
Theses and Dissertations

Spring 2013

The effect of referent similarity and phonological similarity on concurrent word learning

Libo Zhao
University of Iowa

Copyright 2013 Libo Zhao

This dissertation is available at Iowa Research Online: <http://ir.uiowa.edu/etd/2419>

Recommended Citation

Zhao, Libo. "The effect of referent similarity and phonological similarity on concurrent word learning." PhD (Doctor of Philosophy) thesis, University of Iowa, 2013.
<http://ir.uiowa.edu/etd/2419>.

Follow this and additional works at: <http://ir.uiowa.edu/etd>

 Part of the [Psychology Commons](#)

THE EFFECT OF REFERENT SIMILARITY AND PHONOLOGICAL SIMILARITY
ON CONCURRENT WORD LEARNING

by
Libo Zhao

An Abstract

Of a thesis submitted in partial fulfillment
of the requirements for the Doctor of
Philosophy degree in Psychology
in the Graduate College of
The University of Iowa

May 2013

Thesis Supervisor: Associate Professor Prahlad Gupta

ABSTRACT

Similarity has been regarded as a primary means by which lexical representations are organized, and hence an important determinant of processing interactions between lexical items. A central question on lexical-semantic similarity is how it influences lexical processing. There have been fewer investigations, however, on how lexical-semantic similarity might influence novel word learning. This dissertation work aimed to fill this gap by addressing one kind of lexical-semantic similarity, similarity among the novel words that are being learned concurrently (concurrent similarity), on the learning of phonological word forms. Importantly, it aimed to use tests that eliminated real time processing confounds at test so as to provide convincing evidence on whether learning was indeed affected by similarity.

The first part of the dissertation addressed the effect of concurrent *referent* similarity on the learning of phonological word forms. Experiment 1 used a naming test to provide evidence on the direction of the effect. Experiment 2 and Experiment 3 used a stem completion test and a recognition from mispronunciation test that controlled for real time processing between conditions. Then a 4-layer Hebbian Normalized Recurrent Network was also developed to examine mechanistically whether learning was affected in behavioral simulations. Consistently across the three experiments and the simulation, a detrimental effect of referent similarity on the phonological word form learning was revealed.

The second part of the dissertation addressed the effect of *phonological* similarity, and specifically, cohort similarity, on the learning of phonological word forms. A test of recognition from mispronunciation of partial words was developed to control for real time processing between conditions so as to isolate the effect of learning. We examined the effect of cohort similarity at different syllable positions and found a detrimental effect at the second syllable and non-effect at the third syllable. This is consistent with the

previous finding that competition among cohorts diminishes as the stimulus is received, suggesting that the effect of cohort similarity depends on the status of competition dynamics among cohorts.

The theoretical and methodological implications of this study are discussed.

Abstract Approved: _____
Thesis Supervisor

Title and Department

Date

THE EFFECT OF REFERENT SIMILARITY AND PHONOLOGICAL SIMILARITY
ON CONCURRENT WORD LEARNING

by
Libo Zhao

A thesis submitted in partial fulfillment
of the requirements for the Doctor of
Philosophy degree in Psychology
in the Graduate College of
The University of Iowa

May 2013

Thesis Supervisor: Associate Professor Prahlad Gupta

Copyright by

LIBO ZHAO

2013

All Rights Reserved

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

PH.D. THESIS

This is to certify that the Ph.D. thesis of

Libo Zhao

has been approved by the Examining Committee
for the thesis requirement for the Doctor of Philosophy
degree in Psychology at the May 2013 graduation.

Thesis Committee: _____
Prahlad Gupta, Thesis Supervisor

Bob McMurray

Cathleen Moore

Shaun Vecera

Bruce Tomblin

ACKNOWLEDGMENTS

Here I sincerely thank all the people who have helped me during my graduate school years.

I want first to thank my advisor, Prahlad. Thank you very much for your guidance on almost every aspect of research, from the literature review to manuscript writing, from generating an idea to making experimental designs, and from running an experiment to data analysis, and so on and so forth. Without your advices in the past years, I would not have known any more about research than what I knew when I enrolled in this program in 2006. Thank you very much also for your patience, and for giving me opportunities to explore and to make mistakes.

I also want to thank Bob. Thank you very much for sharing your ideas and providing inputs, which has made a difference to the dissertation work reported here and to my understanding of related fields. Thank you very much also for allocating times to answer my questions, and for your patience in checking my Matlab code every time when I made changes.

Thanks to Shaun. Thank you very much for your guidance and support on the attention project. Without these it would not have been possible at all for me to carry out any research on vision, which has expanded my horizon. Thank you very much also for your help in writing the manuscript on visual attention.

Thanks also to Cathleen, Bruce, and Larissa. Thank you very much for your suggestions and comments as my dissertation and/or research committee, which have helped me to be on track and to make steady progress in my research.

I also want to thank all the professors for teaching great classes that have enriched and shaped my mind, including Prof. Prahlad Gupta, Prof. Bob McMurray, Prof. Larissa Samuelson, Prof. Andrew Hollingworth, Prof. Cathleen Moore, Prof. Karla McGregor,

Prof. Ed Wasserman, Prof. Scott Robinson, Prof. Eliot Hazeltine, and Prof. Daniel Tranel.

Many thanks also to my lab mates, Brandon Abbs, Stephanie Packard, Effie Kapnoula, and Francis Smith, for your help with shaping my ideas and fruitful discussions. Many thanks to Jamie Tisdale for being a great help with computers and technical assistance.

I'd also love to thank some other graduate students who helped me in various ways, including Joe Toscano, Keith Apfelbaum, Lynn Perry, Kimberly Halvorson, Zhen Wu, Sarah Kucker, Josh Cosman, Ian Rasmussen, Hyunkyu Lee, Julie Hsuh, and Weiwei Zhang.

Thanks also to the many practicum students who helped me with data collection and data transcription and coding over the years.

Lastly, thanks also to my family for great support.

TABLE OF CONTENTS

LIST OF TABLES	vi
LIST OF FIGURES	vii
CHAPTER	
CHAPTER 1 INTRODUCTION & OVERVIEW	1
CHAPTER 2 A CONCEPTUAL MODEL ON NOVEL WORD LEARNING	7
2.1 The model	7
2.2 Similarity and its effect on the phonological word form learning	9
2.3 How to measure the phonological word form learning?	11
2.4 Theoretical significance of the dissertation study	13
CHAPTER 3 SEMANTIC SIMILARITY ON SPOKEN WORD PROCESSING AND LEARNING	16
3.1 The Effects of Semantic Similarity in Spoken Word Processing	16
3.1.1 The effect of semantic neighbors	16
3.1.2 Semantic priming	17
3.1.3 Semantic interference and retrieval-induced forgetting	18
3.1.4 Summary	19
3.2 The Effect of Semantic Similarity on Word Learning	20
3.2.1 The effect of semantic neighbors	20
3.2.2 Similarity in associative memory	21
3.2.4 Summary	22
3.3 Overall Summary	22
CHAPTER 4 THE BEHAVIORAL STUDY ON THE EFFECT OF REFERENT SIMILARITY ON CONCURRENT NOVEL WORD LEARNING	24
Experiment 1	26
Methods	27
Results and Discussion	30
Experiment 2	32
Methods	33
Results and Discussion	34
Experiment 3	36
Methods	36
Results and Discussion	38
General Discussion	39
CHAPTER 5 THE SIMULATION STUDY ON THE EFFECT OF CONCURRENT REFERENT SIMILARITY ON WORD LEARNING	42
Architecture of the model	43
Processing in the model	44

Learning in the model.....	47
Testing the model	47
Results and Discussion	48
General Discussion on Chapter 4 and Chapter 5	54
CHAPTER 6 PHONOLOGICAL SIMILARITY AND SPOKEN WORD PROCESSING AND LEARNING.....	57
6.1 The Effects of Phonological Similarity in Spoken Word Processing.....	57
6.1.1 The effect of phonological neighbors.....	57
6.1.2 Phonological similarity effect in immediate serial recall	59
6.1.3 Summary.....	61
6.2 The Effect of Phonological Similarity on Word Learning	62
6.2.1 The effect of phonological neighborhood density and phonotactic probability.....	62
6.2.2 The effects of different types of phonological neighbors.....	65
6.2.3 Infants' word learning using minimal pairs.....	66
6.2.4 Summary on the word learning studies	70
CHAPTER 7 THE EFFECT OF PHONOLOGICAL SIMILARITY ON CONCURRENT NOVEL WORD LEARNING	72
Experiment 4.....	75
Methods	76
Results and Discussion	78
Experiment 5.....	80
Methods	81
Results and Discussion	81
General Discussion	83
CHAPTER 8 GENERAL DISCUSSION AND FUTURE DIRECTIONS	87
APPENDIX A: EXAMPLES OF STIMULI USED IN THE DISSERTATION STUDIES.....	91
APPENDIX B: PARAMETERS USED IN THE 4-LAYER HNR NETWORK IN CHAPTER 5	94
APPENDIX C: FORMULAS USED IN THE 4-LAYER HNR NETWORK.....	95
REFERENCES	98

LIST OF TABLES

Table

Table A-1 Examples of Nonword Stimuli (IPA) Used in the Behavioral Experiments in Chapter 4	91
Table A-2 Examples of Nonword Stimuli (IPA) Used in the Behavioral Experiments in Chapter 7	92
Table B-1 Free parameters in the 4-layer HNR network	94

LIST OF FIGURES

Figure	
Figure 2-1. Examples of random shapes used in Experiment.1. The first row shows the low similarity condition, and the second row shows high similarity condition.....	9
Figure 4-1. Examples of random shapes used in Experiment.1. The first row shows the low similarity condition, and the second row shows high similarity condition.....	27
Figure 4-2. Naming accuracy for the three similarity conditions across seven tests.....	31
Figure 4-3. Stem completion accuracy for the Low- and the High-Similarity conditions.....	35
Figure 4-4. Recognition accuracy for the Low- and the High-Similarity conditions.....	39
Figure 5-1. The architecture of a typical HNRN model.....	44
Figure 5-2. Structure of the 4-layer network. Inputs consist of visual features and word-forms.....	45
Figure 5-3. Proportion correct on naming task for the Low- and High-Similarity conditions simulated by the 4-layer network.....	50
Figure 5-4. Proportion correct on stem completion task for the Low- and High-Similarity conditions simulated by the network.....	50
Figure 5-5. The strength of the connections between the auditory and lexical nodes for the Low- and High-Similarity conditions in the model.....	51
Figure 5-6. The strength of the connections between the visual feature and the visual category nodes for the Low- and High-Similarity conditions in the model.....	53
Figure 5-7. The strength of the connections between the visual category and the lexical nodes for the Low- and High-Similarity conditions in the model.....	54
Figure 7-1. Recognition accuracy in the cohort and the dissimilar conditions.....	79
Figure 7-2. Naming accuracies across the first three epochs in the two conditions.....	80
Figure 7-3. Recognition accuracy in the cohort and the dissimilar conditions.....	82
Figure 7-4. Naming accuracy in the cohort and the dissimilar conditions.....	83
Figure A-1. Examples of alien images used as referents in the experiments reported in Chapter 7.....	93

CHAPTER 1 INTRODUCTION & OVERVIEW

A word consists of a phonological word form (the sound pattern of the word), its referent (meaning; semantics) and associations or links between the two. Often, the word form is regarded as the *label* for the referent. Thus, word learning entails learning three integral aspects of a word: word form, referent, and links between them. This dissertation focuses on the learning of phonological word forms, not on the learning of semantics or on the learning of associations between a word form and its semantics. Focusing on just the phonological aspect can keep the scope of the dissertation manageable, but more importantly, it will contribute to our understanding of an aspect of word learning that is relatively less studied. In particular, this dissertation focuses on word form learning in the situation where a set of novel words learned in the same task (concurrent word learning) are similar in semantics or phonological word forms.

Similarity has been an important variable in the field of language. It has been regarded as a primary dimension by which lexical representations are organized, and hence an important determinant of processing interactions between lexical items (Luce & Pisoni, 1998; Marslen-Wilson, 1989). There is considerable evidence of how lexical processing of a particular word is influenced by that word's phonological neighbors (e.g., Luce & Pisoni, 1998; Dahan, Magnuson & Tanenhaus, 2001; Magnuson, Dixon, Tanenhaus, & Aslin, 2007; Vitevitch, 2002a, 2002b; Vitevitch & Luce, 1998, 1999) and semantic neighbors (e.g., Mirman, 2011; Mirman & Magnuson, 2008, 2009). However, less is known about whether lexical-semantic similarity also influences novel word *learning*. It is a theoretically important question to ask whether the effect of lexical-semantic similarity impacts only real-time lexical processing or whether it also has a long-term effect on learning. This dissertation will focus on the effect of such similarity

on novel word learning with the aim of contributing to understanding of this theoretical issue.

While recent studies have begun to provide evidence suggesting that the similarity of a novel word to known words (in either phonological word form or semantics) may influence its acquisition (e.g., Hoover, Storkel, Hogan, 2010; Storkel & Adolf, 2009; Storkel, Armbruster, & Hogan, 2006; Tomasello, Mannle, & Werdenschlag, 1988), very little is known about the effect on word learning of a different kind of similarity: similarity *to each other*, among novel words that are being learned at roughly the same time. This is a situation that is often encountered by word learners in the real world. For example, even for infants and young children, multiple words being acquired during the same period of time are often similar in semantics, such as *cat* and *dog*, or similar in phonological word forms, such as *cat* and *cap* (Regier, 2005). The present dissertation examines how the learning of the phonological word forms might be affected by concurrent semantic similarity and concurrent phonological similarity respectively, and thus will contribute to a more comprehensive understanding of word learning.

How might the similarity in referents and the similarity in phonological word forms have an effect on concurrent word learning? Although the specific underlying mechanisms might be different for concurrent referent similarity and concurrent phonological similarity, one common mechanism might be the co-activation and competition processes. That is, similarity in word forms or in referents leads to simultaneous activation of the similar words (e.g., Allopenna, Magnuson, & Tanenhaus, 1998; Huettig & Altmann, 2005; McClelland & Elman, 1986). On one hand, this simultaneous activation naturally leads to a blended and ambiguous representation of words in the working memory (Spivey, 2007), which might then lead to laying down

inaccurate representation of novel words that are being learned. Moreover, the simultaneously active words may compete and weaken the representation of each other (e.g., Magnuson, Tanenhaus, Aslin, & Dahan, 2003), which may also hurt learning. On the other hand, simultaneous activation of words also means that the shared information is strengthened (Gupta, Lipinski, & Aktunc, 2005; Vitevitch, 2002a), which may potentially facilitate learning. Therefore, lexical-semantic similarity may lead to two effects in opposite directions, and the overall observed effect can be in either direction depending on the relative size of the two. Thus, as will be discussed further when discussing each of referent similarity and phonological similarity in greater detail, the direction of the effects of concurrent similarities on learning is not clear.

The first part of the dissertation addresses how concurrent *referent* similarity may influence phonological word form learning. In addition to the significance of studying the effect of similarity on word learning in general as discussed above, this question is also important in that it examines the effect of properties of one aspect of a word (referent similarity) on the learning of the other aspect of a word (phonological label), and thus speaks to the theoretical debate on modularity and interactivity between the components of a word. Although there is ample evidence that suggests an interaction between semantics and phonological word forms in known word processing (e.g., Dell, 1986; McClelland & Rumelhart, 1986; Mirman & Chen, 2010; Peterson & Savoy, 1998; Rapp & Goldrick, 2000), little is known about whether semantics also influences the *learning* of phonological word forms. This dissertation work will thus extend our understanding of this issue by focusing on novel word learning.

The second part of the dissertation addresses how concurrent *phonological* similarity may influence phonological word form learning. In thinking about how concurrent phonological similarity might affect learning, an important issue to consider is that the information of a phonological word form is available incrementally over time as

the word unfolds, and consequently its competitor dynamics changes moment-by-moment. Consistent with this, previous studies have revealed that the effect of phonological similarity is different depending on where the similarity lies and it is also different from moment to moment depending on how the competition dynamics unfolds over time (e.g., Allopenna et al., 1998; Magnuson et al., 2007). Therefore, this study will focus on only one kind of similarity, cohort similarity (sharing the word onset) to get a clearer picture of the effect. Since a given word's cohort neighbors become active and compete strongly at an early point of stimulus presentation and gradually die out when more information is received, it can be predicted that any detrimental effect of cohort similarity on learning will decrease over time. Therefore, this dissertation study will examine the effect of cohort similarity separately at different syllable positions. It will thus provide novel evidence to how the temporal evolution of competition dynamics from phonologically similar words may influence novel word *learning* differently at different time points.

This dissertation also emphasizes a methodological issue: To more definitively assess whether referent and/or phonological similarity indeed have an influence on phonological word form learning, new behavioral measures need to be designed. As in measuring any other psychological phenomenon, a behavioral task is what is relied on to make inferences about the status of learning. However, a behavior/performance in a given task is always “multiply determined” (Gathercole, 2006; Gupta, 2008; Gupta & Tisdale, 2009). In the situations where learning is involved, real time processes and long-term learning jointly contribute to task performances (Gupta & Tisdale, 2009; McMurray, Horst, & Samuelson, 2012; Samuelson, Schutte, & Horst, 2009), and thus correct inferences about learning cannot be made without taking into account the nature of the online processing in the task. Because of this, the present dissertation will make special effort to design behavioral measures that can separate effects of real time processes from

effects of learning. In fact, the existing studies that addressed the role of similarity in word forms or in semantics on novel word learning all failed at doing this (e.g., Creel & Dahan, 2010; Hoover et al., 2010; Stager & Werker, 1997; Storkel & Adolf, 2009; Storkel et al., 2006). For example, Storkel & Adolf (2009) concluded that semantic similarity impaired novel word learning based on a detrimental effect in naming. However, it was possible that the poorer naming observed in the high similarity condition was not due to learning (a “better” representation being built), but instead driven by the greater competition from these semantic neighbors that were also active in the moment of testing (via the naming task). Bearing this in mind, this dissertation takes special effort to design tests that control for real-time processing at test across conditions so as to draw convincing conclusions on whether learning is affected by word form similarity and referent similarity. However, isolating learning and on-line processing completely is almost impossible using behavioral tasks. Therefore, we also use computational modeling to provide a pure measure of learning that is independent of real-time processing. A four-layer Hebbian Normalized Recurrent Network is developed to simulate the effects of concurrent referent similarity on novel word learning. (Due to scope limitations, no simulation is carried out for the study on concurrent phonological similarity.)

This dissertation is organized as follows. Chapter 2 will set up a conceptual model on novel word learning, which will be used as the basis to understand the effect of concurrent similarities on the learning of phonological word forms, and also to talk about the theoretical significances of the dissertation study. Chapter 3 will survey the existing studies on how similarity in semantics might influence spoken word processing and novel word learning. Chapter 4 will present three behavioral experiments examining how concurrent referent similarity might influence the learning of a set of words concurrently. Chapter 5 will report the simulation study on the effect of concurrent referent similarity on word learning. Chapter 6 will survey the existing studies on how similarity in

phonological word forms might influence spoken word processing and novel word learning. Chapter 7 will present two behavioral experiments on how phonological similarity might influence the learning of a set of words concurrently. Chapter 8 will be a general discussion on the two dissertation studies, including the contributions, limitations and future directions.

CHAPTER 2 A CONCEPTUAL MODEL ON NOVEL WORD LEARNING

This chapter aims to delineate a conceptual model on novel word learning, so as to provide a framework to predict how concurrent referent similarity and concurrent phonological similarity might influence the learning of the word forms, and also to discuss the theoretical implications of the dissertation work. Due to these purposes of this, it is sufficient to build a box-and-arrow model without specifying the representations and the mathematical algorithms for the change of activation and learning in detail. However, the model does borrow the concepts in the connectionist framework, including connection weights, activation and activation spreading. The speculation of the real-time processes involved in this model, including co-activation, competition and activation spreading, is based on the empirical findings on the effect of lexical-semantics similarity in lexical processing (will be reviewed in Chapter 3 and Chapter 6). The involvement of these processes is further supported because a number of neural networks that implemented them are able to simulate a wide range of behavioral phenomena in lexical processing and word learning (e.g., Dell, 1986; McClelland & Elman, 1986; McMurray, Horst, & Samuelson, 2012; Mirman & Chen, 2010; Rapp & Goldrick, 2000).

