadas bagus tohes	dahok bidum dibam tapik
putem datim gubam gotum	dabem dodok pituk bohim

Appendix D

EXPERIMENT 2: Syntax Test Items

NE	Full
dubek bigim tagam tabok	dobim bidom godem pigam
gedom datek hituk gutam	puduk hadak bahek tubok
tedak dupus buhes tutok	tigim padam hedem totom
puges dahom pegum gapis	gidak pubek tatuk topik
bipam bitis bogos badem	huhom tetim hegum tahem
hagem tihas dehis hetum	hupos titis botas detus
pabes dupak debuk bogos	pabek dodak gubok buguk
dohak puhus tupis bapek	tidos gipus topis botas
bipuk dohas tugos depek	hohus hages digas pegos
gepok gehum bobam gohek	duduk hohak dihik godek
bepas padok hepik pebus	tetis tepas gubus gupes
padam pupos pidis pigum	bipok bebik peguk pebak
dahuk gidom badim bugak	dobim huhom tatam budem
tehom pabik gobak dipem	hatek hogik pigok gubuk
padum hopos gupes gabam	bidos bidus tahes tagas
dages dabik totuk totas	podam dupom dipem hepim
hotes bepum gagam hibos	huhum popam debom badim
pogis habem dibam detos	dutim gipum topem hegom
poduk hudos bohis digak	bebis pubes hibus botos
gepom tigik tutak hibum	hatis teges hitos petus
hupus potek gatok gobas	getes dobes bubas gahis
hopak getem gabom dihik	bepuk potek petok hibak
puhas gehok dedek tubus	giham dohum butom gadim
gihak bedum tudim tohek	dagim hogem pibum buhim
gitis behom bagam pigus	tihok huduk gopek tadik

CD	SL3
duduk hohas dihik godes	tidos gipus topim botam
hagem tihas dehim hetus	puduk hadak bahes tubos
gepok gehum bobak gohem	bipos bebis peguk pebak
huhom tetis hegum tahes	gidam pubem tatuk topik
padam pupos pidim pigus	pabek dodak gubom bugum
puges dahom pegus gapim	tigim padam hedek totok
bipuk dohas tugok depes	tetik tepak gubus gupes
bipam bitis bogom bades	hupom titim botas detus
bepas padok hepis pebuk	hohus hages digam pegom
dobim huhos tatam budes	dobim bidom godes pigas
tedak dupus buhek tutos	dubek bigik tagam tabom
gedom datek hitum gutak	dohak puhuk tupis bapes
pabes dupak debus bogok	tihom hudum gopek tadik
puhas gehok dedes tubuk	dagik hogek pibum buhim
bidos biduk tahes tagak	hatik tegek hitos petus
dahuk gidom badik bugam	huhuk popak debom badim
hupus potek gatos gobak	tehom pabim gobak dipek
bepuk potem petok hibam	hates hogis pigok gubuk
poduk hudos bohik digas	padum hopom gupes gabas
hotes bepum gagas hibom	dages dabis totuk totak
getes dobem bubas gahim	dutis gipus topem hegom
giham dohuk butom gadik	pogis habes dibam detom
gihak bedum tudik tohem	hopak getek gabom dihim
gepom tigik tutam hibuk	podam dupom dipek hepik
gitis behom bagas pigum	bebis pubes hibuk botok

Appendix E

EXPERIMENT 2: Phonotactic Training Items

ʤushaʤish	sosisiz
ʤishaʤuʤ	susizis
ʤazasosh	shesheshosh
ʤuzaseʤ	shezasash
zaʤuʤoz	zoshuʤus
shosusoʤ	sizezaz
zesusez	shizoziʤ
sishiʤiz	ʤozizosh
shuʤoshosh	ʤaʤishush
sizasas	seshishez
shusisash	shushuʤash
ʤoʤoʤiʤ	ʤiʤishuʤ
ʤisuzoʤ	sheʤoʤish
zazozaz	zasizes
zoʤushez	zushishez
siʤeʤis	shaʤaʤuʤ
seʤeshas	shusazash
ʤesizash	sosheshis
zosizuz	zeʤashoz
sasesus	zazesiz

