

UC Berkeley

Dissertations, Department of Linguistics

Title

Of Sound, Mind, and Body: Neural Explanations for Non-categorical Phonology

Permalink

<https://escholarship.org/uc/item/1kc7d5vh>

Author

Bergen, Benjamin

Publication Date

2001

Of sound, mind, and body:
Neural explanations for non-categorical phonology

by

Benjamin Koppel Bergen

B.A. (University of California, Berkeley) 1996

M.A. (University of California, Berkeley) 1997

A dissertation submitted in partial satisfaction of the requirements for the degree of

Doctor of Philosophy

in

Linguistics

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:

Professor George P. Lakoff, Chair

Professor Sharon Inkelas

Professor Jerome A. Feldman

Fall 2001

The dissertation of Benjamin Koppel Bergen is approved:

Chair

Date

Date

Date

University of California, Berkeley

Fall 2001

Of sound, mind, and body:
Neural explanations for non-categorical phonology

© 2001

by

Benjamin Koppel Bergen

Abstract

Of sound, mind, and body:

Neural explanations for non-categorical phonology

by

Benjamin Koppel Bergen

Doctor of Philosophy in Linguistics

University of California, Berkeley

Professor George P. Lakoff, Chair

Traditional linguistic models are categorical. Recently, though, a number of researchers have begun to study non-categorical human linguistic knowledge (e.g. Bender 2000, Pierrehumbert 2000, Frisch 2001). This new empirical focus has posed significant difficulties for categorical models, which cannot account for many non-categorical phenomena. Rather than trying to fit the non-categorical complexities of language into categorical models, a number of researchers have begun to treat non-categoricity in probabilistic terms (Jurafsky 1996, Abney 1996, Bod 1998). This dissertation demonstrates experimentally that language users have knowledge of non-categorical correlations between phonology and other grammatical, semantic, and social knowledge and that they apply this knowledge to the task of language perception. The thesis also proposes neural explanations for the behavior exhibited in the experiments, and develops neurally plausible, probabilistic computational models to this end.

This first half of this dissertation presents new evidence of the non-categoricity of human linguistic knowledge through two case studies. The first addresses the relation

between sound and meaning, though an experimental investigation of the psychological reality of English phonaesthemes, and shows that these non-categorical sub-morphemic sound-meaning pairings are psychologically real. A second, larger study addresses the multiple factors that non-categorically affect a particular morpho-phonological process in French, called *liaison*. These two studies provide evidence that language users access non-categorical relations between phonological patterns and their phonological, morphological, syntactic, semantic, and social correlates. An additional result of the liaison study is the finding that language users exhibit unconscious knowledge of non-categorical interactions between factors that influence this morpho-phonological process.

While there are general neural explanations for the ability to learn and represent the knowledge suggested by these studies, a formal model can only be produced in a computational architecture. Therefore, in the dissertation's second half, I develop a computational model of non-categorical, cross-modal knowledge using a probabilistic architecture used in Artificial Intelligence research, known as Belief Networks (Pearl 1988). In addition to capturing the generalizations about non-categorical knowledge evidenced by the two case studies, Belief Networks are neurally plausible, making them a sound architecture for a bridging model between neural structure and cognitive and linguistic behavior.

To my family.

Novels get read,

Poems do, too.

Please use my thesis

As an impromptu clobbering device.

- B.K.B

TABLE OF CONTENTS

Chapter 1. Phonology in a mind field	1
1. Overview	1
2. Probability and perception	4
3. A bridge between brain and behavior	7
Chapter 2. Probability and productivity	10
1. Introduction	10
2. Grammatically correlated phonology	11
3. Sound-meaning correlations	24
4. An experimental study of phonaesthemes in language processing	33
5. Future directions	42
Chapter 3. Social factors and interactions in French liaison	45
1. Introduction	45
2. Variability in the language of the community	46
3. Interactions between factors	56
4. French liaison and the liaison corpus	61
5. A test of autonomous factors	92
6. A test of interactions between factors	121
7. Final note	131
Chapter 4. Individual perception of variability	133
1. Introduction	133
2. Experiment 1: Autonomous factors	134
3. Experiment 2: Interactions between factors	145
4. Discussion	156
Chapter 5. Probabilistic computational models of non-categoriality	157
1. Introduction	157
2. Belief Networks	159
3. Liaison in a Belief Network model	164
4. Phonaesthemes in a Belief Network model	175
5. Properties of the models	284
Chapter 6. Neural bases for unruly phonology	191
1. Introduction	191
2. Levels of representation	192
3. Neural explanations for non-categorical phonology	196
4. A neural rendition of Belief Networks	211
Chapter 7. Conclusion	217
1. Debates	217
2. Future directions	218
3. Summary	220
References	222
Appendix: Phonaestheme examples	240

Acknowledgments

Writing a dissertation is always a group process. A group of people who are not the author sit around and hock up theoretical goo until they agree to a sufficient degree for a dissertation topic to be born. Then it's a matter of filling in the goo-blanks. As the street poets say, there's no "I" in "dissertation". That's the problem with street poets - they usually can't spell.