In the following, the basic components and the architecture of the model will be described first, based on which the effect of concurrent referent similarity and concurrent phonological similarity on novel word learning, the issue of measuring the *learning* of novel words, and the theoretical implications of the dissertation study will be discussed.

2.1 The model

The basic components of this model are shown in Figure 2.1. In the bottom there are the input components, with one for the phonological word form and the other for the referent. On the top there is the internal representation of the word that combines the

word form and the referent. In addition, there are two components in the middle that refer to the intermediate representation of the word forms and the intermediate representation of the referents before they are combined as a single word. The intermediate representation of the referents can be thought of as the representation of the visual categories that the raw referent inputs belongs to, and the intermediate representation of the word forms can be conceptualized as some kind of abstract representation of the word forms. Someone might argue that these intermediate layers are not necessary ones computationally, however, they are useful constructs at the description level to talk about the issues involved in this dissertation.

In this model, activation spreads in both directions between every two adjacent components via the bi-directional links, and the flow of activation is proportional to the strengths of the links (i.e., connection weights). The activation spreading occurs in real time, and thus activation (information) at any component of the model will be able to influence all the rest components over a certain number of processing cycles. Another processing characteristics of the model is that things that are co-activate inside a given component (e.g., multiple visual categories, multiple words) compete with each other, suppressing the activation growth of each other.

The model learns words by adjusting the strength of the connection weights between the adjacent layers. We can say a word is learned when the right set of connection weights are acquired, which enables the correct representation of the visual category, of the word form, and that of the word to be activated by a given pair of the phonological input and the referent input.

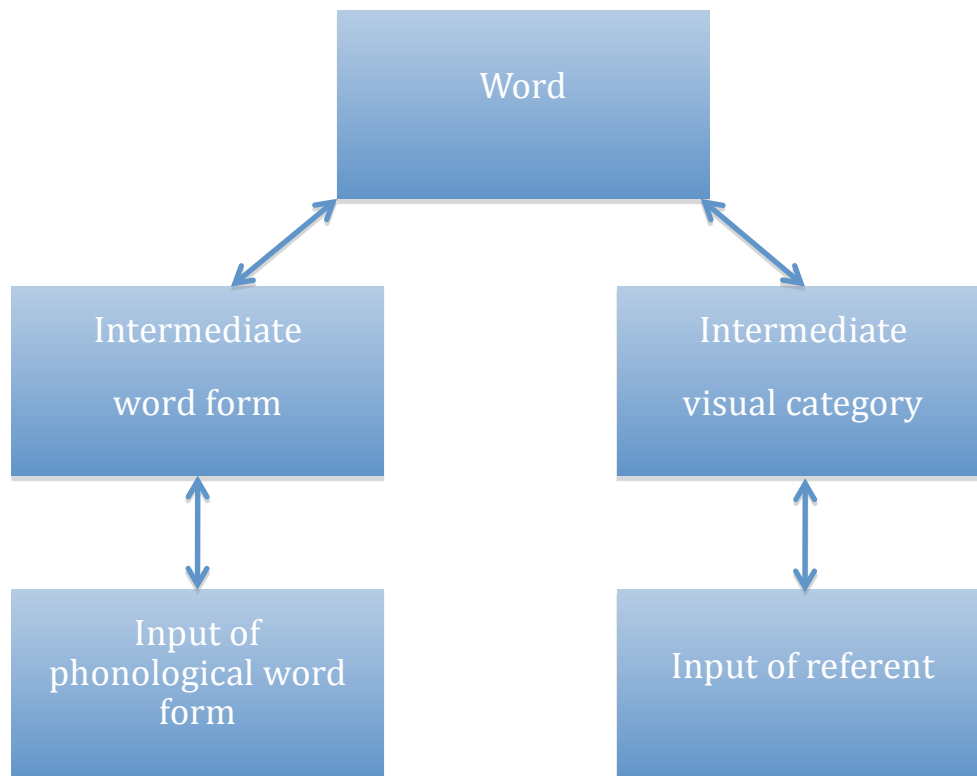


Figure 2.1 The conceptual model of word learning.

2.2 Similarity and its effect on the phonological word form

learning

Concurrent referent similarity and concurrent phonological similarity can be represented as overlaps in the referent input or the phonological input. Since the focus of this dissertation is on the learning of the phonological word forms, in the following I will

analyze how this aspect of word learning might be influenced by these two kinds of similarities respectively in the framework of the model. The predictions below are based on the real time processes implemented in the model, such as co-activation, competition, and activation spreading, combined with the simplest form of associative learning, Hebbian Learning.

First, suppose a set of novel words that are similar in referents are being learned concurrently, for example, words referring to similar-looking animals such as *dog*, *cat*, and *rat*. Since the input of a given referent (e.g., *dog*) overlaps with those of the other referents, multiple visual categories will be co-active at the intermediate layer, making categorizing the visual referent difficult. However, the ambiguity does not stop here. In fact, it will be spread throughout of the model, first from the visual category level to the word level, then to the intermediate word form level and then down to the phonological input level. According to the Hebbian learning rule, things that are active together are wired together, thus undesired connections, those between the correct information at each layer and the incorrect information at the adjacent layer(s), will be strengthened in the model. In addition, things that are co-active inside each component will also compete with each other, and thus the correct information will become less active due to receiving lateral inhibition from the co-active information. Since the increment of the connection strength is proportional to the activation level of the information, the growth of the desired associations in this condition will be reduced. In the model, the learning (knowledge) of the word forms can be indexed by the connections between the phonological input layer and the intermediate phonological word form layer, and the connections between the intermediate phonological word form layer and the word layer. Obviously as analyzed above, it is possible that referent similarity may impair the learning of the word forms, by resulting in less growth of the phonological weights and by leading to spurious associations.

However, the effect can also be in the opposite direction. As will be introduced in the following chapter, overlaps between similar referents can form a chunk (e.g., Goldstone, 2000; Schyns & Rodet, 1997). This will highlight the unique feature of each individual referent, and will make it easier to discriminate among the referents and categorize them (e.g., Tomasello, Mannle, & Werdenschlag, 1988). This means that in the model, the co-activation of the non-target categories will be largely reduced at the visual category level, and so will competition. As discussed above, the activation status of the visual layers will be cascaded to the rest of the model, and thus the word-form related components will also be relatively free of co-activation and competition. This means that the learning of the word forms will be facilitated, as indexed by the faster growth of the desired phonological weights will be faster and less spurious connections.

Therefore, the effect of referent similarity on phonological word form learning can be either facilitative or detrimental, and it is not obvious what the overall direction will be. The same is true for the effect of phonological similarity on the learning of the word forms. On one hand, sharing phonological elements in the word forms will be detrimental to learning, because it leads to ambiguous activation and competition in the intermediate phonological word form layer, which is further cascaded to the entire model. On the other hand, it can be facilitative because the shared part can be used as a chunk, which can reduce the level of co-activation and competition in the model and thus facilitate learning.

2.3 How to measure the phonological word form learning?

Here using this conceptual model I argue that many traditional tests are not adequate to measure the learning of the word forms. I will illustrate this point by using the naming test as an example. Suppose the question of interest is whether referent similarity among a set of words that are being learned concurrently has any impact on the learning of the word forms of these words, and suppose that a detrimental effect is

revealed. However, as will be illustrated below, the real time processing difference between the different similarity conditions in the moment of the naming test can simply explain the effect. When a referent that is similar to other referents is presented to probe for its label, it will activate multiple visual categories in the category level, which then propagate throughout the model, leading to ambiguous activation at all levels, including those of the phonological word forms, where the naming response will be based on. In addition, since the co-active things at each level will compete with each other, the correct information that will help retrieve the desired word form will become weaker due to being inhibited. In contrast, when the referent that is presented is not similar to other ones, there will be less ambiguous activation at the visual category level and elsewhere in the model. And due to free of competitors, the desired information that will support the retrieval of the correct word form will not be suppressed much. Taken together, even though the word forms have been learned to the same extent in the two conditions, having more ambiguous activation and competition will be sufficient to drive poorer naming in the former than in the latter.

It will be the same situation when the question of interest is whether concurrent phonological similarity affects the learning of word forms. In the condition where the word forms are similar to each other, when a referent is presented to probe for the label, the other labels that are similar to it will also be partially active and compete with it. In contrast, in the condition where the word forms are dissimilar to each other, the co-activation of other labels will be less likely to occur and thus the target word form will receive less competition from the other labels. Therefore, having more ambiguous activation and competition will be sufficient to drive poorer naming in the former condition than in the latter one, and the difference in naming performance does not necessarily speak to how the learning of the word forms has been affected by the phonological similarity.

Therefore, in order to tap the effect of concurrent similarity on the learning of word forms, tests that control for real time processing difference between different similarity conditions need to be used. And this is the methodological goal of this dissertation study.

2.4 Theoretical significance of the dissertation study

The first dissertation question addresses the effect of concurrent referent similarity on the learning of the phonological word forms. Thus, it asks whether the property of one part of a word (i.e., the referent) influences the learning of another part of a word (i.e., the word form). If these two components of a word are separate modules that do not talk to each other, then it would be predicted that it is not possible to have this hypothetical effect. This view is supported by the neuropsychological findings that knowledge of word forms or knowledge of semantics can be selectively impaired by focal brain lesions (e.g., Hodges, Patterson, Oxbury, & Funnell, 1992; Tranel, 2006). However, some studies support the view that there is cascaded interaction between semantics and phonological word forms. For example, it was revealed that a prime facilitated the processing of the phonological neighbor of its semantic neighbor (e.g., *sheep* primed *goal*, a phonological neighbor of *goat*) (Peterson & Savoy, 1998). For another example, analysis in spontaneous and experimentally elicited speech errors revealed a higher rate of mixed errors (i.e., the substitute of the target word that is both semantically and phonologically similar to the target) than would have been predicted by a purely feedforward model (for a review, see Rapp & Goldrick, 2000). If during learning semantics and phonological word forms of novel words also interact with each other in real time as these studies suggest, it would be predicted that the properties of semantics should have an impact on the learning of the word forms. This is suggested by the conceptual model described above, because referent similarity leads to ambiguous activation and competition everywhere of the model via cascaded interaction (some

existing theoretical models made similar suggestions, e.g., Gupta, 1996, 2009, Gupta & MacWhinney 1997; Plaut, 1997).

However, as novel word learning are being learned initially, the representations of semantics and the representations of the associative links between semantics and word forms are relatively weak, and thus it is not certain whether semantics could send enough information to have an impact on the learning of the word forms. This is hinted by some existing studies. For example, Plaut & Shallice (1993) used an attractor network and showed that lesions in the part of orthography also led to semantic errors, but not to the same extent as lesions in the part of the semantics itself did, suggesting that the amount of influence from one part of a word to the other is limited, even for known words. Therefore, it is still an open question whether the interaction is strong enough to enable semantics to influence the learning of the word forms even though the semantics and the word forms do interact during novel word learning.

In sum, this first part of the dissertation will inform the interactivity between semantics and word forms. If it turns out that referent similarity has an impact on the learning of the phonological word forms, it will suggest that these two components of a word do interact with each other during learning. This is important because it not only lends support to the theoretical position that there is interaction between semantics and word forms in a word, but also it suggests that this interaction is strong enough to enable the properties of one part of a word (here semantics) to influence the learning of the other part of a word (here the word forms), even when the representations of each component and the association between them is relatively weak.

The second dissertation question addresses the effect of concurrent phonological similarity (specifically cohort similarity) on the learning of the phonological word forms. Whereas the focus is on the similarity in phonology on the learning of the phonological word forms themselves, this question will not speak to the theoretical debate on

modularity vs. interactivity in a word. However, this question will inform the debate on whether newly learned words could engage in competition with similar words. Some studies have shown that newly learned words will not compete with existing words until after at least a night's sleep for the new words to be integrated into the lexicon after some consolidation (e.g., Gaskell & Dumay, 2003; Dumay & Gaskell, 2007, 2012). However, there is some recent evidence from an eye-tracking study that novel words learned after only a brief period of exposure could in fact compete with known words immediately following learning (Kapnoula, Packard, Apfelbaum, McMurray, & Gupta, 2012). An infant word learning study (Swingley & Aslin, 2007) also showed that right after learning newly learned words (e.g., *tog*) are ready to compete with the known word (e.g., *dog*), impeding the recognition of the latter. Although the debate in the literature has been on whether a recently learned novel word could inhibit existing words that are phonologically similar to it, it can be easily generalized to the interaction between novel words that are being learned together. That is, do they compete when they are being learned together and thus influence the learning of each other, or do they only compete when they are consolidated after some sleep and are somewhat integrated with each other? In the conceptual model I propose that competition between the similar-sounding novel words occurs in real time as they are being learned concurrently, and this may lead to a detrimental effect on the learning of these word forms. Thus, if a detrimental effect of cohort similarity is revealed, it will suggest that novel words that are phonologically similar to each other compete with each other as they are being learned together and there is no need of sleep-based consolidation for competition to happen.

CHAPTER 3 SEMANTIC SIMILARITY ON SPOKEN WORD PROCESSING AND LEARNING

This chapter will survey the literature with an aim to set the basis for the dissertation study on concurrent semantic similarity's effect on phonological word form learning. First this chapter will survey the studies on how the interactive processing among words similar in semantics influences spoken word processing, with an emphasis on how the existing evidence might generalize to learning the novel words, and especially learning of the phonological word forms. It will then survey the studies that provided preliminary evidence that the semantic similarity might also influence how a novel word can be acquired. This part of the survey will in particular evaluate to what extent and how well this question has been answered, with respect to the soundness of the methodologies of the existing studies. These two sets of surveys will set the basis for the dissertation question addressing the effect of concurrent referent similarity on the learning of a set of novel words.

3.1 The Effects of Semantic Similarity in Spoken Word

Processing

In this section, three lines of studies will be briefly summarized, including the studies on the effect on semantic neighbors on spoken word processing, the studies on semantic priming, and those on semantic interference and retrieval-induced forgetting. These studies revealed that semantically similar words either facilitate or impede the processing of a target word.

3.1.1 The effect of semantic neighbors

Using the visual world paradigm, Huettig & Altmann (2005) demonstrated that hearing one word (e.g., 'piano') led to more fixations to the visual object that was semantically related to the target (e.g., trumpet) than to the unrelated distractors (e.g.,

boat). In addition, the proportion of looks to the semantically related object was associated with the degree of semantic similarity, suggesting a graded co-activation of the related concepts. Some other studies found that words that are thematically related to the target word produced similar effects (Moore, Laiti, & Chelazzi, 2003; Yee & Sedivy, 2001). Mirman & Magnuson (2008, 2009) manipulated semantic relatedness and showed that close neighbors were more active than the distant neighbors. Moreover, these studies revealed that the close neighbors impaired spoken word recognition while the distant neighbors facilitated it. Chen & Mirman (2012) further showed that a simple computational model with interactive activation and competition implemented could account for these phenomena. In sum, the above findings suggest that semantically similar words tend to be co-active and can influence the processing of each other in either direction. It is an open question, however, whether novel words that are similar in semantics also interact with each other and influence the learning of each other.

3.1.2 Semantic priming

Semantic priming (e.g., Meyer & Schvaneveld, 1971; for a review, see Lucas, 2000) refers to the phenomenon that responding to a target word (e.g., *doctor*) is facilitated if it follows the presentation of a semantically related prime (e.g., *nurse*), (e.g., Meyer & Schvaneveld, 1971; for a review, see Lucas, 2000). It has been suggested that these priming effects can be accounted for by spreading activation processes and/or by increasing the activation of the shared representation (Lucas, 2000). Importantly, when similar semantic representations are primed, they feed activation to the associated phonological representations and prime them as well (e.g., Apfelbaum, Blumstein, & McMurray, 2011; Dell, 1986; McClelland & Elman, 1986; Rapp & Goldrick, 2000; Zwitserlood, 1989). These findings provide interesting suggestion to the effect of semantic similarity on the learning of the novel words. On one hand, these words may also prime each other, not only in semantics but also in phonology, which potentially can

facilitate the learning of each other. On the other hand, priming leads to ambiguous activation in the moment, which can in fact impede learning.

3.1.3 Semantic interference and retrieval-induced forgetting

In contrast to the studies on semantic priming, there are at least two fields of studies that suggest that similarity in semantics leads to detrimental effects. Studies on semantic interference revealed that naming a picture led to the impairment of later naming of other pictures from the same semantic category (e.g., Damian, Vigliocco, & Levelt, 2001; Kroll & Stewart, 1994). Specifically, the semantic interference effect turned out to be accumulative, with more preceding semantically related words producing stronger interference (e.g., Oppenheim, Dell, & Schwartz, 2010). In addition, the interference effect seemed to only accumulate with the retrieval of names from the same category, and seemed to be insensitive to intervening names from other categories or the passage of time (Howard, Nickels, Coltheart, & Cole-Virtue, 2006; Schnur, Schwartz, Brecher, & Hodgson, 2006). Therefore, it is argued that the semantic interference effect is a long-term learning effect rather than transient changes in the activation levels (Howard et al., 2006; Oppenheim et al., 2010).

Similar to research on semantic interference, research on retrieval-induced forgetting revealed that the recent retrieval acts of words impair the later retrieval of semantically related ones (for a review, see Anderson, 2003). Also similar to semantic interference, retrieval-induced forgetting tends to be a long-lasting effect that can persist for as long as a week (Storm, Bjork, & Bjork, 2012).

As for the underlying mechanisms, an inhibition account has been proposed to account for both phenomena (for non-inhibition accounts, see Anderson & Bjork, 1994; Oppenheim et al., 2010). Howard et al. (2006) argued that semantic interference would occur in any speech production system that possess three properties: sharing of semantic activation, competition and priming. Sharing activation leads to parallel activation of

multiple semantically related words during the retrieval of the target. Competition refers to the simultaneously active words competing with each other via lateral inhibition. Priming referred to the changes in the production system that can be carried over to the future retrieval acts. Howard et al. (2006) implemented the three properties in a computational model. Specifically, priming was implemented as strengthening the connection weights between the lexical and the semantic units that were active when a decision was reached at the lexical level. In this way, upon the retrieval of a later target, the previously strengthened targeted that were semantically related would compete more strongly and thus slow down the retrieval process. The inhibition account for retrieval-induced forgetting argues that during the retrieval practice phase, inhibitory control is recruited in order to overcome the competition of the unpracticed items (Anderson, 2003; Anderson & Spellman, 1995). The suppressed items therefore become less accessible during the following final test.

The studies here suggest that when concurrently *learned* novel words are similar in semantics, they may inhibit each other, making it difficult to form accurate and strong semantic representation, which could in turn cascade to phonological representation.

3.1.4 Summary

In sum, existing studies showed that when semantically related *known words* are active simultaneously with the target, they can either compete with it or prime it. Importantly, this interaction can produce a learning effect manifested as the future processing of these words being influenced. It remains to be answered whether *novel words* that are similar in semantics *to each other* also interact with each other and lead to an impact on their learning, and importantly, the learning of not only semantics, but also phonological word forms.

3.2 The Effect of Semantic Similarity on Word Learning

In this section, two areas of studies will be briefly summarized, including the studies on the effect on semantic neighbors on novel word learning and the effect of shape similarity on associative memory.