Appendix F

EXPERIMENT 2: TEST ITEMS

NE	Full	CD	SL3
shazasash	zisozuz	sadʒozidʒ	dʒushesos
shasosodʒ	shudʒodʒidʒ	shisidʒez	shushosaz
soshashus	dʒadʒodʒedʒ	shasodʒos	sesishush
shesezesh	sozozoz	dʒosashes	seseshush
sushidʒus	dʒudʒadʒosh	sushisudʒ	dʒadʒoses
dʒosasesh	shushodʒash	zudʒesadʒ	dʒudʒasos
zudʒedʒas	sesesuz	soshasush	shudʒosiz
dʒosisadʒ	zazezos	seshusish	zisoshudʒ
shuzozush	zezizos	sadʒuzesh	sozoshosh
sadʒodʒiz	sesizuz	shazashas	shashizos
shisizedʒ	zisuzes	shuzoshuz	zezishosh
seshushis	dʒushedʒodʒ	sheseshez	zazeshodʒ
sadʒushez	shashidʒosh	dʒasushas	zesudʒosh
zidʒoshos	dʒudʒushosh	dʒosidʒas	shedʒozas
dʒuzesedʒ	zesusos	dʒozadʒaz	sozedʒudʒ
dʒisizash	sesusas	zishezodʒ	sesudʒadʒ
dʒozazadʒ	dʒidʒushidʒ	shuzadʒus	shoshusaz
dʒasusash	zasisaz	zidʒosedʒ	dʒishuzes
dʒozezadʒ	sozusis	zidʒososh	sozudʒidʒ
shuzasudʒ	zozosaz	sidʒazish	dʒudʒuzoz
zeshashuz	shoshushash	dʒozedʒaz	zasidʒadʒ
zidʒodʒes	sheshashidʒ	zishazedʒ	dʒidʒuziz
zishadʒez	shedʒoshash	zeshazush	zozodʒash
zishedʒoz	dʒashashesh	dʒuzedʒes	dʒashazez
sidʒashiz	sozesus	dʒisishaz	sheshazis

Appendix G

IRB APPROVAL LETTERS

UNIVERSITY OF DELAWARE

RESEARCH OFFICE

210 Hullihen Hall
University of Delaware
Newark, Delaware 19716-1551
Ph: 302/831-2136
Fax: 302/831-2828

DATE: September 24, 2012

TO: Yee King Lai
FROM: University of Delaware IRB

STUDY TITLE: [245773-3] Learning Non-regular Syntactic Patterns

SUBMISSION TYPE: Continuing Review/Progress Report

ACTION: APPROVED
APPROVAL DATE: September 24, 2012
EXPIRATION DATE: September 23, 2013
REVIEW TYPE: Expedited Review

REVIEW CATEGORY: Expedited review category # 7

Thank you for your submission of Continuing Review/Progress Report materials for this research study. The University of Delaware IRB has APPROVED your submission. This approval is based on an appropriate risk/benefit ratio and a study design wherein the risks have been minimized. All research must be conducted in accordance with this approved submission.

This submission has received Expedited Review based on the applicable federal regulation.

Please remember that informed consent is a process beginning with a description of the study and insurance of participant understanding followed by a signed consent form. Informed consent must continue throughout the study via a dialogue between the researcher and research participant. Federal regulations require each participant receive a copy of the signed consent document.

Please note that any revision to previously approved materials must be approved by this office prior to initiation. Please use the appropriate revision forms for this procedure.

All SERIOUS and UNEXPECTED adverse events must be reported to this office. Please use the appropriate adverse event forms for this procedure. All sponsor reporting requirements should also be followed.

Please report all NON-COMPLIANCE issues or COMPLAINTS regarding this study to this office.

Please note that all research records must be retained for a minimum of three years.

Based on the risks, this project requires Continuing Review by this office on an annual basis. Please use the appropriate renewal forms for this procedure.

- 1 - Generated on IRBNet

If you have any questions, please contact Jody-Lynn Berg at (302) 831-1119 or jlberg@udel.edu. Please include your study title and reference number in all correspondence with this office.

University of Delaware
Research Office
210 Hullihen Hall
University of Delaware
Newark, Delaware 19716-1551
Ph: 302/831-2136
Fax: 302/831-2828

DATE:	March 12, 2012
TO:	Yee King Regine Lai
FROM:	University of Delaware IRB
STUDY TITLE:	[165440-5] Learning Long-Distance Phonotactic Patterns
SUBMISSION TYPE:	Continuing Review/Progress Report
ACTION:	APPROVED
APPROVAL DATE:	March 12, 2012
EXPIRATION DATE:	April 11, 2013
REVIEW TYPE:	Expedited Review
REVIEW CATEGORY:	Expedited review category # 7

Thank you for your submission of Continuing Review/Progress Report materials for this research study. The University of Delaware IRB has APPROVED your submission. This approval is based on an appropriate risk/benefit ratio and a study design wherein the risks have been minimized. All research must be conducted in accordance with this approved submission.

This submission has received Expedited Review based on the applicable federal regulation.

Please remember that informed consent is a process beginning with a description of the study and insurance of participant understanding followed by a signed consent form. Informed consent must continue throughout the study via a dialogue between the researcher and research participant. Federal regulations require each participant receive a copy of the signed consent document.

Please note that any revision to previously approved materials must be approved by this office prior to initiation. Please use the appropriate revision forms for this procedure.

All SERIOUS and UNEXPECTED adverse events must be reported to this office. Please use the appropriate adverse event forms for this procedure. All sponsor reporting requirements should also be followed.

Please report all NON-COMPLIANCE issues or COMPLAINTS regarding this study to this office.

Please note that all research records must be retained for a minimum of three years.

Based on the risks, this project requires Continuing Review by this office on an annual basis. Please use the appropriate renewal forms for this procedure.

If you have any questions, please contact Jody-Lynn Berg at (302) 831-1119 or jlberg@udel.edu. Please include your study title and reference number in all correspondence with this office.