I would like to thank many people for tolerating me during the dissertation-writing process. I am normally not a very nice person, and while working on this document, I might have gone a little overboard.

For example, I'd like to apologize to Steve for eating his dog. Steve, I really wasn't thinking about the consequences of my actions, or about childhood traumas you might have endured.

Sorry to Madelaine and Richard for registering a star in your name, and calling it "Fungus Toes Eleventy Million". I was under a lot of stress then, and I temporarily forgot how sub-cutaneous foot mold had almost destroyed your relationship.

To Chisato, I have only two words of remorse: bones knit.

Additionally, I should apologize to Julie, Ashlee, and Nancy for some brash statements I might have made. I don't ACTUALLY think that as women you should be required to write dissertations on cooking terminology.

To Jocelyn and Rick, I convey my most heartfelt congratulations at the newest addition to their family, a delicious baby girl. I promise to keep this one away from my sausage grinder.

Finally, sorry to my parents for the whole “I’m sending you into a rest home when you turn sixty” thing I wrote on the Hanukkah cards. That wasn’t very considerate. Especially since you lose your sense of humor when you get old.

Others, I would like to thank for not ever being born. Like that horrific goat-faced mutant boy not living in my basement. Or the hirsute but polite ninja not occupying space on my bookshelf.

No dissertation acknowledgements would be complete without some glib reference to the thesis gnomes who come out at night and erase connecting sentences and key modifiers in dissertation drafts. Those gnomes really piss me off. But they’re not bad with lentils.

This dissertation hereby puts the “fun” back in “phonology”. It also puts the “inguis” back in “linguistics”.

Seriously, though, thanks to everyone who read this document, and also thanks to my dissertation committee.

Chapter 1. Phonology in a mind field

Outline

1. Overview
2. Probability and perception
3. A bridge between brain and behavior

If you are out to describe the truth, leave elegance to the tailor.

Albert Einstein

1. Overview

How the brain works matters to language. The brain matters to language in an obvious way and in a deep way. Obviously, the brain happens to be the computing device that makes language and the rest of cognition happen - if you lose part of your brain, chances are you'll also lose part of your mind. Less obvious is whether the details of linguistic and other knowledge depend on computational properties of the human brain. This thesis presents evidence that language knowledge and behavior is shaped in a deep way by how the human brain works. The thesis targets one particular aspect of human language that is explained by neural functioning. That aspect is non-categorical knowledge.

Traditional theories of language view linguistic knowledge as inherently categorical. That is, it is made up of rules that are stated in absolute terms, and which are applied deterministically when the appropriate context arises. Recently, though, a number of researchers have begun to investigate the degree to which linguistic knowledge is in fact not categorical (e.g. Bender 2000, Pierrehumbert 2000 and In Press, Frisch 2001).

Language is non-categorical in several ways. First, pieces of linguistic knowledge such as phonological rules are sometimes not categorical; that is, they do not apply across the board, deterministically, but rather have some probability associated with their application (Pierrehumbert In Press). The second way in which language is non-categorical is in the interactions between pieces of linguistic knowledge. Two phonological rules, for example, might not stand in an absolute precedence relation. Rather, one might be assigned some probability of taking precedence over the other (Hayes 2000). Finally, the objects of linguistic knowledge, units like morphemes, are often not classes that can be defined categorically - there for example may be soft edges to phoneme (Jaeger and Ohala 1984) or morpheme categories (Bergen 2000a).

This new empirical focus has posed significant difficulties for categorical models, which cannot account for many non-categorical phenomena. Rather than trying to fit the non-categorical complexities of language into categorical models, a number of researchers have begun to deal with non-categoricity using probabilistic models (Jurafsky 1996, Abney 1996, Dell et al. 1997, Narayanan and Jurafsky 1998, Bod 1998).