3.2.1 The effect of semantic neighbors

Overall, few studies have been done to address the effect of semantic similarity on word learning. There was only one study that examined the effect of the number of semantic neighbors on preschool children's word learning (Storkel & Adolf, 2009). In this study, the number of semantic neighbors was determined by asking participants to generate words that were meaningfully associated with the novel object. Novel objects that differed in the number of semantic neighbors were paired with nonsense words and children were required to learn the nonsense word as the label for the paired novel object. Naming and referent selection tasks showed no effect of semantic neighborhood density right after training. However, after a week the detrimental effect of large number of semantic neighbors emerged. These results suggested that semantic similarity hindered novel word learning. The authors argued that one potential mechanism was that the meanings of those semantically related known words competed with the novel object, impeding the creation and retention of detailed semantic representation of the latter.

In contrast, an opposite effect was observed by Tomasello, Mannle, & Werdenschlag (1988). Young children were taught two novel words in two successive phases. It was found that they appeared to learn the second word more efficiently, as measured in a naming task, when it was semantically similar to the first word. Tomasello et al. (1988) argued that the facilitation effect of semantic similarity derived from the fact that similarity highlighted the characteristic features of the new referent.

Thus these studies revealed mixed effects of semantic similarity on word learning, leaving it unclear whether semantic similarity should facilitate or impede learning of

word forms. Moreover, as discussed in the first chapter, the effect of semantic neighbors might simply arise at test rather than be due to a learning difference. To be specific, it is likely that during the naming and the referent selection tests, the semantically related known words were simultaneously active and competed with the target novel word, leading to a poorer response. If that was the case, the poorer performances of those words with dense neighborhood in these tests were not due to poorer learning, but instead, due to a difference in online processing at test.

3.2.2 Similarity in associative memory

Studies on associative memory using the paired-associate learning paradigms can also inform word learning since both involve arbitrary associations. The difference is that the former field typically uses known words and known images, while the later uses novel ones. Semantic similarity has been found to influence associative memory. Typically, it was found that within-pair similarity facilitated performance and across-pair similarity hindered performance (known as interference) (Nelson, Bajo, McEvoy, & Schreiber, 1989; Thompson-Schill, & Tulving, 1994; McGeoch, 1942; Osgood, 1949; cf., Pantelis, van Vugt, Sekuler, Wilson, & Kahana, 2008). The past studies generally implemented semantic similarity coarsely in terms of semantic category membership, however, a recent study used synthetic faces and achieved parametric manipulation of similarity (Pantelis et al., 2008). Since the faces used were novel ones, this study involved a learning situation closer to that of word learning. Specifically, male faces were created varying systematically along a 37-feature dimension space and were paired with common American names. Participants were required to memorize the face-name pairings and were required to recall the names given the faces as the cue at the test. The recall performance was found to be monotonically decreasing with the number of neighbor faces that the cue face had. In addition, the intrusion errors, incorrectly recalling names of other faces, was more likely to come from the faces that were similar to the cue

face. The authors interpreted that these results reflected two underlying effects: a weakening of the association between a particular face and its name, and partial associations between a name and other faces that were similar to the target face. The authors also took special efforts to rule out the possibility that the similarity effect was due to difficulties in perceptual discrimination so as to be more sure that it was indeed a learning effect. However, the same concern applied to the study on semantic neighborhood density (Storkel & Adolf, 2009) worked here as well. It is likely that at test the similar faces were simultaneously active and interfered with the target face, slowing down the retrieval of the correct name or leading to retrieval of a wrong one.

3.2.4 Summary

Some existing studies provided preliminary evidence (although mixed) that similarity in referents might influence novel word learning. However, the behavioral tests used in these studies involved a real time processing confound at test. Thus, the detrimental effect of semantic similarity in naming could simply be due to more competition and interference in the high similarity condition than in the low similarity condition in the moment of test, and were not able to establish whether it was indeed a learning effect or not.

3.3 Overall Summary

In sum, existing studies on spoken word processing showed that semantically related known words interacted with each other, which not only influenced the performance in the moment, but also lead to a durable learning effect. These studies suggest that when words that have concurrent semantic similarity are being learned together, the interaction among them may also influence how well they could be learned.

To date, no studies have addressed how referent similarity might influence the concurrent learning of a set of novel words. Some studies addressing similar situations

provided preliminary evidence that similarity in referents might influence novel word learning, but due to a methodological defect the evidence revealed in these studies may not indeed represent a learning effect.

This dissertation study will address for the first time how the concurrent learning of a set of novel words might be influenced by referent similarity. More importantly, behavioral tests that carefully control for real time processing across conditions, and a computational modeling study will be carried out so as to provide an even more convincing index of a learning effect that is not contaminated by real-time processing. Chapter 4 will summarize the behavioral research and Chapter 5 will summarize the modeling study.

CHAPTER 4 THE BEHAVIORAL STUDY ON THE EFFECT OF REFERENT SIMILARITY ON CONCURRENT NOVEL WORD LEARNING

As discussed in Chapter 3, semantic similarity leads to interactions between words. Studies on semantic priming suggest when the novel words are being learned currently, the semantically similar ones will be activated. This may lead to two opposite effects on learning: it may facilitate the learning of the words being primed, but may hurt the learning of the prime itself. Critically via cascaded activation (e.g., Dell, 1986; McClelland & Elman, 1986; Peterson & Savoy, 1998; Rapp & Goldrick, 2000), the associated phonological representation will also be influenced, and thus the learning of the phonological word forms can be affected as well.

In contrast, studies on semantic interference and retrieval-induced forgetting suggest that when the novel words with similar referents are being learned, these partially active words may compete with each other, making it difficult to form accurate and strong representations. Similarly as discussed above, not only the learning of the semantics but also the learning of the phonological word forms will be impaired.

In addition, similarity in referents may influence novel word learning mediated by visual categorization process. Word learning involves mapping labels to referents, and thus one requires categorizing the referents as separate entities. To the extent that concurrent semantic similarity creates between-category similarity, it would be expected to lead to increased difficulty in categorization (Goldstone, 1994; Kloos & Sloutsky, 2008; Nosofsky, 1986, 1992; Posner & Keele, 1968; Reed, 1972; Rosch & Mervis, 1975; Palmeri, 1997), which in turn would lead to increased difficulty in word (and presumably word form) learning. However, as Tomasello et al. (1988) suggested, it is also possible that in some cases similarity of the referents makes categorization between them easier, and thus may facilitate word learning.

In sum, the research summarized above suggests that semantic similarity among currently learned words will have an impact on learning of these words, including not only the learning of the semantics, but also the phonological word forms. However, the influence can be either facilitative or detrimental.

The current study took a lead in addressing this question. Experiment 1 aimed to provide more definitive information about the direction of this effect, using the naming test that was also used in previous studies (e.g., Storkel & Adolf, 2009; Tomasello et al., 1998). However, the naming test is not an adequate measure of the learning of phonological word forms. For one thing, completing this test requires recognizing the presented referent, having the right association between the referent and the phonological word form, and having the correct representation of the word form. Thus, if an effect of current semantic similarity is revealed by a naming test, it is not clear whether the effect lies in the phonological word form representation or not. For another, a behavior/performance in a given task is usually “multiply determined” (Gathercole, 2006; Gupta, 2008; Gupta & Tisdale, 2009), including the representation acquired through past learning and the in-the-moment processing (Gupta & Tisdale, 2009; McMurray, Horst, & Samuelson, 2012; McMurray, Horst, Toscano & Samuelson, 2009). An effect of concurrent semantic similarity in naming can potentially come from either source. That is, when a referent similar to other referents is being named, the other referents are activated as well, which in turn activate their phonological word forms. Competition among these co-active word forms will impair the retrieval of the desired label. This is less likely to happen when the words do not have similar referents. Thus, a poorer naming performance in the current semantic similarity condition may simply be driven by these in-the-moment processes at test, and have nothing to do with what the learners know about the words. With these considerations in mind, Experiment 2 and Experiment

3 used tests that aimed both to be a pure measure of the phonological word forms and to control for online processing between different similar levels.

In sum, three experiments were conducted. In all of them, we employed a learning paradigm in which participants simultaneously learned a set of nonsense words, and mapped them to a set of novel visual referents (Gupta 2003; Abbs & Gupta, 2008). We manipulated semantic similarity while holding phonological similarity constant, thus examining whether a semantic manipulation would influence purely phonological performance (e.g., knowledge of the word form). Experiment 1 used the naming test. Although this test cannot provide unambiguous evidence that phonological word learning is influenced, it helps reveal the direction of the effect because it is the most commonly used test in existing studies. Experiment 2 and Experiment 3 used specially designed tests, the stem completion test and the recognition from mispronunciation test respectively, to provide more convincing evidence that the current referent similarity influences the phonological word form learning.

Experiment 1

The aim of this experiment was to examine whether referent similarity has any influence on recall of concurrently learned novel words in a naming task. As discussed in the introduction, the naming task is not adequate to isolate the effect of concurrent referent similarity on phonological word form learning. It is nevertheless a canonical test of word form knowledge, and we therefore employed it in Experiment 1 to obtain an indication of at least the direction of any effect of concurrent similarity on overall observed performance.

In this experiment, participants were required to learn a set of nonsense words as the labels of visually presented referents (nonsense shapes) in one of two conditions that differed in terms of referent similarity. The first condition used sets of shapes that were highly similar to each other (High-Similarity), and the second condition used sets of

shapes that were dissimilar to each other (Low-Similarity). A naming test was administered periodically at the end of each of the seven learning blocks to measure how participants' performance grew in each of the conditions.

Methods

Participants

Thirty undergraduate students at the University of Iowa were recruited from the department subject pool and were granted course credit for participation. All participants were native speakers of English who reported having normal hearing and normal or corrected-to-normal vision. They were randomly assigned to one of two conditions: High-Similarity and Low-Similarity.

Materials

Pictures of nonsense shapes from a corpus developed by Vanderplas and Garvin (1959) were used as referents of novel phonological word forms. Specifically, five shapes of low-associative values from this corpus were chosen to minimize the influence from semantic knowledge. As these five shapes were very different from each other, they were used as referents for the Low-Similarity condition. Using these shapes as the templates, five sets of referents of high similarity were created. A set of High-Similarity shapes consisted of one shape of the Low-Similarity condition and four modified exemplars of it that differed slightly from each other and from the template shape (see Figure 4-1 for examples).

The pair-wise pixel similarity of five shapes within each set was indexed by Euclidean distance. The averages of this Euclidean distance in the Low- and High-Similarity conditions were 0.42 and 0.28 respectively ($t(54) = 8.41, p < .001$). In addition, to match these conditions on the associative values, a verification task was run on 18 participants. In this task, one shape was presented at a time on the computer screen for

three seconds, and participants were required to indicate whether the shape reminded them of some real object or real situation. The average rates of “yes” responses were 0.60 and 0.58 respectively for the Low- and High-Similarity conditions, which were not statistically different ($p > .30$).

Two-syllable CV-CVC nonwords were used as the labels of the shapes. A corpus of 47 nonwords was created with a constraint that neither of the two component syllables, was a real word. Items were stressed either on the first syllable or the second syllable. These nonwords were digitally recorded as digitized sounds by a female native speaker of English at 44.1 k Hz sampling rate and amplitude normalized.

Three sets of 47 nonword stimuli were then created from the corpus. In each set, five items were chosen randomly from the corpus to be the learning targets (with the constraint that they should not share the initial consonant and they should be very distinct from each other). Each target word was paired with a nonsense shape as its referent. The remaining 42 nonwords were distributed across the seven learning blocks as non-target fillers, six in each block. These were intended to minimize the possibility that participants could rehearse the small number of target nonwords across trials. Finally, to complete the stimulus set we randomly paired a set of five referents with the five target nonwords in one set of nonword stimuli. For the High-Similarity condition, 15 sets were created in total by crossing the three sets of nonword stimuli and the five sets of referents. For the Low-Similarity condition, the base shapes were assigned to the target nonwords in each set of auditory stimuli in a pseudorandom manner to make sure that each nonword was associated with all five shapes, thus also creating 15 stimulus sets in total.




Figure 4-1. Examples of random shapes used in Experiment 1. The first row shows the low similarity condition, and the second row shows high similarity condition.

Procedure

The experiment consisted of seven *epochs*, with each consisting of a learning block and a naming test immediately following it.

There were 11 trials in a learning block, consisting of five target trials interspersed with six foil trials. The target stimuli re-occurred in each of the seven learning blocks, but the foil stimuli in each block were unique. On both kinds of trials, a nonword was presented auditorily through headphones, and participants were required to repeat it aloud into a microphone as quickly and as accurately as they could. However, only on the target trials, a nonword was accompanied by a picture of a nonsense shape on the computer display and participants were asked to memorize the nonword as the label of that shape.

A naming test consisted of five trials. In each trial, an image of a nonsense shape was presented and participants were required to recall its name aloud. The order of the five targets was randomized in each of the seven learning and test blocks.

Reliability

A trained coder who was blind to the study's purpose coded the naming accuracy at the whole-word level. That is, a response was coded as being correct if all the five phonemes were produced correctly. Another independent trained coder coded 67% of the data, and the point-by-point agreement between them was 96.7%.

Results and Discussion

Naming accuracy was first subjected to arcsine transformation and then subjected to a 2(similarity: low, and high) by 7(block: 1-7) mixed ANOVA (see Figure 4-2)¹. A significant main effect of similarity was revealed, $F_1(1, 28) = 7.05, \eta_p^2 = .20, p < .05, F_2(1, 24) = 12.74, \eta_p^2 = .35, p < .01$, showing that naming accuracy decreased as the referent similarity level increased. A significant main effect of epoch and a significant interaction between similarity and epoch were also revealed, with $F_1(6, 168) = 32.97, p < .001, \eta_p^2 = .54, F_2(6, 144) = 33.72, p < .001, \eta_p^2 = .59$, and $F_1(6, 168) = 2.68, p < .05, \eta_p^2 = .09, F_2(6, 144) = 2.43, p < .05, \eta_p^2 = .09$, respectively. Simple effect analysis showed that for both of the Low-Similarity and the High-Similarity conditions, the naming accuracy increased significantly across the seven epochs, $F_1(6, 84) = 21.95, p < .001, \eta_p^2 = .61, F_2(6, 84) = 24.97, p < .001, \eta_p^2 = .64$, and $F_2(6, 84) = 13.23, p < .001, \eta_p^2 = .49, F_2(6, 84) = 13.06, p < .001, \eta_p^2 = .49$. However, the Low-Similarity condition outperformed the High-Similarity condition for all epochs except for epoch 1 and epoch 5 ($ps < .05$)

¹ In an initial analysis, the set of nonwords was also included as a between-subjects factor. Neither the main effect of nonword set nor the interaction between nonword set and similarity was significant. Thus, the analysis reported here collapses across nonword sets.

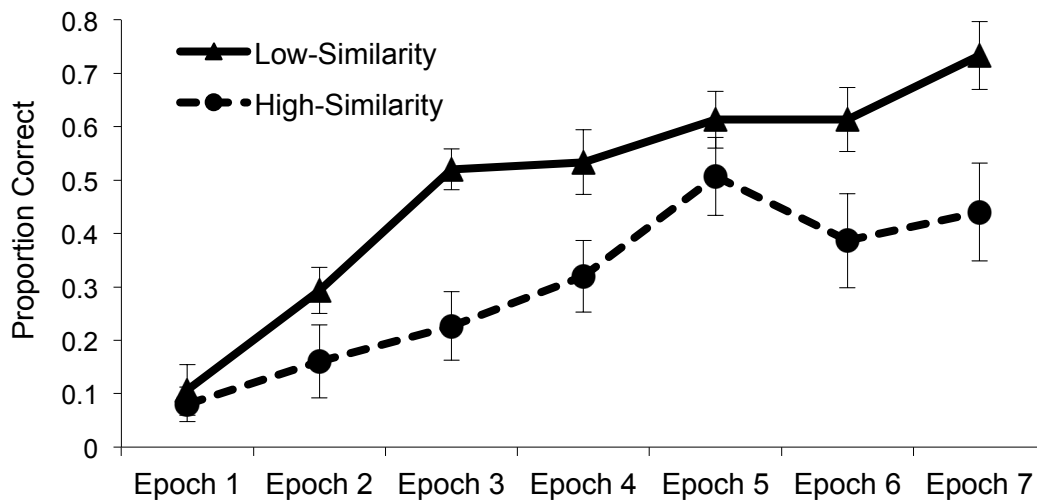


Figure 4-2. Naming accuracy for the three similarity conditions across seven tests.

Clearly, the results suggest that participants' performance is worse for the concurrently learned novel words with similar referents, when the word form needs to be retrieved given the referent. To my knowledge, this provides the first direct evidence regarding the effect of specifically concurrent referent similarity on word learning. This is important because naming is a very widely used measure of lexical knowledge, in addition to being one of the most common "word-use" situations. However, as discussed

in the Introduction, the locus of this effect is not necessarily in the learning of *phonological word forms* per se. It could instead be due to increased competition and interference at test rather than due to having poorer word form representations. In addition, even if the results of the naming test truly reflect a learning effect, it is unclear which aspect of word learning was affected – whether learning of the referent, the word form, or the association between the two. In the following two experiments, we attempted to overcome these shortcomings of the naming test by adopting two different tests, the stem completion test and the recognition from mispronunciation test, so as to examine whether phonological word form learning is indeed affected by concurrent referent similarity.

Experiment 2

In this experiment, the learning phase was followed not by a naming test but by a stem completion task to measure phonological word form learning. In this test, participants were given the first syllable of a phonological label as the cue and were required to recall the whole word form that they had learned. Because it probed the phonological word form based on a fragment of it (and with no visual referent present), this was a relatively pure phonological test, not requiring the participant to access either the referent or the referent-word form link. In addition, since the referent was not presented as in a naming test, the likelihood of competition and interference induced by the co-activation of multiple referents at test should be greatly reduced even when the referents were similar. With this effect of online processing at test largely controlled for across the conditions, any effect of referent similarity on observed performance should more truly reflect an effect on the learning of the phonological word forms.

Methods

Participants

Eighty undergraduate students at the University of Iowa were recruited from the department subject pool. They were granted course credit for participation. All participants were native speakers of English who reported having normal hearing and normal or corrected-to-normal vision. They were randomly assigned to one of two conditions: Low-Similarity and High-Similarity.

Materials

The pictures of nonsense shapes were the same as those in Experiment 1. Two-syllable CV-CVC nonwords were used as the labels of the shapes. A different corpus of 47 nonwords were made with the same constraint that neither of the two syllables in an item was a real word. These nonwords were stressed either on the first syllable or the second syllable. The nonwords and their first syllables were recorded as digitized sounds by a female native speaker of English and amplitude normalized.

Two sets of nonword stimuli were created in the same way as those in Experiment 1. They were combined with sets of nonsense shapes in the same way as well, creating 10 different complete stimuli sets in total for both the Low- and the High-Similarity conditions.

Reliability

A trained coder who was blind to the study's purpose coded stem completion accuracy at the whole-word level. That is, a response was coded as being correct if all the five phonemes were produced correctly (including the first syllable that was given). Another independent trained coder coded 64% of the data, and the point-by-point agreement between them was 96.2%.

Tasks and Procedures

The task consisted of seven learning blocks and one stem completion test at the end of the seventh block. The learning blocks were exactly the same as those in Experiment 1. However, naming tests were not administered at the end of each block, to avoid inducing referent competition among the highly similar referents that might then affect learning. In the stem completion test at the end of training, participants heard the first syllable of a target nonword and were required to produce the whole nonword that began with it. Each stem completion trial started with the probe sentence “Please say the nonword that begins with...” that remained on the computer display for 1500ms. Then the first syllable of one nonword was presented by headphone. There was no time limit for participants to recall. If participants indicated that they did not know, they were encouraged to guess. The next trial started after a response was made (spontaneous response or prompted guess). Although participants were informed at the beginning of the experiment that they would be tested at the end of the experiment, specific instructions for the stem completion test were not given until the end of the seventh learning block.

Results and Discussion

Figure 4-3 shows the accuracy during the stem completion trials (after training) as a function of similarity. As before, semantic similarity decreased performance. This was confirmed with a one-way ANOVA^{2,3} which found that the main effect of referent similarity was significant, $F_1(1,78) = 4.20, p < .05, \eta_p^2 = .05$, $F_2(1,8) = 12.13, p < .01, \eta_p^2 = .60$.

² In an initial analysis, the set of nonwords was also included as a between-subjects factor. Neither the main effect of nonword set nor the interaction between nonword set and similarity was significant. Thus, the analysis reported here collapses across nonword sets.

³ Arcsine transformation was applied to the accuracy data here before they were subjected to ANOVA.




Figure 4-3. Stem completion accuracy for the Low- and the High-Similarity conditions.

This result showed that concurrent referent similarity impaired stem completion performance for newly learned novel words. Because in the stem completion test, no referent was presented, it presumably eliminated at-test competition and interference differences at the semantic level between the similarity levels. Thus, this finding provides the first evidence that concurrent referent similarity impairs the learning of novel words. In addition, because performing the stem completion test requires accessing the phonological word form representation, and should not require accessing semantics or the form-referent link, this finding suggests strongly that it is specifically building representations of the phonological word forms that are impaired by concurrent referent

similarity. Therefore, this experiment provides the first empirical evidence that semantic variables can impact the learning of the phonological word forms.

Experiment 3

This experiment used a test of recognition from mispronunciation, to provide converging evidence that concurrent referent similarity had a detrimental effect on the learning of phonological word forms. This test required participants to recognize a target word form from two mispronounced versions of it that differed by only one phoneme. Like stem completion, this was a relatively pure phonological test because performing it should require accessing only the phonological labels. In addition, because the referent was not presented at test as it was in a naming test, the likelihood for confounding by at-test effects arising from differential competition in the different similarity conditions was again very low. Therefore, an effect of referent similarity revealed by this test should again more truly reflect a learning difference among the conditions. Finally, there is a question as to whether semantic similarity impairs learning *in general* (e.g., subjects get confused by the similar referents and “give up”) or whether its effects are isolated to the specific words that are similar. To address this we adopted a within-subject design in Experiment 3.

Methods

Participants

Sixteen undergraduate students at the University of Iowa were recruited from the department subject pool. They were granted course credit for participation. All participants were native speakers of English who reported having normal hearing and normal or corrected-to-normal vision. Each subject participated in both the Low and High Similarity conditions.

Design, Task and Procedures

Experiment 3 used a within-subjects design both to increase power, and to determine if the effects of semantic similarity were isolated to the specific words, as noted above. The structure of the task was also modified to avoid participant fatigue. The task consisted of three learning blocks and a recognition-from-mispronunciation test at the end of training. The learning blocks were exactly the same as those in Experiment 1 except that there were ten targets (five for each condition) and 11 fillers in each block.

In the recognition test, each trial started with a fixation cross at the center of the screen for 500ms followed by a blank screen of 500ms. Then, three nonwords were presented in a sequence from headphones. One of them was the correct form of a target word form and the other two were its mispronounced versions. The presentation order of the three words was randomized. At the end of the third word, participants were prompted by the question on the screen “Which One?” to indicate which one was the correct one by pressing a number key (“1”, “2”, or “3” respectively). Although participants were informed at the beginning of the experiment that they would be tested at the end, they were not given specific instructions for the recognition test until the end of the third learning block.

Materials

The pictures of nonsense shapes used were the same as those in Experiment 1 except that only two sets of each of the High- and Low-Similarity shapes were used (five sets were used in Experiments 1 and 2).

For the nonword stimuli, we first created two pools of two-syllable CV-CVC nonwords, one consisting of 15 nonwords and the other consisting of 120 nonwords, all incorporating the constraint that neither of the two syllables in an item was a real word⁴.

⁴ The stimuli in the first of these pools were also in the corpus of 47 stimuli used in Experiment 1, and the stimuli in the larger pool had some overlap with stimuli used in Experiment 2

We also created two mispronounced versions of each of the fifteen stimuli in the first pool, each differing from the original stimulus by one phoneme, for use as lures in the recognition-from-mispronunciation test. Two sets of nonword stimuli for the learning phase were then constructed by choosing ten targets at random from the first pool, and 33 filler nonwords (11 for each block) at random from the second pool, for a total of 43 nonwords in each set. Stimuli for the accompanying recognition-from-mispronunciation test consisted of the ten targets chosen for that set together with their mispronounced lures.

Within each set of 43 nonwords, the ten targets were randomly divided into two sub-sets of five, with one sub-set assigned to be labels of one set of the High-Similarity pictures and the other set assigned to the Low-Similarity pictures. The pairings between sub-sets of nonwords and sets of the High- and Low-Similarity pictures were counterbalanced across participants.

All the nonwords were digitally recorded at 44.1 kHz, by a female native speaker of English and amplitude normalized.

Results and Discussion

The recognition accuracy (see Figure 4-4) was first subjected to arcsine transformation and then to a repeated measures ANOVA with accuracy as the fixed effect⁵. Recognition accuracy in the Low-Similarity condition was found to be significantly higher than that in the High-Similarity condition, $F_1(1,15) = 12.13$, $p < .01$, $\eta^2 = .60$, $F_2(1,12) = 8.39$, $p < .05$, $\eta^2 = .41$. It suggests that listeners were much more accurate for the low-similarity words than for the high-similarity words.

⁵ In an initial analysis, the set of nonwords was also included as a between-subjects factor. Neither the main effect of nonword set nor the interaction between nonword set and similarity was significant. Thus, the analysis reported here collapsed across nonword sets.




Figure 4-4. Recognition accuracy for the Low- and the High-Similarity conditions.

This result indicates that referent similarity among a set of concurrently learned words impairs performance in recognizing a novel word form from highly confusable phonological distractors. Because this recognition test largely eliminated the at-test processing difference among the different similarity conditions and because it was a relatively pure phonological test, it provides converging evidence that concurrent referent similarity impairs the learning of the phonological word forms. Moreover, because these effects were observed within participants, this suggests that the locus of the effect is restricted to the specific words being learned.

General Discussion

This study provided the first evidence that the referent similarity among concurrently learned novel words impaired not only how the newly learned words could

behave in the naming situation but also the learning of the phonological word forms of these words. Since all previous studies on semantic similarity addressed corpus similarity (similarity to existing words in the lexicon; e.g., Storkel & Adolf, 2009), by shifting gear to concurrent semantic similarity (similarity among the words being learned concurrently), this study enriches our understanding of how semantic similarity might influence novel word learning.

This study also for the first time provided direct evidence on the influence of a semantic property on the learning of phonological word forms. This extends our understanding that there is interactive activation between semantics and phonology from the domain of processing to learning. At the same time, it also suggests that the interaction between semantics and phonology not only leads to an immediate effect in real time processing but also lead to a learning effect that is on a more long-term time scale.

Our ability to draw these conclusions rests upon our use of the stem completion test and the recognition from mispronunciation test. Compared to traditional tests such as naming, these two tests do not present the referent and thus largely eliminate the involvement of semantic processing, and also controlled for real time processing difference between the conditions. Therefore, this study taps the effect of interest (i.e., learning) more closely than almost all existing studies on related topics. In this sense, it makes an important methodological contribution to the field of word learning.

However, the stem completion and recognition from mispronunciation tests may not completely eliminate the interpretational difficulties that a naming test has. Although the target referent is not presented during these two tests, nevertheless, the phonological label (in the recognition test) or its initial syllable (in the stem completion test) *itself* may partially activate its referent, which in turn may lead to the co-activation of the similar referents and their labels. If that is the case, the retrieval of the target label in the

conditions with similar referents would have still suffered more competition and interference from the other labels, and thus, we still could not be certain whether the effect of similarity on these tests indeed represent a learning difference. In addition, if the stem completion and the recognition tests do lead to activation of referents, it is also possible that the effects observed in these tests really reflect learning differences at the *semantic* level.

We therefore followed up the above behavioral investigation using the computational modeling methodology with a hope that a computational model can factor out real-time processing and learning and provide a pure measure of phonological word form learning.

CHAPTER 5 THE SIMULATION STUDY ON THE EFFECT OF CONCURRENT REFERENT SIMILARITY ON WORD LEARNING

We developed a computational model of this problem within the Hebbian Normalized Recurrence Network (HNRN) framework. This network combining the approaches of connectionist and dynamic systems has been used to successfully simulate various phenomena of word learning under referential ambiguity (McMurray, Horst, Toscano, & Samuelson, 2009; McMurray Horst, & Samuelson, 2012). The great merit of this model is that it incorporates both on-line competition and learning, and has a way to factor them out computationally. Specifically, if a performance difference in the model is accompanied by a difference in connection weight strengths, then at least part of the observed performance difference must be due to these differential weight strengths rather than due to differences in competition. Thus differing weight strengths provide an unambiguous index of differential learning. This was crucial for our purpose because the auditory-to-lexical weights could be analyzed as a way to compare phonological word form learning across similarity conditions independent of on-line competition processes.

Typical HNRN models of word learning have three layers: a referent layer, an auditory word form layer, and a lexical layer (see Figure 5-1). A node in the referent layer refers to an object and one in the auditory layer refers to a phonological label. Nodes in the lexical layer are abstract internal representations of words that link referents and labels. Over training, the model learns the correct set of weights between the auditory and referent inputs and the lexical units so that a single lexical unit connects a referent and its corresponding label. In the present work, we augmented this architecture with a fourth layer as shown in Figure 5-2. The addition was a visual feature layer that fed input to the localist referent layer. This enabled referents to be represented as distributed activation patterns across feature nodes, so that the similarity among them could be manipulated by varying the amount of pattern overlap. We used this model to simulate

the learning of a set of new words with either high or low concurrent semantic similarity. As a means of validating the model and linking it to observed behavioral performance, we tested its ability to simulate the naming and the stem completions results of Experiment 1 and 2. The auditory-to-lexical weights underlying phonological word forms in the model were then analyzed. If these weights differed as a function of the manipulation, it would establish that concurrent semantic similarity can have an effect that is unambiguously located in phonological word form learning.

Architecture of the model

The current model had four layers: a visual feature layer (20 units), a visual category layer (400 units), an auditory word form layer ($20 \times 3 = 60$ units, as explained below), and a lexical layer (500 units)⁶. Each of the 20 nodes in the visual feature layer represented an artificial visual feature and an activation pattern across these nodes represented a referent. Each of 20 referents had a single unique defining feature and 15 non-defining features that were randomly chosen from the remaining 19 features. Thus one referent's non-defining features were the defining features of the other referents, and different referents thus shared a certain number of non-defining features. In the feature vector for a given referent, the defining feature always had an activation of 1.0. The activation level of the non-defining features provided a way of manipulating referent similarity: an activation level of 0.3 for non-defining features constituted a relatively low level of similarity between referents, while an activation level of 0.8 for non-defining features constituted a relatively high level of referent similarity.

⁶ Having this large number of decision units is critical to the success of the model because it minimizes the chance that any given unit is chosen as the lexical unit for more than two words.

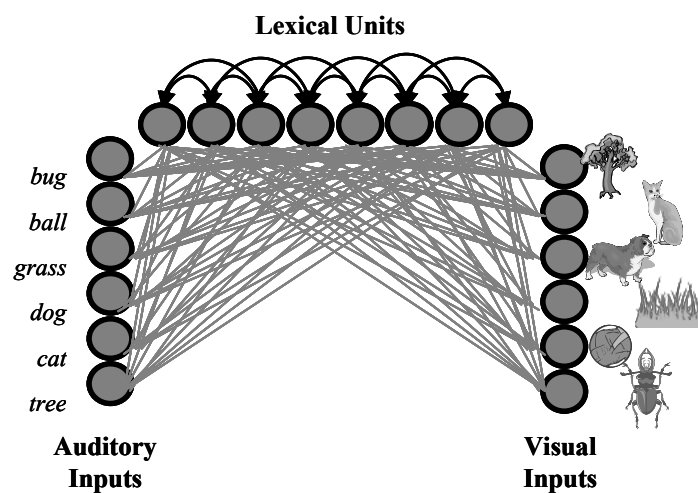


Figure 5-1. The architecture of a typical HNRN model. Adopted from McMurray, Zhao, Kucker, & Samuelson (2012).

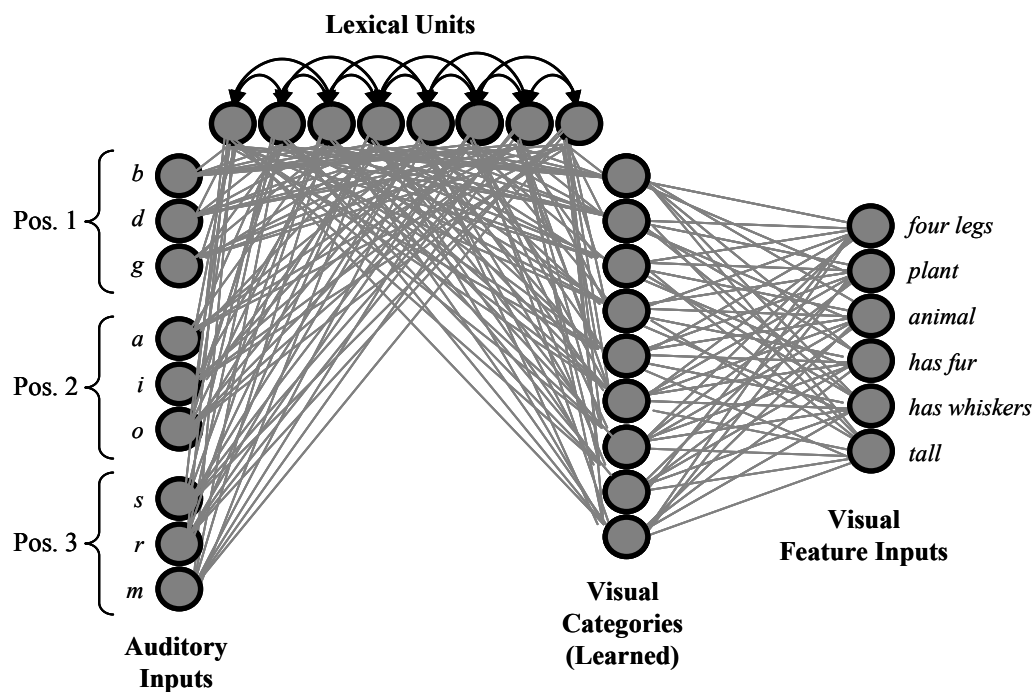


Figure 5-2. Structure of the 4-layer network. Inputs consist of visual features and word-forms. Visual categories are learned. Adopted from McMurray, Zhao, Kucker, & Samuelson (submitted).

Each node in the visual category layer represented a visual category. In the auditory layer, three banks of nodes were used, with each bank representing a phoneme position and each unit in a bank representing an artificial phoneme. A word form was

represented by a distributed representation across the three banks, with only one unit active in each bank. The word form representations were made completely orthogonal to each other. This provided a reasonable approximation to the nonwords used in the behavioral experiments, which were designed to be very dissimilar to each other. More importantly, it allowed an examination of the effect of referent similarity independent of phonological similarity. Each lexical node was a localist unit, representing an abstract word that binds a visual category and a word form.

Processing in the model

Each cycle of processing in the model proceeded as follows: input activation fed forward from the auditory layer and the visual feature layer (via the visual category layer) to the lexical layer, and then fed back from the lexical layer to the auditory layer and the visual feature layer (via the visual category layer) (for the values of the parameters, see Table B-1 in Appendix B; for formulas, see Appendix C). On each cycle, the activations of nodes in each layer were normalized so that they summed to 1.0. Inhibition was applied at all layers so that the most active node suppressed the less active ones to some extent at each step of processing. Preliminary examinations of the parameter space suggested that using different levels of inhibition at the different levels yielded the best performance. Thus, the category layer used $\omega=2.0$, the lexical layer used $\Psi=1.5$, and a very low level of inhibition was applied at the visual feature and auditory layers⁷.

Activation cycled back and forth in the model in this manner until the lexical layer settled (the derivative from cycle to cycle is smaller to a very low threshold). The number of cycles ranged from hundreds or thousands at the beginning of the training to less than 10 at the end.

⁷ A systematic parameter search suggests that having a high level of competition at the category level and a moderate level of competition at the lexical level produces the best results.

Learning in the model

Hebbian learning occurred at the end of each cycle. The association strength of two nodes from adjacent layers increased if both nodes were active, decayed if only one node was active, and remained unchanged if both were inactive. At each step of processing, inhibition shaped the activation at each layer by making the most active node suppress the competitors. It therefore also shaped learning incrementally by causing the associations between the nodes that won the competition on each cycle to become stronger and the rest weaker.

Initially all the feature nodes were fully interconnected to all the visual category nodes, and all auditory nodes and all visual category nodes were fully interconnected to all lexical units. Eventually with sufficient training, the model learned to categorize the feature patterns by associating the most defining feature in a pattern with only one category node, so that one visual category node came to represent one set of visual features. In addition, with training, one lexical node came to link the three auditory nodes representing a word form with the category node that bound the corresponding visual feature pattern.

Testing the model

To test the model, we used an analogue of the naming test and an analogue of the stem completion test. In the naming test, the referent was given (as an input pattern in the visual feature layer) to probe for the auditory label. In the stem completion test, two nodes of a label (in two of the auditory banks) were activated to probe for the third. If learning in the model is valid, it should be able to perform these two tests and should produce referent similarity effects similar to those found in the behavioral experiments. To determine whether these effects have a true locus in the learning of the phonological labels, the connection weights between the lexical and the auditory nodes were analyzed. The strength of these connection weights is a pure index of the phonological word form

learning that accumulates from training. If these weights in the high similarity condition were lower than those in the low similarity condition, it would indicate convincingly that referent similarity truly impairs this learning.

Results and Discussion

Both the naming and the stem completion performance were poorer for the high referent similarity level (Figure 5-3 and 4-4). Thus, the model successfully captured the behavioral results. However, this does not necessarily mean that the locus of these referent similarity effects lies in the learning of the auditory labels. The effect in the naming test could simply be driven by processes occurring at test. When objects are similar, multiple visual category nodes will be active simultaneously and compete with each other. Further, these category nodes will spread activation to the associated lexical nodes, inducing a higher level of competition at the lexical level. The same process will then happen at the auditory layer when the lexical activation spreads to it, eventually leading to slower and less accurate naming. For the stem completion test, although there is no visual input provided externally, the auditory input (the stem) will be spread to the visual category and the visual feature layers. Then the internally activated visual representation will spread back to the auditory layer. Thus, as in the naming test, poorer stem completion for the high similarity condition could just be due to the increased level of representation ambiguity and competition at test. We addressed this concern by lesioning the network at test. Specifically, the connections from the visual feature layer to the visual category layer and those from the visual category layer to the lexical layer were lesioned, leaving only the connections from the auditory layer to the lexical layer intact. The results on stem completion showed that the same pattern as those when the visual part of the network was intact. Therefore, the differences in the auditory weights between the high-similarity and the low similarity conditions are sufficient to drive the effect in the stem completion test, and the contribution from the referents is not necessary.

However, we also found that higher referent similarity led to poorer *learning* of the auditory weights that link auditory representations to lexical representations. As can be seen in Figure 5-5, these weights in the High-Similarity condition started to grow at a later point than those in the Low-Similarity condition, and thus at a given point before they both reached the asymptote level, the former weights were weaker than those in the latter. Because in this model connection weights represent learning that is independent of real time processing, these results indicate unambiguously that referent similarity indeed impairs the learning of the phonological labels. This in turn indicates that the effects of referent similarity on the phonological performance observed in simulated naming and stem completion have a locus in differential phonological learning. Given that the model's simulated naming and stem completion captures the broad patterns observed in the behavioral results, this suggests that the observed behavioral effects of referent similarity are also likely to have at least a partial locus that is truly in phonological learning. Thus, taken together, the behavioral and the computational evidence provide convincing evidence that referent similarity, a semantic variable, influences the learning of another element of a word, the phonological word form.