This dissertation focuses on non-categorical correlations between phonology and other grammatical, semantic, and social knowledge. Its major empirical contribution is to demonstrate experimentally that language users pick up on these non-categorical correlations and apply them to the task of language perception. The thesis also articulates specific neural explanations for the behavior exhibited in the experiments, and develops neurally plausible computational models to this end.

This first half of this dissertation presents new evidence of the non-categoricity of cross-domain human linguistic knowledge through two case studies. The first addresses the relation between sound and meaning, though an experimental study of the psychological

reality of English *phonaesthemes* (Firth 1930). A second, larger study addresses the multiple factors that non-categorically affect a particular morpho-phonological process in French, known as *liaison* (Tranel 1981, Selkirk 1984). These two studies provide evidence that language users access non-categorical relations between phonological patterns and their phonological, morphological, syntactic, semantic, and social correlates. An additional result of the liaison study is the finding that language users exhibit unconscious knowledge of non-categorical interactions between factors that influence this morpho-phonological process. A result of the phonaestheme experiment is that morphemes as defined strictly are not the only sub-lexical material that can probabilistically pair form and meaning.

There are general neural explanations for the ability to learn and represent the knowledge suggested by these studies. But a formal model linking linguistic knowledge and its neural explanation can only tractably be produced in some computational architecture. Therefore, in the dissertation's second half, I develop a computational model of non-categorical, cross-modal knowledge using a probabilistic architecture used in Artificial Intelligence research, known as Belief Networks (Pearl 1988, Jensen 1996). Belief Networks are able to capture the generalizations about non-categorical knowledge evidenced by the two case studies in this thesis. But they are also neurally plausible, making them a sound basis for a bridging model between neural structure and cognitive and linguistic behavior. Belief Network models are sufficiently flexible to also deal with canonical, categorical linguistic generalizations.

2. Probability and perception

The first half of this thesis asks to what extent language users pick up on non-categorical pairings between phonological patterns and other linguistic and extralinguistic patterns.

When listening to language, hearers are confronted with all sorts of variability in the phonology of the input. It was generally held that variation was inherently uncorrelated, or free (e.g. Hubbel 1950), until statistical sociolinguistic methods began to show systematic social correlates of phonological variation (e.g. Labov 1966). The last thirty years have taught us that most variability correlates with other facts that hearers may have direct or indirect access to. These factors can be social or linguistic. More recently, as we will see below, hearers have been shown through psycholinguistic experimentation to possess knowledge of these correlations.

There is strong evidence that phonological knowledge displays non-categorical correlations with morphological or syntactic knowledge. For example, English verbs tend to have front vowels, as in *bleed* and *fill*, while nouns tend to have back vowels, as in *blood* and *foal* (Sereno 1994). Another case is the tendency for English disyllabic verbs to display word-final stress, while nouns more frequently have word-initial stress. For example, consider the contrasting pronunciations of *convert*, *convict*, and *record* (Sherman 1975). These asymmetries are slight statistical tendencies. And yet, a large number of studies have shown that language users make unconscious use of the asymmetries during language perception and production. Hearers respond more quickly to words whose phonological features best match their morphosyntactic category (Sereno 1994). Speakers are more likely to produce novel words of a given morphosyntactic category if those words have the phonological characteristics predominantly shared by words of that class (Kelly and Bock 1988). These studies suggest

that human knowledge about sound is not independent of syntactic or morphological knowledge. They also show that while some relations between phonology and morphosyntax may be categorical, language users are also able to pick up on non-categorical ones.

The goal of Chapter 2 is to first survey the evidence for human knowledge of cross-modal non-categorical linguistic correlations, and then to investigate how phonological knowledge is also non-categorically related to semantic knowledge. A short but nevertheless compelling line of research has unearthed a number of non-categorical relations between sound and meaning. Examples include the relation between the phonology of first names and biological sex (Cassidy et al. 1999), and between semantic complexity and word length (Kelly et al. 1990). This research has also shown that language users once again pick up on these correlations and incorporate them into their language perception and production systems.

In Chapter 2, I augment the range of known probabilistic sound-meaning pairings through an experimental study of the role of *phonaesthemes* in human language processing. Phonaesthemes are sub-morphemic sound-meaning pairings, exemplified in words such as *glow*, *glisten*, and *gleam*. These words, and a host of others in English, share the complex onset *gl-* and meanings related to ‘VISION’ and ‘LIGHT’. Another phonaestheme is the rime of *slap*, *tap*, *rap*, and *snap*, words which share a common semantics of ‘TWO SURFACES COMING TOGETHER, CAUSING A BURST’. In the perception experiment described in Chapter 2, phonaesthemes exhibit facilitory priming effects. These effects are different from semantic and phonological priming, and as such indicate that individuals pick up on phonaesthemes’ statistical predominance and make use of it during language processing. This result supplies further evidence that non-categorical correlations between form and meaning are psychologically real, even when they have no grammatical status.