Figure 5-3. Proportion correct on naming task for the Low- and High-Similarity conditions simulated by the 4-layer network.




Figure 5-4. Proportion correct on stem completion task for the Low- and High-Similarity conditions simulated by the network.

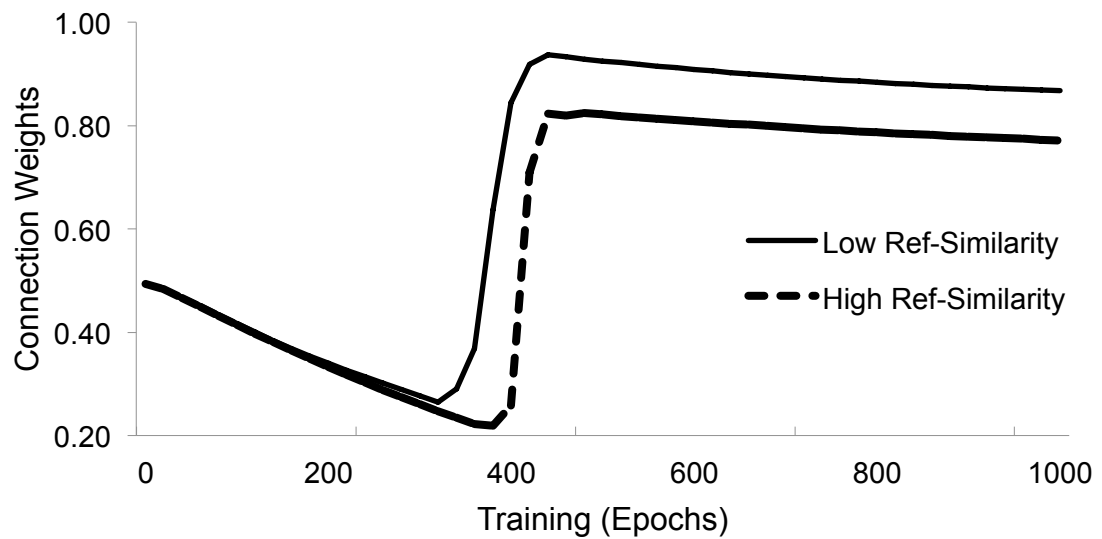


Figure 5-5. The strength of the connections between the auditory and lexical nodes for the Low- and High-Similarity conditions in the model.

Why does referent similarity impair phonological learning? The reason that the auditory weights are weaker in the high similarity condition is likely to be the spreading activation dynamics that occur during learning. In the model, activation can spread between the adjacent layers in two directions. Thus, the ambiguous representation at the feature layer can spread to the visual category layer, the lexical layer, and even to the auditory layer, leading to ambiguous activations in these layers. The simultaneously active result in many spurious associations being formed, reducing the quality of the learned representations.

The reason that the auditory weights are delayed in its development by referent similarity is possible due to a developmental sequencing mechanism. We found that visual similarity delayed the development of the feature-to-visual category weights ($F \rightarrow Vc$ weights)⁸ (see Figure 5-6) and the visual category-to-lexical weights ($Vc \rightarrow L$ weights)⁹ (see Figure 5-7), in addition to the auditory-to-lexical ($A \rightarrow L$ weights). More importantly, there was a precise time locking between these weights: the $F \rightarrow Vc$ weights self-organized consistently before the $Vc \rightarrow L$ weights and the $A \rightarrow L$ weights, and the latter two self-organized almost simultaneously once the former was organized. Thus, the categorization of referents seems to be a pre-cursor in word learning in that once the referents are categorized, the referents and labels can then rapidly link together.

⁸ Similar to the auditory-to-lexical weights, the feature-to-visual category weights refer to the connection strengths between the defining feature of each of the referent representation and the visual category node that was selected to associated with it.

⁹ The the visual category-to-lexical weights refer to the the visual category node that was selected to associated with the most-defining feature and the lexical node that was selected to link both that visual category node and the three auditory nodes of the associated label.




Figure 5-6. The strength of the connections between the visual feature and the visual category nodes for the Low- and High-Similarity conditions in the model.




Figure 5-7. The strength of the connections between the visual category and the lexical nodes for the Low- and High-Similarity conditions in the model.

General Discussion on Chapter 4 and Chapter 5

In this study we addressed the effect of referent similarity in a concurrent word-learning situation. We revealed a detrimental effect of this similarity in the phonological word form learning consistently across three behavioral experiments and one simulation. These results extend our understanding of lexical-semantic similarity from its well-

known influences on lexical processing to an effect on learning. It is important to know that semantic similarity not only influences real time lexical processing, but also leaves an impact on the internal lexical representations that may impact processing and learning in the future.

Our finding is important also because it suggests that variation in *semantics* influences the learning of *phonology*. While there have been previous studies that suggested an effect of semantic variability on learning, they either used tests such as naming from which no clear conclusions could be drawn (Capone & McGregor, 2005), or used training tasks that differed in other aspects in addition to a difference in semantic properties (Leach & Samuel, 2007). Aided by carefully designed behavioral tests and computational modeling, the present study provided the first convincing evidence of the effect of a semantic variable on phonological learning. This finding lends further support to the findings and proposals that there is interactive activation between semantics and phonology (e.g., Apfelbaum, Blumstein, & McMurray, 2011; Dell, 1986; McClelland & Elman, 1986; Rapp & Goldrick, 2000), and extends them in an important way. That is, semantics and phonology interact not only in the processing but also the *learning* of words, and that this interaction is strong enough to make one part of a word influence the learning of the other part. Our HNRN model suggests that the interaction occurs both in developmental time (in the way of developmental sequencing) and in the real time (activation spreading shapes Hebbian learning to occur).

An additional contribution of the present work is our methodological approach to isolating the learning effect from contamination by online processing. This approach is, we believe, relevant to future studies of word learning and probably studies of learning in other domains as well. As we have highlighted throughout this thesis, a given task behavior can be driven by multiple sources, including a learning source and online processing sources. Thus, in order to conclude that there is a learning effect underlying an

observed performance difference, the online processing difference between conditions must be controlled for.

This is consistent with other recent work emphasizing this distinction. For instance, a recent study of fast mapping showed that although infants could solve the referential uncertainty after minimal exposures, they showed no retention of the correct mapping after only five minutes (Horst & Samuelson, 2008); the authors suggested that fast mapping may not represent learning but may instead represent in-the-moment inference process. This is supported by computational studies that factored out online processing and learning and showed a superior performance of fast mapping but only a tiny amount of learning (in terms of changes in the connection weights; McMurray et al., 2012; McMurray, Horst et al., 2009).

Finally, it is worth emphasizing that this study is quite preliminary with regard to the *direction* of the effect of concurrent semantic similarity on phonological word form learning. Although we found a detrimental effect of concurrent similarity consistently across the behavioral tests and the simulation, it is possible that in some cases concurrent referent similarity might facilitate phonological word form learning. As discussed in the introduction, concurrent referent similarity may also make visual categorization of the referents easier if the overlapped part could form a chunk. Further research is needed to determine under what conditions concurrent referent similarity helps phonological word form learning and under what conditions it hinders it.

CHAPTER 6 PHONOLOGICAL SIMILARITY AND SPOKEN WORD PROCESSING AND LEARNING

This chapter will survey the literature with an aim to set the basis for the dissertation study examining the effect of concurrent phonological similarity on word form learning. This chapter will first survey the studies on how phonological similarity influences spoken word processing. As will be illustrated below, a growing theme in this field is to take into account of the temporal distribution of phonological similarity and to examine closely how the competition dynamics of phonological neighbors evolve over time. This chapter will also summarize the studies on phonological similarity's effect on novel word learning and evaluate their quality in methodologies.

6.1 The Effects of Phonological Similarity in Spoken Word

Processing

This section will survey two lines of research, those on the effect of phonological neighbors on the spoken word processing, and those on the phonological similarity effect in immediate serial recall.

6.1.1 The effect of phonological neighbors

A rich body of studies has showed that processing of a given word depends not only on its own properties, but also on the other words that are phonologically similar to it. Luce & Pisoni (1998) proposed a construct called neighborhood density to capture one aspect of the similarity matrix of phonological representations. A phonological neighbor is a word that differs from the target word by a single phoneme (phonological neighbors have also been defined in different ways, for example, see Magnuson et al. (2003)). Neighborhood density refers to the total number of phonological neighbors, which if weighted by the frequency of each neighbor, becomes frequency-weighted neighborhood density. Generally, neighborhood density was found to be detrimental in spoken word

recognition, implemented by tasks such as lexical decision, auditory shadowing and same-different judgment (e.g., Luce & Large, 2001; Vitevitch & Luce, 1998, 1999). In contrast, a facilitative effect of neighborhood density was seen in word production, usually implemented by a picture naming task (e.g., Vitevitch, 2002a; Vitevitch & Sommers, 2003).

So far, two theories have been proposed to account for this discrepancy, both resorting to real-time processes such as the parallel and interactive activation of the phonologically similar words. The level of processing theory (Vitevitch & Luce, 1998, 1999) assumes that the recognition tasks emphasize the processing at the lexical level, and thus words with many neighbors suffer more lexical competition. In contrast, the production tasks emphasize the processing at the sub-lexical level, and thus words with many neighbors receive more activation support due to sharing phonological segments with the neighbors. In a similar vein, a recent interactive activation and competition account (Chen & Mirman, 2012, see also Dell & Gordon, 2003) argues that the recognition tasks are phonologically driven and thus the phonological neighbors are strongest competitors. However, the production tasks are semantically driven and thus the strongest competitors are the semantic neighbors, not the phonological neighbors. Instead, the weakly active phonological neighbors boost activation of the target, helping it to overcome the competition from the semantic neighbors.

Some studies took into account the temporal distribution of phonological neighbors and examined separately the effects of sub-types of neighbors that overlap with the target at different locations. In particular, onset neighbors were the main focus, motivated especially by the theoretical proposals that the word onset is more important in lexical access (e.g., Marslen-Wilson, 1987; McClelland & Elman, 1986). In both a shadowing and a lexical decision task, Vitevitch (2002b) demonstrated that dense onset neighborhood led to poorer word recognition compared to sparse onset neighborhood

when the overall density was controlled for. Importantly, a detrimental effect of onset neighbors was also seen in a production task (Vitevitch et al., 2004), contradicting the finding that general neighbors facilitate production (e.g., Vitevitch, 2002a). Some studies, harnessing the eye-tracking technique, provided information about the continuous activation of the target and its phonological competitors as the acoustic information unfolds. It is generally found that a cohort neighbor competed with the target at the early time window of stimulus presentation and then died out (e.g., Allopenna et al., 1998).

A recent study calculated frequency-weighted cohort density in addition to the frequency-weighted neighborhood density (Magnuson et al., 2007). It was shown that onset density had an early inhibitory effect on the activation of the target in a recognition task. However, overall neighborhood density had an early facilitatory effect and a late inhibitory effect. Closer examination of the stimuli revealed that the low neighborhood density stimuli contained more cohort neighbors than those high density ones. It means that the high neighborhood density words suffered less competition at the beginning of the stimulus presentation but more competition at the later point, when all the neighbors, especially the rhyme neighbors were also active and competed for recognition. In sum, these recent studies suggest that the effect of neighbors depend on the temporal distribution of similarity and that the cohort neighbors are especially harmful in spoken word processing. This is a point we will return to in examining the effects of phonological similarity in word learning.

6.1.2 Phonological similarity effect in immediate serial recall

Studies on the phonological similarity effect (PSE) in immediate serial recall provided converging evidence that similar words influenced the processing of each other and the influence was dependent on the location of the phonological overlap (e.g., Fallon, Groves, & Tehan, 1999; for a review, see Gupta, Lipinski, & Aktunc, 2005). Gupta et al.

(2005) provided a systematic examination of PSE as a function of type of similarity. In this study, alliterative similarity (cohort similarity) referred to sharing the beginning part among the words (e.g., bib big bill bin bit), rhyme similarity referred to sharing the end part (e.g., bale male pale kale sale), and canonical similarity referred to sharing phonemes in random positions (e.g., cab gab fad gag nan). The results showed that the rhyme similarity was facilitative, the canonical similarity was detrimental, and alliterative similarity had no effect. Gupta et al. (2005) argued that there were two processes going on in serial recall: 1) using the overlapped the segment as the category cue to aid recall (also suggested by Fallon et al. (1999)), and 2) interference due to the overlap. The final effect of the phonological similarity was the sum of the two opposing effects. The rhyme similarity was beneficial probably because rhymed items not only possessed a category cue to aid recall but also were easy to discriminate; alliterative produced no effect probably because the effect of having a category cue and the difficulty in discriminating the items canceled out each other; canonical impaired item recall probably because it provided no category cue but only obstacle to discrimination.

Similar patterns as those in Gupta et al. (2005) were seen in nonwords as well. Using three-syllable nonwords, Service & Maury (2003) compared the effects of the phonological similarity on immediate serial recall when the items shared the first, the second or the last syllable. Error rate as the function of overlapped position revealed a general pattern: a helpful effect on the recall of the redundant syllables and a harmful effect on all the other syllables. However, when the nonwords shared the first syllable, the beneficial effect on recalling the first syllable was less than the harmful effect on the other two syllables; when they shared the last syllable, the condition was the opposite, the large beneficial effect on recalling the last syllable was able to compensate for the small harmful effect on the other two syllables; when the items shared the middle syllable, the helpful and the harmful effects balanced out. Service & Maury (2003) also argued that

since the beginning part of a word carried the largest weight to a word's unique identity, sharing the first syllable led to most confusion among the items.

6.1.3 Summary

In sum, the existing studies establish that a spoken word interacts with words that are phonologically similar to it. A general pattern across studies of different research fields was that words that shared initial sounds with it (cohorts) were more detrimental than those sharing the final ones (rhymes). This is consistent with the sequential nature of spoken word processing (e.g., Gupta et al., 2005; Marslen-Wilson, 1987; McClelland & Elman, 1986; Sevald & Dell, 1994). The incremental unfolding of the input determines that the cohorts are active early when the target has not been highly active, while the rhymes are active at a later point when the target is very active. Thus, the former competes strongly with the target but does not receive much competition from it, while the latter compete less strongly with the target but receives strong suppression from it. Related to this, another important new finding is that the effect of phonological neighbors on processing a given word evolves in real-time, presumably depending on how competitor dynamics changes over time as the acoustic input is received (Magnuson et al., 2007). These have important implications on the studies addressing the effect of phonological similarity on novel word learning. First, it is likely that different kinds of phonological similarity (e.g., cohorts vs. rhymes) may have different effects (qualitatively and/or quantitatively) on novel word learning. Second, the effects of a given kind of phonological similarity on novel word learning of novel words might also be different at different positions of a novel word depending on how the competition it receives from the similar words changes over time.

In addition, the existing studies also suggested that phonological similarity can be beneficial by providing a categorical cue for retrieval (e.g., Gupta et al., 2005) or making the phonological segments to be recalled more active and thus more accessible (e.g.,

Vitevitch, 2002a). Therefore, the final effect of phonological similarity depends on the relative size of the two opposing effects. This may apply to the effect of phonological similarity on novel word *learning* as well.

6.2 The Effect of Phonological Similarity on Word

Learning

This section surveys two lines of research, those on the effect of phonological neighborhood density and phonotactic probability on novel word learning, and those on infants word learning studies using minimal word pairs.

6.2.1 The effect of phonological neighborhood density and phonotactic probability

Some studies have addressed whether phonological similarity to known words influences the learning of novel words. In addition to phonological neighborhood density, the effect of phonotactic probability was also addressed. Since phonotactic probability refers to the probability of a certain phonological segment occurring in the entire lexicon, it can be regarded as the overall similarity to all the words in the lexicon. Most studies in this vein did not manipulate neighborhood density and phonotactic probability independently. Thus, those studies that claimed to address the effect of phonotactic probability may in fact have tapped the effect of neighborhood density or the joint effect of both factors, and vice versa. However, the advantage of these studies might be that they represented the word learning situation in real life more closely since these two variables are naturally associated.

The empirical studies consistently reported that high phonotactic probability (high neighborhood density at the same time) facilitated novel word learning of infants (Graf Estes, Edwards, & Saffran, 2011), pre-schoolers (Storkel, 2001) and school-age children (Storkel & Rogers, 2000). It was interpreted that the novel words with frequent segments

received more support from the already established representations, and thus would be more active in short-term storage and also be easier to be consolidated into long-term memory. In addition, due to the ease of encoding the phonological word forms, there would be more processing capacity left to encode semantics and the association between the word forms and the referents (Graf Estes et al., 2011; Storkel, 2001).

However, the studies that claimed to study the effect of neighborhood density found less consistent results. Gupta, Newman, Samuelson, & Tisdale (2010) taught toddlers aged 20 and 24 months a high density nonword “wat” (/wæt/) and a low density one “fowk” (/faʊk/), with each paired with a novel referent. A looking-while-listening measure revealed that the 1.5-year-old toddlers learned the high-density word better than the low-density one. A seemingly contradictory effect was found by Swingley & Aslin (2007), who compared word learning between novel words that were highly similar to familiar words (novel neighbors) and those that were not similar to any familiar ones (non-neighbors). An example of the novel neighbors is “tog”, which is similar to “dog”, and an example of non-neighbors is “meb”, which does not have any word neighbors known to participants in this study. The proportions of fixation revealed that infants did not learn the novel neighbor as well as they learned the non-neighbors. In addition, learning a novel neighbor seemed to have impaired the recognition of the familiar neighbor. Thus, contrary to the finding of Gupta et al. (2010), this study suggested that having a known neighbor impaired the learning of a novel word. Gupta et al. (2010) argued that this inconsistency lay in the structure of the neighborhood. The novel word in Swingley & Aslin (2007) had a single strong word neighbor while the novel word in Gupta et al. (2010) had multiple moderate neighbors. Using a connectionist model, Gupta et al. (2010) demonstrated that due to the well-established mapping from its phonological word form to its semantic region, a single well-learned neighbor tend to entrap the novel word form to its own semantic region. However, when there were a number of neighbors,

these word forms mapped to broader regions of the semantic space. Thus, when there was the need to map a similar novel word form to its own semantic representation, the broad range of activation in the semantic space instead served a scaffolding function.

There have been some studies that manipulated lexical neighborhood and phonotactic probability orthogonally (Storkel, Armbruster, & Hogan, 2006; Storkel & Hogan, 2010; Storkel & Lee, 2011). Across these studies, a common finding is the disadvantage of high phonotactic probability. This pattern is the opposite of the findings in the above studies where the phonotactic probability was confounded with neighborhood density. It is likely that the effect of phonotactic probability observed in the former set of studies was contaminated by the effect driven by neighborhood density. The effects of neighborhood density revealed by Storkel and her colleagues showed no clear pattern and seemed to be dependent on the learning context and the learning-test interval. In addition, the authors provided interpretations that were essentially post-hoc.

In sum, the existing studies on the effect of phonological neighborhood density and phonotactic probability showed a messy picture. Most of them confounded these two variables, and for those that did not, the findings were mixed and the interpretations were not convincing. More importantly, the effects might arise at test, driven by on-line interactions between novel words and its known word neighbors rather than driven by learning. For example, the fewer fixations to the novel neighbor (e.g., ‘tog’) in Swingley & Aslin (2007) might be due to competition from the known neighbor (e.g., ‘dog’) rather than due to being learned less well. Therefore, studies with clear a priori predictions and tests that eliminate on-line processing confound at test are needed, so as to reveal a clearer picture of how similarity to known words might influence the learning of real words.

6.2.2 The effects of different types of phonological neighbors

Some studies took into consideration the temporal distribution of phonological similarity and examined the effects of different types of phonological neighbors on the novel word learning and recognition (Creel & Dahan, 2010; Magnuson et al., 2003). These studies provided evidence that the overlap at the word onset might be more harmful than the overlap at the end. Magnuson et al. (2003) designed an artificial lexicon that consisted of novel words that were neighbors to each other. In addition, the type and the frequency of neighbors were precisely controlled. A visual world paradigm was adopted as the learning task, in which four pictures were presented along with the auditory word form: the referent of the target, the referent of the cohort neighbor, the referent of the rhyme neighbor, and the referent of the unrelated control. Feedback was given after participants picked a picture as the referent of the word form. The eye-tracking method revealed that at the first day of training both rhyme and cohort neighbors received an equivalent proportion of fixations and showed a similar time course. With more exposure, however, the activation of rhyme neighbors diminished and showed delayed onset compared to that of the cohort neighbors. This finding can be interpreted as a learning effect. That is, the cohort neighbors were learned better than the rhyme neighbors and the difference in their representations led to the difference in their activation levels. However, it can also be explained by on-line processing dynamics. That is, as revealed by the previous non-learning studies (e.g., Allopenna et al., 1998; Magnuson et al., 2007; Vitevitch, 2002b), the cohort neighbors became active at the earlier point of stimulus presentation when the target itself was not strong enough to suppress it.

Creel & Dahan (2010) examined the effect of different types of phonological similarity on novel word learning. Each novel word was paired with a picture of a

nonsense shape. In the recognition test, a picture of a shape was presented together with either its label or a different word form, and participants were asked to report whether the pairing was original or rearranged. Critically, for the rearranged pairs, the lure label overlapped with the original label of the shape either at the initial sounds (“joop” and “joob”) or at the final sounds (“joop” and “choop”). Greater false alarm rate was found for the onset-overlapping condition than for the offset-overlapping condition. It was further ruled out that the difficulty for the onset match condition was due to discrimination difficulty. According to the authors, the results suggested that learning two initial overlapping words was easier than learning two end overlapping words, because the onset of a word was given more attention. However, this effect could also be driven by real-time processing. That is, the onset-matched lure may have activated the target label to a greater extent than the end-matched lure, thus leading to higher rate of false alarms.

In sum, both Magnuson et al. (2003) and Creel & Dahan (2010) suggested that the types of phonological similarity might matter in novel word learning. However, both studies could not rule out the possibility that these effects were due to some real time processes at test, such as co-activation and competition.

6.2.3 Infants’ word learning using minimal pairs

Studies on infants’ word learning using minimal pairs have typically been framed as examining whether infants are able to use the phonemic distinctions in word learning tasks. However, since these studies used phonologically similar words, they also spoke to the question of whether phonological similarity influences word learning.

Stager & Werker (1997) used the word-object association task to investigate whether 14-month-olds could map similar sounding word forms (“bih” and “dih”) to distinct objects. In this task, infants were first habituated to two novel word-object combinations. Following habituation, infants were tested with a switching paradigm that

included two types of trials: “same” trials, in which the original word-object pairs were presented, and “switch” trials, in which the learned words and objects were presented in novel pairings. The results showed that 14-month-olds looked equally long to the same and the switch trials, indicating that they failed to learn these words. This failure of learning similar sounding words was in contrast with the ability of infants at this age to discriminate between phonemic contrasts in their native language (for a review, see Saffran, Werker, & Werner, 2006). To explain this paradox, different accounts were proposed. In the following, I classify the various accounts into three categories: resource limitation, information, and competition.

6.2.3.1 Resource Limitation

The gist of this account is that the failure revealed by the switch test lies in difficulty of the word learning task. Stager & Werker (1997) argued that the word-object association task requires more processing than simple discrimination, such as referent discrimination and making associations between word forms and referents. The increased processing demands in word learning thus tax the limited processing capacity of young infants, preventing them from using their full phonological sensitivity. It could also be understood that learning similar word forms at the same time increases the demands for phonological discrimination and thus hinders word learning (learning of the referents, the word forms, and their associations). Fennel (2012) tested the resource limitation account directly. The authors hypothesized that if the failure to learn the highly similar novel words was because of the overwhelming demands of the word learning task, then familiarizing infants with the object referents should reduce the overall processing demands and thus should improve their performance. The results showed that only those 14-month-old infants who were pre-exposed to the object successfully detected the mismatch between the referent and the word for the switch trials, lending support to the resource limitation account.

6.2.3.2 Informational Account

This information account was my summary from a hybrid of studies addressing the above-mentioned paradox in different ways (Rost & McMurray, 2009; Thiessen, 2007). I would like to argue that these studies suggest one common mechanism: information that is useful to distinguish between similar word forms is helpful to map them onto different referents.

Rost & McMurray (2009) argued that 14-month-olds might not have grasped the phonetic distinctions fully even though they are able to discriminate between phonemic contrasts in certain tasks. In addition, this immature phonemic representation coupled with a single exemplar for a label might be the reason why the 14-month-olds failed at mapping phonologically similar words to distinct referents. However, if multiple exemplars of labels were used, the variations in the tokens may provide useful information to augment or to maintain the representations of phoneme categories. Indeed, when recordings from multiple talkers for the same words were used, 14-month-old infants were able to detect the sound change for the switch trials. The follow-up study revealed that it was the variations in the phonologically irrelevant indexical and suprasegmental aspects that helped (Rost & McMurray, 2010).

Thiessen (2007) argued that experiencing the phonemic contrasts in different lexical contexts would help infants grasp the distinctions and regard them as functional distinctions that differentiate meanings. In this study, after learning “dawgoo” and “tawgoo”, 15-month-olds failed to detect that “daw” was different from “taw”. However, when “daw” and “taw” were presented in different lexical contexts (daw, dawbow, and tawgoo), infants recognized the difference. Thus, experiencing “daw” in different contexts (daw and dawbow) and experiencing “daw” and “taw” in different contexts (dawbow and tawgoo) enhanced infants’ ability to distinguish between “daw” and “taw”. The authors argued that this effect was similar to acquired distinctiveness (Lawrence,

1949). During the switch test, the presence of “daw” activated the lexical contexts where it appeared in (daw and dawbow). Since these contexts were different from where “taw” appeared in (tawgoo), it was less likely for “taw” to be highly activated and then to cause confusion (Rost & McMurray, 2009).

To sum up, these studies converged on the suggestion that information useful for distinguishing between similar word forms is helpful to map them onto different referents, no matter where the distinctive information comes from.

6.2.3.3 Competition and co-activation

Swingley & Aslin (2007) compared novel word learning when the phonological label was a neighbor to a familiar word (novel neighbors) versus not (non-neighbors). Strictly speaking, this study is different from the above studies using minimal pairs because it concerns phonological similarity between a novel word (the known word neighbor) and a known word (the novel neighbor). However, it bears deep similarity to these studies because these two words presented together during learning and were contrasted directly during testing. The results showed that having a known word neighbor impaired the recognition performance of the novel word. In addition, the performance of the known word neighbor was impaired as well. The authors suggested that the competition between the novel word and its known word neighbor was the underlying mechanism. However, the authors were vague about when competition occurs and whether competition affects the test performance or learning, or both. Yoshida, Fennell, Swingley, & Werker (2009) proposed explicitly that some process that happens at the switch test masks the learning which could otherwise be shown. Specifically, Yoshida et al. (2009) argued that due to the phonetic overlap, the correct label is also active in a switched trial, making this kind of trial comparable to a non-switched trial (for similar discussions, see Apfelbaum & McMurray (2011)). In fact, Yoshida et al. (2009) showed that when both referents were presented with a label, 14-months showed more looking

towards the correct referent, suggesting that they successfully encoded the phonetic difference and mapped the two similar labels to distinct objects.

6.2.3.4 Summary

Studies in this section showed that learning similar sounding words at the same time was a challenging task for young infants. The challenge may come from multiple sources: increased demand for processing capacity, less developed phonemic categories, competition between the similar words, and some real time processes happening at test. It is important to note that although most studies in this field have contributed to a failure at learning due to limited resources or improper learning of the phonological categories, some studies have linked this failure to the processes occurring in the switch test.

6.2.4 Summary on the word learning studies

The existing studies have provided preliminary evidence that phonological similarity might influence novel word learning. However, there was a big confound in most of these studies, making it unwarranted to conclude that phonological similarity under study indeed influenced the learning of the novel words. As demonstrated by Yoshida et al. (2009), the effect of phonological similarity observed in these novel word learning studies could simply arise from a real-time processing difference between conditions at testing, rather than be driven by how well these words have been learned. Thus, most studies so far on this topic have ignored the extensive evidence that phonological similarity influences real-time processing of spoken words (as reviewed in the previous section) and mistakenly assuming that the effect of phonological similarity at test can be cleanly attributed to learning. The same as the case of examining semantic similarity's effect on novel word learning, this is another instance of a broader assumption that test performance is a direct read-out of underlying knowledge, and that real-time processes are not differentially affected by the empirical manipulations under

study. This dissertation strives to tap the effect of phonological similarity on *learning*, by designing a test in which the online processing is controlled for between the conditions. In the following chapter, the empirical research on this question will be summarized.

CHAPTER 7 THE EFFECT OF PHONOLOGICAL SIMILARITY ON CONCURRENT NOVEL WORD LEARNING

The central question to be addressed in the second part of the dissertation is how phonological similarity might influence the learning of a set of novel words concurrently. To our knowledge, none of the existing studies has addressed the same question. Some studies examined on a different kind of similarity, i.e., similarity to existing words (e.g., studies by Storkel and her colleagues). Others examined the concurrent word learning situation, but used tests that tap things different from how learning is affected by similarity, e.g., the co-activation of the neighbors (Magnuson et al., 2003) and confusability between phonologically similar words (e.g., Creel & Dahan, 2010, Stager & Werker, 1997). As discussed in Chapter 5, some temporal dimensions are important to consider in this study because the competitor dynamics that may influence novel word learning might be different depending on these factors. One dimension is at which position in a word the similarity lies. Previous research showed that sharing the initial sounds (cohorts) is more detrimental than sharing the final sounds (rhymes) in spoken word recognition (e.g., Allopenna et al., 1998; Magnuson, Dixon, Tanenhaus, & Aslin, 2007; Vitevitch, 2002), production (e.g., Vitevitch et al., 2004), and also immediate serial recall (e.g., Sevald and Dell, 1994; for a review, see Gupta et al., 2005), presumably because the cohorts are stronger competitors when the target is relatively weak in activation while the rhymes start to compete when the target is highly active. Therefore, when it comes to examine the effect to phonological similarity on novel word learning, the picture will be clearer if only one kind of similarity is focused on in one study. Here as a starting point we focused on the cohort similarity.

The second dimension to consider is at which position in a word that learning is examined. One study found the effect of phonological neighbors on spoken word recognition is different at different positions in a word, depending on how the competitor

dynamics that a word receives evolves over time as the acoustic input is received (Magnuson et al., 2007). A recent study (Apfelbaum & McMurray, 2012) showed that more spurious associations between words and objects were acquired when the object was presented at the same time with the word (the synchronous condition) than when the object was presented after a delay (the delayed condition). This is presumably because at the beginning of a word's presentation, the cohorts are partially active, while after a delay they are less active and suppressed by the strongly active target. These findings suggest that learning occurs in the real time rather than waiting for the competition to be resolved, and that learning occurs depends on the status of the lexical competition at a given moment. Based on these findings, this dissertation study used nonwords with three syllables and examined the effect of cohort similarity at the middle and the final syllable separately. We predict that the detrimental effect of cohort similarity will decrease from the middle syllable to the third syllable as cohort competition decreases over time. Relatedly, the accuracy of nonword repetition was found to be lower at the middle of a nonword than at the beginning and the end (Gupta, 2005; Gupta, Lipinski, Abbs, & Lin, 2005), suggesting that short-term memory of the middle syllables is poorer than that of rest ones. In a three-syllable novel word, this may make the middle syllable of a word in a cohort condition even more vulnerable to competition from and the third syllable more resistant to it.

However, it needs to be noted that sharing initial sounds can also be beneficial, because shared segments may not only provide a category cue for encoding and retrieval (e.g., Gupta et al., 2005), but also boost the activation of the cohorts via activation spreading (e.g., Vitevitch, 2002a). Therefore, at a given time position the effect of cohort similarity will be the sum of these beneficial effects and the detrimental effect from cohort competition. Since there is no known basis to estimate the relative size of these two opposing effects, this study does not make strong predictions on the direction of the

net effect. However, as there is a strong evidence to expect that the detrimental effect decreases over time, we predict that the effect of cohort similarity is more likely to be detrimental at the middle syllable than at the end syllable (although this is likely to be wrong if in fact the cuing effect decreases in a large amount from the second to the third syllable).

To actually tap the learning effect, as discussed in the previous chapters, a critical potential confound — real-time processing difference at test — must be avoided. In fact, this confound was present in all the existing studies addressing similar questions, and thus the effect of phonological similarity they observed could simply arise from the real-time processing difference at test. For example, a novel word with a dense neighborhood could show a poorer naming performance due to suffering greater competition from the neighbors at test, even though it has been learned equally well as the words with a sparse neighborhood. Thus, a crucial goal of this study is to develop a test of word learning that controls for real-time effects between different conditions of phonological similarity and thus represents solely learning effects.

We developed a test of recognition from mispronunciation (e.g., Leach & Samuel, 2007). The most important feature of our version of this test is that it focused on the rest of the nonword without the first syllable (e.g., /neɪbʊd/ in /kɪ'neɪbʊd/). In this test, the correct partial word was presented in a sequence together with two fillers (e.g., /peɪbʊd/, /neɪtʊd/) that differed from the target partial word by only one phoneme, and participants were asked to indicate which stimulus sounded most familiar to them. Having the first syllable removed should to a large extent prevent the co-activation of the other words in the cohort condition and thus control for competition between the two conditions.

Therefore, this study used this recognition from mispronunciation test to examine the effect of the cohort similarity on the novel word learning. Two experiments were conducted. Experiment 4 examined the effect at the middle syllable (the second syllable

in this experiment), and Experiment 5 examined the effect at the end syllable (the third syllable).

Experiment 4

This experiment aimed to examine the effect of cohort similarity, on the learning of the middle syllable of the novel words that are concurrently learned. As discussed in the introduction, sharing the word onset can lead to two opposing effects, with a beneficial effect being due to the shared segments boosting the activation of the cohorts and the shared segments being used as the categorical cue for encoding and retrieval, and a detrimental effect due to competition from cohorts. Since there is no known basis to compare the relative size of these two effects, we do not hold a strong prediction on the final effect. However, there is a solid basis to expect that a detrimental effect of cohort similarity is more likely to be observed at this position because of higher level of competition as this segment is being received, coupled with the weak short-term storage of it.

In this experiment, cohort similarity was manipulated between two groups of participants. For each condition, participants learned a set of three-syllable nonwords, with each paired with an alien picture as its referent. To encourage learning, they received a naming test periodically (e.g., Karpicke & Roediger, 2008). The naming test was simply used to facilitate learning rather than to measure it, since as discussed earlier an effect of cohort similarity in naming can be due to difference in real-time processing at test. The learning of the novel words was instead evaluated by the recognition from mispronunciation test at the end of training. As introduced earlier, this test focused on the partial word after the first syllable, thus it largely eliminated the at-test co-activation of the other words in the cohort condition and controlled for interference between the two conditions. Thus any effect of similarity should be a relatively pure learning effect. To test the learning at the second syllable, the two fillers were thus mispronounced at the

first phoneme of the second syllable (e.g., /se bud/, /pe bud/ for /ne bud/ in /kI'ne bud/). If cohort similarity affects learning, recognition accuracy should differ across similarity conditions, with the tentative prediction being for lower recognition accuracy in the cohort condition.

Methods

Participants

Thirty-two undergraduate students at the University of Iowa were recruited from the department subject pool. They were randomly assigned to one of two conditions: cohort and control.

Stimuli

For the targets to learn, there were two sets of eight three-syllable nonwords (CV-CV-CVC structure) that were within-set cohorts, and two sets of eight dissimilar control nonwords that were yoked to those in the cohort sets (for examples, see Table A-2 in Appendix A). Specifically, for the paired sets, each cohort nonword had a counterpart control that had the same second and third syllables. As will be described later, this is to achieve the best match for the partial words used in the recognition from mispronunciation test between the two conditions. The eight nonwords in a cohort set all started with the same initial consonant, and there were some that shared one or two more phonemes. The eight nonwords in a control set shared no initial consonant, but similar to those in the yoked cohort set, some of them shared the first vowel. The cohorts and the controls were matched on positional probability, biphone probability, and neighborhood density (all $p_s > .30$). In addition to the targets, thirty-six nonwords of the same structure were randomly chosen as fillers from a pool of ninety.

A set of eight partial words was extracted from a yoked pair of target sets. A partial word consisted of the second and the third syllables of a three-syllable target. In

addition, for each partial word, two foils were created that differed from the original one at the first phoneme.

Drawings of humanoid aliens were used as referents (see Gupta et al., 2005) (for examples, see Figure A-1 in Appendix A).

Procedures

The learning paradigm

The task consisted of four learning epochs with each followed by a naming test except after the fourth epoch. In a learning epoch, there were eight targets interspersed with nine foils. The targets were presented in every learning epoch, but the foils were unique to a given block. On both kinds of trials, a nonword was presented auditorily, and the participant's task was to repeat it aloud. However, only on target trials, the nonword was accompanied by a picture of an alien in the middle of the computer display, and participant were asked to memorize the nonword as the label of the alien. Each naming test consisted of eight trials. In each trial, an alien image was presented and participants were required to recall its name. The order of the targets was randomized in each of the learning and the testing epochs.

The recognition from mispronunciation test

There were eight trials in total, with each testing one nonword. In a trial, one original partial nonword and its two foils were presented in a random order. Participants were asked to identify which one of the three sounded most familiar by pressing one of the three designated keys ("1", "2", and "3"). The interval between the two adjacent nonwords was 1000ms, and the inter-trial interval was 2000ms. The presentation order of the eight trials was randomized across participants.

Results and Discussion

The recognition accuracy (see Figure 7-1) was first subjected to arcsine transformation and then to a One-Way ANOVA¹⁰. Recognition accuracy in the cohort condition was found to be significantly lower than that in the dissimilar condition, $F_1(1,30) = 5.02$, $\eta_p^2 = .14$, $p < .05$, $F_2(1,15) = 6.99$, $\eta_p^2 = .32$, $p < .05$. This indicates that cohort similarity among a set of concurrently learned words impairs performance in recognizing a novel word form from highly confusable phonological distractors that differ from it at the second syllable. Because this recognition test used the partial words without the initial sounds, it largely eliminated the at-test interference between the cohorts and control for this processing difference between the two conditions. Thus, this finding suggests that concurrent referent similarity impairs the learning of the second syllables of phonological word forms, making them less discriminable from distractors. This suggests that the competition from cohorts at this position produced stronger effect on the learning of this segment, thus overriding any beneficial effect due to categorical cuing and activation boosting from sharing the onset.

¹⁰ In an initial analysis, the set of nonwords was also included as a between-subjects factor. Neither the main effect of nonword set nor the interaction between nonword set and similarity was significant. Thus, the analysis reported here collapses across nonword sets.




Figure 7-1. Recognition accuracy in the cohort and the dissimilar conditions.

Although naming test could not provide convincing evidence on learning, it could provide information about how a novel word behaves in this canonical word use situation. Thus, naming accuracy (see Figure 7-2) was also analyzed using a 2 (condition: cohort, dissimilar) by 3 (epoch, 1-3) Mixed ANOVA model. There was a significant effect of epoch, a marginally significant effect of condition, and a significant interaction between these variables, $F(1,30) = 13.86$, $\eta_p^2 = .32$, $p = .001$, $F(1,30) = 3.18$, $\eta_p^2 = .10$, $p = .085$, and $F(1,30) = 5.61$, $\eta_p^2 = .16$, $p < .05$ respectively. Simple effect analysis revealed that although the two conditions did not differ at the first two epochs (both $p > .50$), at the third epoch, the naming accuracy of the cohort condition was significantly lower than that of the dissimilar condition, $F(1,30) = 5.61$, $\eta_p^2 = .16$, $p < .05$. This suggests that the ability of producing newly learned labels is impaired by cohort similarity. As discussed

before, it is difficult to tell whether this naming effect is due to a learning difference or a difference in real time processing. However, the finding from the recognition test suggests that the learning difference is one potential contributor to this effect.