Non-categorical effects are not limited to the language of the individual, however. They also surface in the influence of social factors on the selection of sounds to express meaning. The case of French liaison consonants, introduced in Chapter 3, is a good example of such sociolinguistic factors, which have been well studied. Liaison consonants are word-final consonants like the final /z/ of *les* ‘the’, which is pronounced in *les heures* /lezoer/ ‘the hours’, but not in *les minutes* /leminyt/ ‘the minutes’. The probability that liaison consonants will be produced depends on a broad range of factors. Among these are phonological ones, like the character of the following segment; syntactic ones, like the relation between the liaison word and the following word, semantic ones, like whether the liaison consonant bears meaning, and sociolinguistic ones, like the age of the speaker (Ashby 1981).

What distinguishes the study of liaison in Chapter 3 from other work on phonological variability is the attention it pays to *interactions* between factors. Some factors influencing the production of French optional final consonants are *autonomous*. By autonomous, I mean that these factors each contribute probabilistically to the realization of the final consonants, regardless of the values of other factors. In other words, the contribution of one factor, like gender, can be calculated without reference to the contribution of other factors, like phonological environment. This autonomous type of effect has been well-studied in the sociolinguistic literature. Other factors influencing the pronunciation of liaison consonants, though, interact - they cannot be understood without also taking into account a number of other factors. For example, while older speakers are more likely to produce liaison consonants in general, when the liaison consonant appears in an adverb, like the final consonant of *trop* ‘too much’, this trend is less strong reversed. A number of such interactions among factors emerge through the statistical analysis in Chapter 3 of a large corpus of spoken French.

Of course, the demonstration of interacting factors in a multi-speaker corpus in no way implies that factors interact in the language of the individual. Chapter 4 reports on a perception experiment that tests whether individual speakers make use of probabilistically interacting factors. The experiment is structured around tokens taken from the corpus described in Chapter 3. In the experiment, native French speakers are presented with a sequence of two words in French, such as *trop envie* ‘too much want’ with a potential liaison consonant (in the example, the final /p/ of *trop*). These stimuli vary by: (1) whether or not the liaison consonant is pronounced, (2) the age of the speaker, and (3) the grammatical class of the liaison word. The experiment shows that speakers are unconsciously aware of interactions between age and liaison word grammatical class. This is the first experimental evidence that I know of for human knowledge of interacting factors on any linguistic variable.

The first part of this thesis, Chapters 2, 3, and 4, demonstrates that phonological knowledge is closely tied to a wide range of other types of knowledge in a probabilistic manner, and that some of these probabilistic contributions interact. The second part, Chapters 5, and 6, identifies explanations for these phenomena in the functioning of the human brain.

3. A bridge between brain and behavior

Humans know and make use of non-categorical influences on phonology from grammatical, semantic, and social sources, and interactions between these factors. These phenomena can only be captured by a restricted class of modeling architectures. Among these architectures are a number of structured and probabilistic models that are both descriptively adequate and

are also neurally plausible, thus bridging between the use of non-categorical, cross-modal knowledge during language perception and its neural explanation.

Chapter 5 develops a computational architecture for modeling probabilistic, interacting factors between modes of linguistic and other knowledge. It does so at the level of a computational system, running on a digital computer, constrained such that it can be realized at the level of the neurobiology responsible for human linguistic behavior. The computational level machinery is a restricted version of Belief Networks (BNs), which are models of probabilistic causal knowledge in a graphical form. BNs can describe independent and interacting probabilistic contributions. BNs can provide an important basis for modeling probability in phonology, and for providing an interface between aspects of individuals' language and the language of their community.

Chapter 6 develops neural-level explanations for the interacting and independent probabilistic cross-domain knowledge demonstrated by the studies in Chapters 2 and 4. I first survey the connectionist literature and demonstrate that the biological mechanisms thought to be responsible for probabilistic, cross-domain behavioral influences can be explained by how they are learned. Associative learning is explained at the neural level, from a connectionist perspective. I then describe a neural model of the acquisition of interacting probabilistic knowledge. Finally, I complete the loop by showing the restricted sort of BN model described in Chapter 5 to be neurally plausible.

Natural language understanding and speech recognition software already makes use of probability inside modules, and incorporating cross-domain probabilistic mechanisms like the ones described here could improve them. Just as the introduction of probability into physics in the twentieth century has led to monumental advances in that field, according

probability a role in linguistics yields insight into the organization of language and lends power to linguistic models.

In sum, this thesis identifies some deep ways in which the brain matters to language. Achieving this requires four steps. First, the thesis identifies a set of non-categorical phenomena in language, including interaction effects, that can not be described or explained by categorical models and documents their psychological reality. Second, it develops probabilistic computational models of these cognitive and linguistic behaviors in considerable detail. Third, it shows that the behaviors are in fact completely predictable on the basis of general properties of neural systems. Finally, it shows how appropriately constructed computational models can serve as bridges between brain function and cognition, by defining detailed mappings from the models to cognitive behaviors and from the models to the neural structure responsible for the behaviors.

Chapter 2. Probability and productivity

Outline

1. Introduction
2. Grammatically correlated phonology
3. Sound-meaning correlations
4. An experimental study of phonaesthemes in language processing
5. Future directions

Not only does God definitely play dice, but He sometimes confuses us by throwing them where they can't be seen.

Stephen Hawking

1. Introduction

It's well known that correlations exist between phonological patterns on the one hand and morphological and syntax ones on the other. Not all of these correlations are *categorical*, though. Oftentimes, a particular part of speech or grammatical construction will only tend to correlate with a phonological pattern. In this chapter, I will first review existing documentation that these *non-categorical* correlations are used by language users during language processing. I will then move on to show that there are correlations between phonology and meaning that also have a psychological status. English phonaesthemes, which up to the present haven't been shown to be psychologically real, will be shown in a priming experiment to play a part in language processing. These results show that language users extract probabilistic associations from their knowledge of linguistic distributions, whether or not those associations are not productive or grammatical.

2. Grammatically correlated phonology

Phonological generalizations are often correlated with other grammatical factors: *phono-syntactic* and *phono-morphological* generalizations. It's not just that correlations exist between phonological and grammatical aspects of linguistic structures. More importantly, language users make use of these correlations when processing language. Of particular interest are *non-categorical* cross-modal generalizations, which stand in contrast with *categorical ones*.

Categorical grammatical category restrictions on the distribution of phonological elements can be found in a number of languages. In most dialects of English, for example, word-initial [ð] is exclusively restricted to function words, while its voiceless counterpart [θ] occurs only in content words. Thus, the voiced interdental fricative [ð] begins *this* and *that*, while [θ] is the onset of *thick* and *thin*. This distribution is categorical in that the voicing of the initial interdental fricative is deterministically tied to function/content status of the word in question.¹ Categorical generalizations like this one have been long recognized as essential to linguistic models. Chomsky and Halle's (1968) Morpheme Structure Constraints are an implementation of such constraints.

By contrast with categorical ones, *non-categorical* phonosyntactic generalizations have been less well incorporated into linguistic models. These involve a more complex relationship between phonological and other grammatical knowledge. Specifically, a phonological feature or set of features correlates probabilistically, and not deterministically, with a morphosyntactic property. We will be looking here at several classes of non-

¹The function/content distinction isn't actually entirely categorical. When the word *this* is used as a noun in the context of Java programming, it usually takes a word-initial voiced [ð]. And Sharon Inkelas (p.c.) points out that *through* is another, more common exception to this rule, indicating that the generalization probably holds only of word-initial, pre-vocalic interdental fricatives.

categorical phonosyntactic generalizations, some of which serve a morphological function, others of which correlate phonological features with parts of speech, and a final set which embody correlations between syntactic constructions and phonological properties of their constituents.

Morphological function

A small set of English strong past tense verb forms adhere to a shared phonological schema (Bybee and Moder 1983). This particular set of forms, exemplified by the forms in Figure 1, relates a range of present tense forms with a single set of phonological features in the past tense.

The relationship between these present and past tense forms seems to be best analyzed not in terms of static (Bochner 1982, Hayes 1998) or derivational (Rummelhart et al. 1986, MacWhinney and Leinbach 1991) relations. Unlike the regular verb forms in Figure 2, where there is a direct mapping between present and past tense phonology (e.g. /o^w/ in the present corresponds with /u^w/ in the strong past, the present tense forms of this special class in Figure 1 have widely ranging phonology. A number of different present tense vowels, including /I/, /i/, /aj/, and /æ/ are mapped to a common past tense target: /U/.