Figure 7-2. Naming accuracies across the first three epochs in the two conditions.

Experiment 5

This experiment aimed to examine the effect of cohort similarity, on the learning of the final syllable of the three-syllable novel words that are concurrently learned. Again

we do not have strong predictions on the effect, but do expect a lower chance to observe a detrimental effect as seen in Experiment 4 since competition among cohorts is expected to be largely resolved and short-term storage of this segment is better.

Methods

Participants

Thirty-two undergraduate students at the University of Iowa were recruited from the department subject pool. They were randomly assigned to one of two conditions: cohort and control.

Stimuli

The same as those in Experiment 4 except that the foils differed from the original one by one phoneme at the 3rd syllable.

Procedures

The same as those in Experiment 4.

Results and Discussion

The recognition accuracy (see Figure 7-3) was first subjected to arcsine transformation and then to a One-Way ANOVA. Recognition accuracies of the two conditions did not differ significantly from each other, $F_1(1,30) = .11$, $\eta_p^2 < .01$, $p > .70$, $F_2(1,15) = .53$, $\eta_p^2 = .03$, $p > .40$. This suggests that cohort similarity does not have an impact on the learning of the last syllable of the phonological word forms, presumably because any beneficial effects and the detrimental effect cancelled out each other.




Figure 7-3. Recognition accuracy in the cohort and the dissimilar conditions.

The naming accuracy (see Figure 7-4) was subjected to a 2 (condition: cohort, dissimilar) by 3 (epoch, 1-3) Mixed ANOVA. There was a significant effect of epoch, a marginally significant effect of condition, and a significant effect of the interaction between these variables, $F(1,30) = 22.70$, $\eta_p^2 = .43$, $p < .001$, $F(1,30) = 3.64$, $\eta_p^2 = .11$, $p = .07$, and $F(1,30) = 8.53$, $\eta_p^2 = .11$, $p < .01$, respectively. Simple effect analysis revealed that although the two conditions did not differ at the first two epochs (both p s $> .70$), at the third epoch, the naming accuracy of the cohort condition was significantly lower than that of the dissimilar condition, $F(1,30) = 6.55$, $\eta_p^2 = .18$, $p < .05$. The results from the naming test replicated those in Experiment 4. Since no difference was observed in the recognition test between the two conditions, it suggests that the cohort detriment in the naming test may only be driven by a real time processing difference.

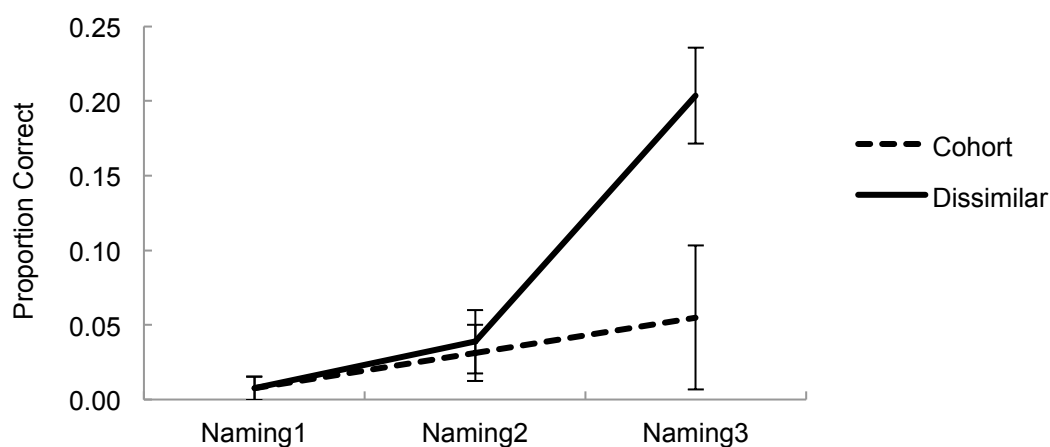


Figure 7-4. Naming accuracy in the cohort and the dissimilar conditions.

General Discussion

Experiments 4 and 5 aimed to examine the effect of cohort similarity on the learning of phonological word forms for a set of concurrently learned words. Here we define learning as an increase in strength of the representation acquired (or knowledge), which is not equivalent to test performance that is usually multiply determined (the factors that influence it can be both learning and real time processing). Although we are not the first to make a distinction between learning and test performance, we for the first time carried out an empirical study under this understanding. Specifically, we designed a

recognition from mispronunciation test that controlled for real-time processing between the cohort condition and the control condition so that the performance difference between these conditions can truly reflect the effect on phonological word form learning.

This study is also novel in that it for the first time examined the effect of the phonological similarity on the novel word learning at different positions in a word. It provided preliminary evidence that the effect of cohort similarity varies across the syllables, a detrimental effect at the middle syllable and non-effect at the final syllable. This finding is consistent with the evidence that competition from cohort neighbors is stronger at word onset and diminishes as more information about the word form is received. Thus, this finding suggests that novel words that are similar in phonology are engaging in competition as they are being learned together, and sleep-based consolidation is not necessary for competition between them to occur. This is consistent to a recent finding that competition between a newly learned word and a phonologically similar known word occurs right away after only a brief exposure (Kapnoula, Packard, Apfelbaum, McMurray, & Gupta, 2012), and extends it to the situation of concurrent word learning. This finding suggests that lexical inhibition does not require consolidated representation and even partially established representations are capable of doing it.

This finding is consistent with another recent finding that word-referent associations acquired depend on the status of competition between phonologically similar words (Apfelbaum & McMurray, 2012). Both studies suggest that word learning occurs in the real time, depending on the status real-time processing right at that moment. These two studies also imply that future studies on examining the effect of phonological similarity in word learning should consider the evolution of real time processing (such as competition) over time and that different constituent segments in a word form should be examined separately so as to capture the specific effects of similarity at different temporal positions in a word.

However, this study is only a preliminary exploration. For one thing, different syllables were used for the second and the third syllables of the target nonwords, and thus it might be the differences in acoustic and phonetic properties between the two sets of syllables that made the effect of cohort similarity position dependent. Thus, one future direction is to use exactly same syllables across the positions so as to rule out this alternative explanation.

A simulation study is another important future direction. The same as the behavioral tests used to examine the effect of referent similarity, the recognition test here also could not completely match for real time processing between the cohort and the control conditions. It is likely that the partial word form presented at test led to activation of the whole word, which then led to co-activation of its cohorts. If that is the case, then the detrimental effect of cohort as revealed by the recognition test might still not be due to learning, but instead due to higher competition in the moment of test. As demonstrated in Chapter 5, a computational model, such as the HNRN, is able to isolate learning and on-line processing and thus provide a pure measure of learning. In the near future, another modified version of HNRN will be developed to provide more convincing evidence on the effect of cohort similarity on novel word learning. The most critical modification of this model will be to incorporate the sequential nature of phonological processing. One way to implement it is to present the input of phonemes or syllables at different positions at different time points. For example, the first phoneme can be presented together with the referent input, the second phoneme can be presented after ten cycles of processing, and the third one can be presented after twenty cycles of processing. Cohort similarity can be easily manipulated as sharing the initial phoneme between a set of words, and the dissimilar control words will have distinct initial phones. Importantly, the learning of each phoneme could be indexed separately by the connection weights between a given phoneme and the word unit that is chosen to bind the word form and the

referent, which can then be compared between the cohort condition and the dissimilar condition.

CHAPTER 8 GENERAL DISCUSSION AND FUTURE DIRECTIONS

This dissertation work addressed the effect of referent similarity and the effect of phonological similarity on the learning of a set of words that are being learned concurrently. In particular, the phonological word form learning was the focus. Chapter 4 and Chapter 5 provided behavioral and computational evidence that converge on the suggestion that the learning of the phonological word forms is impaired by concurrent referent similarity. Chapter 7 addressed one kind of phonological similarity, cohort similarity, and provided behavioral evidence that it influenced the learning of the word forms, and furthermore, influenced learning differently at different syllable positions.

Empirically, this work expands our understanding of word learning and processing in two major ways. First, this study focused on a situation where similarity lies in the novel words that are being learned concurrently, rather than between a novel word and the existing words that were the main focus in previous studies. Therefore, it extends our understanding of the effect of lexical-semantic similarity on novel word learning from one situation to another. Second, this study addressed the effect of lexical-semantic similarity on novel word learning rather than lexical processing, the main focus of the existing research. Thus, it expands our understanding of the effect of lexical-semantic similarity from processing to the learning, suggesting that it not only impacts how the real time lexical processing works, but also leaves a footprint in the internal representation that will lead to an influence in future processing and learning.

Closely tied to the empirical contributions, this study also makes its theoretical and methodological contributions. Here in this work we define learning as representation (or knowledge) acquired from experience and we questioned the assumption that learning can be read out directly from behavioral tests. This is because behavioral performances in a certain task are usually multiply determined (Gathercole, 2006; Gupta, 2008; Gupta & Tisdale, 2009), including both learning and real-time processing factors (McMurray et

al., 2009, 2012). Therefore, real time processes need to be characterized properly so as to make accurate inferences about learning (Horst & Samuelson, 2008; McMurray et al., 2012; Samuelson, Schutte, & Horst, 2009). Due to the above theoretical thinking, we made careful analysis on the real time processing that is potentially involved in behavioral tests, and based on this analysis we designed those tests that can control for real time processing between conditions. In addition, we also made use of computational modeling to provide unambiguous index of learning that is not contaminated by real time processing (although only for the study on the referent similarity). This enabled us, for the first time, to draw convincing conclusions on whether learning itself is affected or not by the lexical-semantic similarity. Our effort here echoes a handful of existing studies on word learning that try to estimate learning (knowledge acquired) more accurately by characterizing real time processing properly. Samuelson et al. (2009) showed that children's ability of generalizing a noun to novel exemplars depend on the specific test used (succeed in a forced choice task, but not in a yes/no task). Similarly, Yoshida et al. (2009) used a visual s choice task instead of a switch-trial test and found that 14 months were capable of mapping similar-sounding labels to distinct referents, suggesting that infants' ability of doing so was masked by the processing demands in the switch test.

The first dissertation study revealed that referent similarity influenced the learning of the phonological word forms. This suggests that the property of one part of a word is capable of influencing the learning of the other part of it. This lends support to the interactive rather than modular viewpoint of the relationship between different components in a word, and extends it to the situation of novel word learning. Thus, the finding here suggests that not only semantics and word forms interact with each other, but also they probably start to interact as soon as there are some partial representations of each component and the associations between them. In addition, this ever-going

interaction is capable to drive the properties of one component to influence the learning of the other.

The second dissertation study revealed that cohort similarity influenced the learning of the phonological word forms and that it influenced it differently at different syllable positions. This finding is consistent with the hypothesis that competition occurs between these novel words that being learned concurrently, but not consistent with the proposal that competition between words will not occur until after some sleep-based consolidation. This finding extends the existing evidence that competition between a newly learned word and phonologically similar known words (e.g., Kapnoula et al., 2012; Swingley & Aslin, 2007) to the competition between novel words whose representations are being built up simultaneously, suggesting that similar-sounding words are becoming integrated with each other in a real-time manner as they are being learned concurrently or as one novel neighbor is being learned.

Some important implications for future studies on word learning can be drawn from this dissertation work. For one thing, this study suggests that we need to discard the assumption that learning can be directly read out from test performance, and that we need to design tests that either characterize real time processes properly or control for them across conditions. Only by doing this, we can be more confident that we are addressing learning as we are claiming. For the other, computational modeling is a very useful tool in studying word learning. Specifically, it explicitly implements multiple processes and also learning so that it enables inference about learning and a real-time processing independently from each other, and also enables observations about how they interact with each other (as in Gupta & Tisdale, 2009; McMurray et al., 2012).

As for major future directions, it is important to re-examine the effect of lexical- semantics similarity between a novel word and its known word neighbors given that the existing studies failed to rule out real time processing confound at test. It will provide us

more convincing evidence to this issue if carefully designed behavioral tests could be used, and ideally also a computational modeling investigation. The second important direction is to examine the influence of phonological similarity on not only the learning of the word forms, but also referents, as a counterpart to our examination of the effect of referents on phonological word form learning addressed in Chapter 4 and Chapter 5. Investigation in this line will provide us richer understanding about the interaction between semantics and phonological word forms during the course of learning. The third one is to follow up the investigation in Chapter 7 to use different kinds of phonological similarity, to expand on examination of how the evolution of real time processing over time may lead to differential effects of word learning at different temporal positions. The fourth one is to examine the effect of lexical-semantic similarity in word learning in a different language. Chinese is a non-alphabetic language that has unique linguistic properties, including image-based orthography, more homophones and fewer syllables in a word. It will be interesting to see how the effect lexical-semantic similarity may affect the learning of novel words differently due to the specific characteristics of this language.

APPENDIX A: EXAMPLES OF STIMULI USED IN THE
DISSERTATION STUDIES

Table A-1 Examples of Nonword Stimuli (IPA) Used in the Behavioral Experiments in Chapter 4

Nonwords	Mis-pronunciations (Exp.3 only)
'fɔɪdʒɛp	'fɔɪtʃɛp
pʌ'kaz	pʌ'gaz
la'dɛg	la'tɛg
'mæfɪs	'mædɪs
zɛ'kɪk	zɛ'gɪk

Table A-2 Examples of Nonword Stimuli (IPA) Used in the Behavioral Experiments in Chapter 7

Cohort Words	Dissimilar words	Partial words	2 nd syl mispronounced	3rd syl mispronounced
'kɛbɛgɑz	'mæbɛgɑz	'bɛgɑz	'hɛgɑz	'bɛtɑz
kɛ'bəʊvɪp	dɪ'bəʊvɪp	'bəʊvɪp	'zəʊvɪp	'bəʊvɪt
'kɛmudəʊk	'kɛmudəʊk	mu'dəʊk	tu'dəʊk	mu'bəʊk
kɪ'sætɪg	pu'sætɪg	'sætɪg	'sætɪg	'sækɪg
kɪ'neɪbʊd	hi'neɪbʊd	'neɪbʊd	'peɪbʊd	'neɪtʊd
kɪ'nʌgɛʃ	bɪ'nʌgɛʃ	'nʌgɛʃ	'lʌgɛʃ	'nʌgɛs
'keɪlɑdɛg	'neɪlɑdɛg	lɑ'dɛg	nɑ'dɛg	lɑ'mɛg
ku'pəʊfʌt	sɪ'pəʊfʌt	pəʊ'fʌt	nəʊ'fʌt	pəʊ'gʌt



Figure A-1. Examples of alien images used as referents in the experiments reported in Chapter 7.

APPENDIX B: PARAMETERS USED IN THE 4-LAYER HNR
NETWORK IN CHAPTER 4

Table B-1: Free parameters in the 4-layer HNR network

Parameter	Value	Symbol
Number of Words	20	n
Visual Category Units	400	m
Lexical Units	500	p
Initial Weight Size (range of random values)	0-.25	
Learning Rate	.00025	η
Decay Rate	.5	δ
Feedforward Temperature	.01	τ_f
Feedback Temperature	2	τ_b
Stability Point	1e-12	
Input inhibition	1.05	ι
Internal inhibition at the visual category layer	2	ω
Internal inhibition at the lexical layer	1.5	Ψ

APPENDIX C: FORMULAS USED IN THE 4-LAYER HNR
NETWORK

At the beginning of a training trial, the auditory and visual feature inputs are given, which are three auditory units (representing the word form) along with one visual feature pattern (representing the referent). Activation in each of the auditory banks and the visual feature layer is then normalized to sum to 1.0.

$$f_x = f_x / \sum f_j \quad (1)$$

Here, f_x refers to visual feature unit x and n is the number of words (equal to the number of features). The same equation is used to normalize the auditory inputs. Next, activation in the feature layer is passed to the visual category layer, L , along the weighted connections.

$$v_y = v_y + \tau_f \left(\sum f_x \cdot w_{xy} \right) \quad (2)$$

Here, v_y refers to category unit y , f_x refers to feature unit x , w_{xy} refer to the weight connecting feature unit x to category unit y . τ_f is a temperature parameter controlling how fast activation accumulates at the category layer. Next, inhibition at the category layer is implemented using (3), where m is the number of category units, and ω represents the degree of inhibition.

$$v_y = v_y^\omega / \sum v_y^\omega \quad (3)$$

Next, activation is passed to the lexical layer, L , from the visual category layer and the three banks of the auditory layer.

$$l_y = l_y + \tau_f \left(\sum_{x=1 \dots n} a_{1x} \cdot z_{1xy} + \sum_{x=1 \dots n} a_{2x} \cdot z_{2xy} + \sum_{x=1 \dots n} a_{3x} \cdot z_{3xy} + \sum_{x=1 \dots n} v_x \cdot u_{xy} \right) \quad (4)$$

Here, L_y refers to lexical unit y , a_{1x} , a_{2x} , and a_{3x} refers to auditory unit x in the first, the second and the third auditory bank respectively, v_x refers to visual category unit x , z_{1xy} , z_{2xy} , and z_{3xy} refers to the weights connecting auditory unit x in each of the banks

to the lexical unit y , and u_{xy} refers to the connections from the unit x in the visual category layer to the lexical layer. Next, inhibition at the lexical layer is implemented using the formulas (3) but with different parameters (p is the number of lexical units, and ψ represents the degree of inhibition).

Then feedback spreads back down from lexical units to the visual category layer and the auditory layer.

$$v_x = v_x + \tau_b \cdot v_x \sum l_y \cdot u_{xy} \quad (5)$$

Here inhibition ($\omega=2$) is applied at the visual category layer (using formulas 3), and they are normalized. The same equation is used to normalize the auditory layer (with a small amount of inhibition, $\iota=1.05$).

Next, feedback spreads down from category units to the feature layer using (5), where m is the number of category units, and w_{xy} refer to the weight connecting feature unit x to category unit y . A small amount of inhibition ($\iota=1.05$) is the applied at the feature layer using (3), and they are normalized.

Finally, at the end of the iteration, a small amount of learning occurs at all three sets of weights using (6),

$$\Delta w_{xy} = \eta \cdot f_x v_y (1 - w_{xy}) - \delta (1 - f_x) \cdot v_y \cdot w_{xy} - \delta f_x (1 - v_y) w_{xy} \quad (6)$$

Here F_x represents feature unit x , v_y represents category unit y , and w_{xy} represents the connection between these two units. Δw_{xy} represents the amount of weight change at the end of a processing cycle. η represents the learning rate (a very small amount of learning will occur, $\eta=.00025$), and δ represents the decay rate ($\delta=.5$).