<i>Present</i>	<i>Past</i>
spin	spun
cling	clung
hang	hung
slink	slunk
stick	stuck
strike	struck
sneak	snuck
dig	dug

Figure 1 - English schematic strong verb morphology

Present	Past	Past Participle	Examples
/o ^w /	/u ^w /	/o ^w n/	blow, flow, know, throw
/aj/	/o ^w /	/I/-/● n/	drive, ride, write, rise
/er/	/or/	/orn/	wear, tear, bear, swear
/ijC/	/ECt/	/ECt/	creep, weep, kneel, feel
/ij/	/E/	/E/	bleed, feed, lead, read

Figure 2 - English regular strong verbal morphology

But the differences between the present tense forms of the strange class in Figure 1 are complemented by a set of phonological similarities between these words. For example, many share a final velar stop or velar nasal. Many have an initial cluster, starting in /s/.

Bybee and Moder (1993) analyze this class as representative of a past tense schema, which

describes the prototypical phonological features of verbs in this class (Figure 3). This description is called schematic because the features described in Figure 3 are not necessary or sufficient to describe the words that undergo this alternation.

Features of a past-tense schema
(a) a final velar nasal
(b) an initial consonant cluster that begins with /s/
(c) a vowel /ɪ/, which has an effect only in conjunction with the preceding elements

Figure 3 - Properties of a productive English past tense schema (Bybee and Moder 1983)

Two types of evidence support the psychological reality of this schema. The first is diachronic. Progressively, since the period of Old English, the class of verbs sharing this alternation has tripled in size (Jespersen 1942). This lexical diffusion (Labov 1994, Kiparsky 1995) suggests that in a large number of chronologically and geographically separate instances, English language users have applied this schema to words that it had previously not been systematically applied to. This sort of historical productivity is often taken as evidence of linguistic knowledge (e.g. Blust 1988).

The second type of evidence for the psychological reality of this schema involves language production. Bybee and Slobin (1982) report on an experiment, which supports the productivity of the schema through a tendency by both children and adults to erroneously produce past tenses that conform to the schema for verbs that do not canonically enter into this alternation. In addition, Bybee and Moder (1983) report on a neologism experiment, in which subjects demonstrated knowledge of the schema through their production of past

tenses for novel verbs. The more similar the novel verb was to the prototype described by the schema, the more likely it was to attract an appropriate strong past tense form.

The existence of associations between a schematic phonological form and a morpho-syntactic class is exemplary of non-categorical relations. When a language user aims to produce a past tense form, there is some probability that the word they chose will have a form associated with the schema in Figure 2. Likewise, when a hearer identifies a phonological form similar to that schema, they can conclude with only some certainty that the word in question is a past tense form.

Nouns and Verbs

The English lexicon displays subtle but significant asymmetries in the distribution of phonological features across grammatical categories. Sereno (1994) demonstrated this fact in a survey of the Brown Corpus (Francis and Kucera 1982). She classified verbs by the front/backness of their stressed vowel: front vowels were [i], [I], [e], [E], and [ae], while back vowels were all others, including central vowels. Sereno found frequent English verbs to more often have front vowels than back vowels, while she found frequent nouns to have more back vowels than front vowels. (Figure 4).

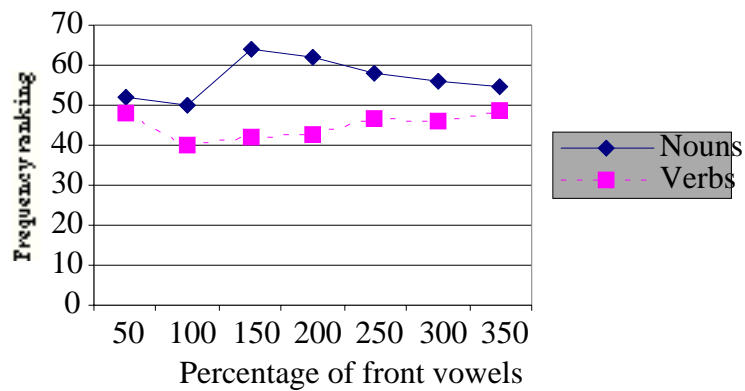


Figure 4 - Front vowels in frequent English words (Sereno & Jongman 1990)

This distributional asymmetry is of little interest unless it can be shown to play a part in linguistic knowledge. Sereno hypothesized that speakers might make use of this asymmetry when processing language. After all, when trying to extract lexical and semantic information from the speech signal, hearers might use any indicators they can get. Sereno (1994) and Sereno and Jongman (1990) then went on to demonstrate that language users do indeed use knowledge of these asymmetrical phonosyntactic generalizations during perception. For example, Sereno (1994) asked subjects to determine as quickly as possible whether a word they were presented with was a noun or a verb. This study yielded the following observations:

- Nouns with back vowels, even infrequent ones, are categorized significantly faster on average (61msec) than are nouns with front vowels.
- Verbs with front vowels, even infrequent ones, are categorized significantly faster on average (7 msec) than are verbs with back vowels.

The startling fact about these findings is that subjects extend the vowel distribution generalization, which holds most centrally of the most frequent words, to relatively infrequent words as well. In other words, individuals pick up on this phonological asymmetry, and under certain circumstances generalize the processing shortcut it allows to the entirety of the lexicon. Just as in the case of the schematic strong past tense described above, language user knowledge of associations between part of speech and vowel quality is non-categorical in nature.

Not only segmental, but also prosodic lexical content correlates non-categorically with grammatical class. English nouns and verbs diverge in their stress patterns. While disyllabic nouns predominantly take stress on their first syllable (trochaic stress), verbs tend to have stress on their second syllable (iambic stress). This distinction is most clearly exemplified by related noun-verb homograph pairs, such as *record*, *permit*, and *compound*. For each of these pairs (and a host of others), the noun is pronounced with trochaic stress (*record*, *permit*, and *compound*), which the verb has iambic stress (*record*, *permit*, and *compound*). This stress contrast in noun-verb pairs is so prevalent that, as reported by Sherman (1975), there are no such homographs which display the reverse pattern: a noun with iambic stress and a verb with trochaic stress.

Indeed, semantically related homographs like these have been argued to be derivationally related (Hayes 1980, Burzio 1994). While this stress-syntax correlation is seemingly categorical among contrastive pairs, it is non-categorical among the noun and verb populations at large. Kelly and Bock (1988) analyzed a random sample of over 3,000 disyllabic nouns and 1,000 disyllabic verbs, and found that 94% of the nouns had word-initial (trochaic) stress, while 69% of verbs had word-final (iambic) stress.

A series of experiments by Michael Kelly and his colleagues have demonstrated that language users capitalize on this asymmetry in both language production and perception. Kelly and Bock (1988) showed that individuals make use of stress patterns when producing novel words. They had subjects pronounce disyllabic non-words in sentences that framed them as either verbs or nouns, and found that they were significantly more likely to give the words initial stress if they were functioning as nouns. Language users also seem to make use of this information during perception. Kelly (1988) showed that when presented with a novel disyllabic word, subjects strongly tended to classify and use it as a noun if it had initial stress. A final piece of evidence comes from a processing experiment by Kawamoto, Farrar, and Overbeek (1990), described in Kelly (1992). In this study, subjects were asked to rapidly classify words by their grammatical category. Subjects classified nouns significantly more quickly if they had word-initial stress than if they had word-final stress, while the reverse was true for verbs.

Not only lexical stress, but also word length correlates non-categorically with part of speech in English. Cassidy and Kelly (1991) found that verbs tend to be longer than nouns, in both adult-adult and adult-child speech. Faced with this observation, Cassidy and Kelly wondered whether adults or children make use of this available information.

Two experiments tested this knowledge. In the first, adults heard novel mono-, di-, or tri-syllabic words, and were asked to use them in sentences. It was hypothesized that if subjects had internalized the correlation between increased word length and increased probability that a word was a noun, then the longer the test token was, the more likely it should be to be used as a noun. This was precisely what Cassidy and Kelly found - adult subjects were about twice as likely to use monosyllabic words as verbs than disyllabic words, even when stress differences in polysyllabic words were controlled for. They found a similar

effect in a second experiment with preschool-aged children, who identified monosyllabic words with actions significantly more often than they did polysyllabic words.

Before moving on to other grammatical categories that can be distinguished on phonological grounds, I should point out that a number of other indicators for the distinction between nouns and verbs have been suggested. These include the longer temporal duration of verbs relative to their homonymic nouns, the greater number of average phonemes in nouns than in verbs (controlling for number of syllables), and the greater tendency for nouns to have nasal consonants (Kelly 1996). (See also Smith 1997 for a discussion of noun-specific properties from an OT perspective.)