REFERENCES

- Abbs, B, Gupta, P, & Khetarpal, N. (2008). Is overt repetition critical to word learning? An examination of the role of overt repetition in word learning with and without semantics. *App Psycholinguist*, 29, 627-667.
- Apfelbaum, K. S. & McMurray, B. (2011). Using variability to guide dimensional weighting: associative mechanisms in early word learning. *Cognitive Science*, 35, 1105-1138.
- Apfelbaum, K. S. & McMurray, B. (2012). When will you learn? Interactions of real-time lexical processing and word learning. Poster presented at the 53rd Annual Meeting of the Psychonomic Society, Minneapolis, MN.
- Apfelbaum, K., Blumstein, S.E., & McMurray, B. (2011). Semantic priming is affected by real-time phonological competition: Evidence for continuous cascading systems. *Psychonomic Bulletin & Review*, 18, 141-149.
- Allopenna, P. D, Magnuson, J.S. & Tanenhaus, M.K. (1998). Tracking the time course of spoken word recognition: evidence for continuous mapping models. *Journal of Memory and Language*, 38, 419-439.
- Anderson, M.C. (2003). Rethinking interference theory: Executive control and the mechanisms of forgetting. *Journal of Memory and Language*, 49, 415-445.
- Anderson, M.C., Bjork, R.A., & Bjork, E.L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 20, 1063-1087.
- Anderson, M.C., Bjork, E.L., & Bjork, R.A. (2000). Retrieval-induced forgetting: Evidence for a recall-specific mechanism. *Psychonomic Bulletin and Review*, 7, 522-530.
- Anderson, M.C. & Green, C. (2001). Suppressing unwanted memories by executive control. *Nature*, 410, 131-134.
- Anderson, M.C., Green, C., & McCulloch, K.C. (2000). Similarity and inhibition in long-term memory: Evidence for a two-factor model. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 26, 1141-1159.
- Anderson, M.C., & McCulloch, K.C. (1999). Integration as a general boundary condition on retrieval-induced forgetting. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25, 608-629.
- Anderson, M.C., & Spellman, B.A. (1995). On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review*, 102, 68-100.
- Baddeley, A. D., Gathercole, S. E. & Papagno, C. 1998 The phonological loop as a language learning device. *Psychol. Rev.*, 105, 158-173.

- Bäuml, K.H. (1998). Strong items get suppressed, weak items do not: The role of item strength in output interference. *Psychonomic Bulletin & Review*, 5, 459–463.
- Bäuml, K.H. (2002). Semantic generation can cause episodic forgetting. *Psychological Science*, 13, 357–361.
- Bloom, P., & Markson, L. (1998). Intention and analogy in children's naming of pictorial representations. *Psychological Science*, 9(3), 200–204.
- Booth, A. E., & Waxman, S. R. (2002). Word learning is 'smart': Evidence that conceptual information affects preschoolers' extension of novel words. *Cognition*, 84, B11–B22.
- Capone, N., & McGregor, K. K. (2005). The effect of semantic representation on toddler's word retrieval. *Journal of Speech, Language, and Hearing Research*, 48, 1468–1480.
- Chen, Q., & Mirman, D. (2012). Competition and cooperation among similar representations: Toward a unified account of facilitative and inhibitory effects of lexical neighbors. *Psychological Review*, 119(2), 417–430.
- Kloos, H., & Sloutsky, V. M. (2008). What's behind different kinds of kinds: Effects of statistical density on learning and representation of categories. *Journal of Experimental Psychology: General*, 137, 52–72.
- Craik, F.I.M., & Lockhart, R.S. (1972). "Levels of processing: A framework for memory research". *Journal of Verbal Learning & Verbal Behavior*, 11(6), 671–84.
- Creel, S. C., & Dahan, D. (2010). The influence of the temporal structure of spoken words on paired-associate learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36, 110–122.
- Dahan, D., Magnuson, J. S., Tanenhaus, M. K., & Hogan, E. M. (2001). Subcategorical mismatches and the time course of lexical access: Evidence for lexical competition. *Language and Cognitive Processes*, 16, 507–534.
- Damian, M. F., Vigliocco, G., & Levelt, W. J. (2001). Effects of semantic context in the naming of pictures and words. *Cognition*, 81(3), B77–86.
- Davis, M.H., & Gaskell, M.G. (2009) A complementary systems account of word learning: neural and behavioural evidence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1536), 3773–3800.
- Dell, G.S. (1986). A Spreading Activation Theory of Retrieval in Sentence Production, *Psychological Review*, 93, 283–321.
- Dumay, N. & Gaskell, M. G. (2007). Sleep-associated changes in the mental representation of spoken words. *Psychological Science*, 18, 35–39.
- Dumay, N., & Gaskell, M. G. (2012). Overnight lexical consolidation revealed by speech segmentation. *Cognition*, 123, 119–132.

- Fallon, A. B., Groves, K., & Tehan, G. (1999). Phonological similarity and trace degradation in the serial recall task: When CAT helps RAT, but not MAN. *International Journal of Psychology*, 34, 301-307.
- Fennell, C. T. (2012). Object familiarity enhances infants' use of phonetic detail in novel words. *Infancy*, 17, 339-353.
- Freeman, J.H. (2010). Developmental neurobiology of cerebellar learning. In Blumberg, M.S., Freeman, J.H., & Robinson, S.R. (Eds.). *Oxford Handbook of Developmental Behavioral Neuroscience*. Oxford University Press. Gaskell, M. G., & Dumay, N. (2003). Lexical competition and the acquisition of novel words, *Cognition*, 89, 105-132.
- Gathercole, S. E. (2006). Nonword repetition and word learning: The nature of the relationship. *Applied Psycholinguistics*, 27, 513-543.
- Gathercole, S. E., Service, E., Hitch, G. J., Adams, A.-M. & Martin, A. J. (1999) Phonological short-term memory and vocabulary development: further evidence on the nature of the relationship. *Appl. Cogn. Psychol.*, 13, 65-77.
- Goldstone, R. L. (1994). The role of similarity in categorization: Providing a groundwork. *Cognition*, 52, 125-157.
- Goldstone, R. L. (2000). Unitization during Category Learning. *Journal of Experimental Psychology: Human Perception and Performance*, 26, 86-112.
- Gordon, J.K., & Dell, G.S. (2003). Learning to divide the labor: An account of deficits in light and heavy verb production. *Cognitive Science*, 27,1-40.
- Graf Estes, K., Edwards, J., & Saffran, J.R. (2011). Phonotactic constraints on infant word learning. *Infancy*, 16, 180-197.
- Gupta, P. (1996). Word learning and verbal short-term memory: A computational account. In G. W. Cottrell (Ed.), *Proceedings of the Eighteenth Annual Meeting of the Cognitive Science Society*, pages 189-194. Hillsdale, NJ, Lawrence Erlbaum.
- Gupta, P. (2005). Primacy and Recency in Nonword Repetition. *Memory*, 13, 318-324
- Gupta, P. (2008). The role of computational models in investigating typical and pathological behaviors. *Seminars in Speech and Language*, 29(3), 211-225.
- Gupta, P.(2009). A computational model of nonword repetition, immediate serial recall, and nonword learning. In Thorn, A & Page, M (Ed.) *Interactions Between Short-Term and Long-Term Memory in the Verbal Domain*, pages: 108-135. New York: Psychology Press.
- Gupta, P., Lipinski, J., and Aktunc, E. (2005). Re-Examining the Phonological Similarity Effect in Immediate Serial Recall: The Roles of Type of Similarity, Category Cueing, and Item Recall. *Memory and Cognition*, 33 (6), 1001-1016.

- Gupta, P., Lipinski, J., Abbs, B. and Lin, P-H. (2005). Serial Position Effects in Nonword Repetition. *Journal of Memory and Language*, 53, 141-162.
- Gupta, P., Newman, R., & Samuelson, L. (2010) Wat the Fowk? Dense phonological neighborhoods help both children and connectionist models learn new words.
- Gupta, P. & Tisdale, J. (2009). Does phonological short-term memory causally determine vocabulary learning? Toward a computational resolution of the debate. *J. Mem. Lang.*, 61, 481-502.
- Gupta, P. & MacWhinney, B. (1997). Vocabulary acquisition and verbal short-term memory: Computational and neural bases. *Brain and Language*, 59, 267-333.
- Hebb, D.O. (1949). *The organization of behavior*. New York: Wiley & Sons.
- Hodges, J. R., Patterson, K., Oxbury, S., & Funnell, E. (1992). Semantic dementia: Progressive fluent aphasia with temporal lobe atrophy. *Brain*, 115, 1783-1806.
- Hoover, J. R., Storkel, H. L., & Hogan, T. P. (2010). A cross-sectional comparison of the effects of phonotactic probability and neighborhood density on word learning by preschool children. *Journal of Memory and Language*, 63, 100-116.
- Horst, J.S., McMurray B., & Samuelson, L.K. (2006). Online processing is essential for learning: Understanding fast mapping and word learning in a dynamic connectionist architecture. In R Sun (Ed.), *Proceedings of the Twenty-Eighth Annual Conference of the Cognitive Science Society* (pp. 339-344) LEA.
- Horst, J. S. & Samuelson, L. K. (2008). Fast Mapping But Poor Retention in 24-Month-Old Infants. *Infancy*, 13(2), 128-157.
- Howard, D., Nickels, L., Coltheart, M., & Cole-Virtue, J. (2006). Cumulative semantic inhibition in picture naming: Experimental and computational studies. *Cognition*, 100(3), 464-482.
- Huetig, F. & Altmann, G.T.M. (2005). Word meaning and the control of eye fixation: semantic competitor effects and the visual world paradigm. *Cognition*, 96(1), 23-32.
- Karpicke, J.D. and Roediger, H.L., III. (2008) .The critical importance of retrieval for learning. *Science*, 15, 966-968.
- Kloos, H., & Sloutsky, V. M. (2008). What's behind different kinds of kinds: Effects of statistical density on learning and representation of categories. *Journal of Experimental Psychology: General*, 137, 52-72.
- Kroll, J. F., & Stewart, E. (1994). Category interference in translation and picture naming: Evidence for asymmetric connections between bilingual memory representations. *Journal of Memory and Language*, 33, 149-174.
- Lipinski, J., and Gupta, P. (2005). Does Neighborhood Density Influence the Repetition of Nonwords? Separating the Effects of Density and Duration. *Journal of Memory and Language*, 52, 171-192.

- Luce, P. A., & Large, N. (2001). Phonotactics, neighborhood density, and entropy in spoken word recognition. *Language and Cognitive Processes*, 16, 565-581.
- Luce, P. A., & Pisoni, D. B. (1998). Recognizing spoken words: The neighborhood activation model. *Ear and Hearing*, 19(1), 1-36.
- Magnuson, J.S., Tanenhaus, M.K., Aslin, R.N. & Dahan, D. (2003). The time course of spoken word learning and recognition: Studies with artificial lexicons. *Journal of Experimental Psychology: General*, 132, 202-227.
- Magnuson, J. S., Dixon, J. A., Tanenhaus, M. K., & Aslin, R. N. (2007). The dynamics of lexical competition during spoken word recognition. *Cognitive Science*, 31, 133-156.
- Marslen-Wilson, W. (1987). Functional parallelism in spoken word recognition. *Cognition*, 25, 71-102.
- McClelland, J., & Elman, J. (1986). The TRACE model of speech perception. *Cognitive Psychology*, 18(1), 1-86.
- McMurray, B., Horst, J., and Samuelson (2012). Using your lexicon at two timescales: Investigating the interplay of word learning and recognition. *Psychological Review*, 119, 831-877.
- McMurray, B., Horst, J., Toscano, J., and Samuelson (2009). Towards an integration of connectionist learning and dynamical systems processing: case studies in speech and lexical development. In Spencer, J., and Thoms, M., and McClelland, J. (Eds). *Toward a Unified Theory of Development: Connectionism and Dynamic Systems Theory Re-Considered*. London: Oxford University.
- McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior Research Methods*, 37, 547-559.
- Merriman, W. E., & Bowman, L. L. (1989). The mutual exclusivity bias in children's word learning. *Monographs of the Society for Research in Child Development*, 54 (3-4, Serial No.220).
- Meyer, D.E., & Schvaneveldt, R.W. (1971). Facilitation in Recognizing Pairs of Words: Evidence of a Dependence Between Retrieval Operations. *Journal of Experimental Psychology*, 90 (2): 227-234.
- Mirman, D. (2011). Effects of near and distant semantic neighbors on word production. *Cognitive, Affective, and Behavioral Neuroscience*, 11(1), 32-43.
- Mirman, D. and Magnuson, J.S. (2008). Attractor dynamics and semantic neighborhood density: Processing is slowed by near neighbors and speeded by distant neighbors. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 34(1), 65-79.
- Mirman, D. and Magnuson, J.S. (2009b). Dynamics of activation of semantically similar concepts during spoken word recognition. *Memory & Cognition*, 37(7), 1026-1039.
- Moores, E., Laiti, L., & Chelazzi, L. (2003). Associative knowledge controls deployment of visual selective attention. *Nature Neuroscience*, 6(2), 182-189.

- Morris, C. D., Bransford, J. D., & Franks, J. J. (1977). Levels of processing versus transfer appropriate processing. *Journal of Verbal Learning and Verbal Behavior*, 16, 519-533.
- Moscovitch, M. & Craik, F.I. (1976). Depth of processing, retrieval cues, and uniqueness of encoding as factors in recall. *Journal of Verbal Learning & Verbal Behavior*, 15(4), 447-458.
- Nosofsky, R.M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39-57.
- Nosofsky, R.M. (1992). Exemplar-based approach to relating categorization, identification, and recognition. In F.G. Ashby (Ed.), *Multidimensional models of perception and cognition*. Hillsdale, NJ: Erlbaum.
- Oppenheim, G. M., Dell, G. S., & Schwartz, M. F. (2010). The dark side of incremental learning: A model of cumulative semantic interference during lexical access in speech production. *Cognition*, 114(2), 227-252.
- O'Reilly, R.C. & Munakata, Y. (2000). *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*. Cambridge, MA: MIT Press.
- Palmeri, T. J. (1997). Exemplar similarity and the development of automaticity. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 23, 324-354.
- Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin*, 116, 220-244.
- Pantelis, P.C., van Vugt, M.K., Sekuler, R., Wilson, H.R., & Kahana, M.J. (2008). Why are some people's names easier to learn than others? The effects of similarity on memory for face-name associations. *Memory & Cognition*, 36, 1182-1195.
- Peterson, R.R and Savoy, P. (1998). Lexical selection and phonological encoding during language production: Evidence for cascaded processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(3), 539-557.
- Plaut, D. C. (1997). Structure and function in the lexical system: Insights from distributed models of word reading and lexical decision. *Language and Cognitive Processes*, 12, 767-808.
- Plaut, D.C., and Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, 10, 377-500.
- Posner, M.I., & Keele, S.W. (1968). On the genesis of abstract idea. *Journal of Experimental Psychology*, 77, 353-363.
- Rapp, B. & Goldrick, M. (2000). Discreteness and interactivity in spoken word production. *Psychological Review*, 107, 460-499.
- Regier, T. (2005). The emergence of words: Attentional learning in form and meaning. *Cognitive Science*, 29, 819-865.

- Reed, S.K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, 3, 382-407.
- Rosch, E., & Mervis, C.B. (1975). Family resemblance: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573-605.
- Rost, G., & MacMurray, B. (2009). Speaker variability augments phonological processing in early word learning. *Developmental Science*, 12, 339-349.
- Rost, G., & MacMurray, B. (2010). Finding the Signal by Adding Noise: The Role of Noncontrastive Phonetic Variability in Early Word Learning. *Infancy*, 15, 608-635.
- Samuelson, L.K., Schutte, A.R. & Horst, J.S. (2009). The dynamic nature of knowledge: Insights from a dynamic field model of children's novel noun generalizations. *Cognition*, 110, 322-345.
- Schnur, T. T., Schwartz, M. F., Brecher, A., & Hodgson, C. (2006). Semantic interference during blocked-cyclic naming: Evidence from aphasia. *Journal of Memory and Language*, 54(2), 199-227.
- Schöner, G. (2008). Dynamical systems approaches to cognition. In J.P. Spencer, M.S. Thomas, & J.L. McClelland (Eds.) *Toward a Unified Theory of Development: Connectionism and Dynamic Systems Theory Re-Considered*. New York: Oxford University Press.
- Service, E., & Maury, S. (2003). All parts of an item are not equal: Effects of phonological redundancy on immediate recall. *Memory & Cognition*, 31, 273-284.
- Sevold, C. A., & Dell, G. S. (1994). The sequential cuing effect in speech production. *Cognition*, 53, 91-127.
- Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, 106 (3), 1558-1568.
- Spivey, M. J. (2007). *The continuity of mind*. New York: Oxford University Press.
- Spivey, M., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors: Thinking with your hands. *Proceedings of the National Academy of Sciences*, 102, 10393-10398.
- Stager, C. L., & Werker, J. F. (1997). Infants listen for more phonetic detail in speech perception than in word learning tasks. *Nature*, 388(6640), 381-382.
- Storkel, H.L. (2001). Learning new words: Phonotactic probability in language development. *Journal of Speech, Language, and Hearing Research*, 44, 1321-1337.
- Storkel, H.L. & Adlof, S.M. (2009b). The effect of semantic set size on word learning by preschool children. *Journal of Speech, Language, and Hearing Research*, 52, 306-320.
- Storkel, H.L., Armbruster, J., & Hogan, T.P. (2006). Differentiating phonotactic probability and neighborhood density in adult word learning. *Journal of Speech, Language, and Hearing Research*, 49, 1175-1192.

- Storkel, H.L. & Rogers, M.A. (2000). The effect of probabilistic phonotactics on lexical acquisition. *Clinical Linguistics & Phonetics*, 14, 407-425.
- Storm, B. C., Bjork, E. L., & Bjork, R. A. (2012). On the durability of retrieval-induced forgetting. *Journal of Cognitive Psychology*, 24, 617-629.
- Swingle, D. & Aslin, R. N. (2007). Lexical competition in young children's word learning. *Cognitive Psychology*, 54, 99-132.
- Theissen, E. D. (2007). The effect of distributional information on children's use of phonemic contrasts. *Journal of Memory and Language*, 56, 16-34.
- Thelen, E., & Smith, L.B. (1994). *A dynamic systems approach to the development of perception and action*. Cambridge, MA: MIT Press.
- Tomasello, M. (2000). The social-pragmatic theory of word learning. *Pragmatics*, 10, 401-14.
- Tranel, D. (2006). *Impaired naming of unique landmarks is associated with left temporal polar damage*. *Neuropsychology*, 20(1), 1-10.
- Vanderplas, J. M., & Garvin, E. A. (1959a). The association value of random shapes. *Journal of Experimental Psychology*, 57, 147-154.
- Veling, H., & van Knippenberg, A. (2004). Remembering can cause inhibition: Retrieval-induced inhibition as cue independent process. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 315-318.
- Vigliocco, G., Vinson, D. P., Damian, M. F., & Levelt, W. J. M. (2002). Semantic distance effects on object and action naming. *Cognition*, 85(3), B61-B69.
- Vitevitch, M S. (2002a). The influence of phonological similarity neighborhoods on speech production. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 28, 735-747.
- Vitevitch, M.S. (2002b). Influence of onset-density on spoken word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 28, 270-278.
- Vitevitch, M.S. (2007). The spread of the phonological neighborhood influences spoken word recognition. *Memory & Cognition*, 35, 166-175.
- Vitevitch, M. S., & Luce, P. A. (1998). When words compete: Levels of processing in spoken word perception. *Psychological Science*, 9, 325-329.
- Vitevitch, M. S., & Luce, P. A. (1999). Probabilistic phonotactics and spoken word recognition. *Journal of Memory and Language*, 40, 374-408.
- Vitevitch, M.S., Armbruster, J. & Chu, S. (2004). Sub-lexical and lexical representations in speech production: Effects of phonotactic probability and onset density. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 30, 514-529.

- Vitevitch, M.S. & Sommers, M. (2003). The facilitative influence of phonological similarity and neighborhood frequency in speech production. *Memory & Cognition*, 31, 491-504.
- Yee, E., & Sedivy, J. (2001). Using eye movements to track the spread of semantic activation during spoken word recognition. Paper presented to the 13th annual CUNY sentence processing conference, Philadelphia.
- Yoshida, K., Fennell, C., Swingley, D., & Werker, J.F. (2009). 14-month-olds learn similar-sounding words. *Developmental Science*, 12, 412-418.
- Yu, C., & Smith, L. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science*, 18 (5), 414-420.
- Zwitserslood, P. (1989). The locus of the effects of sentential-semantic content in spoken-word processing. *Cognition*, 32(1), 25-64.