Grammatical gender

The noun-verb schism is not the only one that correlates with phonological features. Grammatical gender also seems to have phonological associates in a number of languages. Grammatical gender is a morphosyntactic grouping of nouns of a language into two or more classes, which have different linguistic behavior. For example, in French, nouns are assigned either masculine or feminine gender, and articles and adjectives modifying these nouns must bear surface markings indicating that gender. *Chaise* ‘chair’ is feminine, so it is modified by the feminine definite article *la* ‘the’, while *mur* ‘wall’ is masculine, and takes the masculine article *le* ‘the’. As these examples demonstrate, linguistic gender is not strictly predictable on a semantic basis, although a large body of research has provided evidence that gender categories are motivated by general cognitive principles (Zubin and Köpcke 1986).

But there is also a phonological component to linguistic gender systems. Given the discussion of nouns and verbs above, it should be unsurprising that non-categorical

correlations exist between words of a given gender class and their phonology. It should be equally unsurprising that language users make use of this information. These correlations and their role in language processing have been documented for a number of languages, including French (Tucker et al. 1968), Russian (Propova 1973), Latvian (Ruke-Dravina 1973), German (MacWhinney 1978), and Hebrew (Levy 1983). French, the earliest documented, provides a clear example of this work.

French words with particular endings tend to have either masculine or feminine gender. For example, words ending in *-illion* [ijõ] tend to be masculine, such as *million* ‘million’ and *pavillion* ‘pavillion’. Words ending in *-tion* [sjõ] tend to be feminine, like *action* ‘action’, *motion* ‘motion’, and *lotion* ‘lotion’. Neither of these is particularly productive. Tucker et al. (1968) asked a large number of 8-16 year old French speakers to choose the gender of novel words, which terminated in these and other endings. Their answers tended to follow along the lines of the distributions of those endings in the lexicon. So words ending in *-illion* were significantly more likely to be categorized as masculine than as feminine, while those ending in *-tion* had a much higher likelihood as being classified as feminine. Tucker et al.’s results also indicated that the initial syllable of a noun may have an effect in marking grammatical gender, especially when the ending is an ambiguous cue. Research on the other languages mentioned above has yielded similar results.

Function and content words

We have seen so far how non-categorical correlations can link specific morphosyntactic categories, like the strong past tense, or a particular linguistic gender, or even more general classes, like verbs or nouns, with phonological features. But statistical pairings also

distinguish among more abstract linguistic classes, like function words and content words. Function words belong to closed (unextendable) grammatical classes, such as prepositions, pronouns, and determiners, while content words belong to open (extendable) grammatical classes, like nouns, adjectives, and main verbs. While there are only about 300 function words in English, the rest of the lexicon consists of content words. Function and content words are distinguished by indicators other than extendability, though. The most salient of these is frequency; the 50 most frequent words in Kucera and Francis's frequency count (1967) are function words, while the great majority of content words occur fewer than 5 times out of a million words. Function and content words are acquired at different paces, as well; the conspicuous lack of early function words gives children's speech its 'telegraphic' quality (Radford 1990). Finally, numerous processing differences distinguish function from content words, both in normals and in impaired adults (c.f. the survey in Morgan et al. 1996).

The phonology of syntactic constructions

We have seen numerous examples of statistical correlations between grammatical classes and the phonological content of words of those classes: probabilistic morphosyntactic generalizations. But syntactic constructions, larger than the word, seem also to have statistical phonological correlates.

A particularly well-studied case is the English ditransitive construction (Partee 1965, Fillmore 1968, Anderson 1971, Goldberg 1994). The ditransitive is characterized as taking two objects, neither of which is a prepositional complement. In general, this syntactic structure evokes a giving event, in which the *giver* is expressed by the subject, the *recipient* is expressed by the first object, and the *theme* is expressed by the second object (1a). Both the

verbs that can occur in this construction and the order of its arguments have been subject to phonological scrutiny.

In terms of verbal restrictions, Pinker (1989) argues that phonological constraints apply to the verbs that can occur with ditransitive. In particular, he argues that shorter words are preferred to longer words, which might explain the strangeness of (1c), relative to (1b).

- (1) (a) Subject + Verb + Direct-Object(Recipient) + Direct-Object(Theme)
(b) Ricky slipped me a fiver.
(c) ?Ricky extended me a fiver.

When expressing propositions like those in (1) above, speakers can actually choose between two syntactic constructions, the Ditransitive and Caused Motion (Goldberg 1994). The latter, exemplified in (2b), includes a direct object indicating the Theme and an indirect object denoting the Recipient (3). Arnold et al. (2000) demonstrate that in selecting between the Ditransitive and Caused Motion constructions, speakers take the relative *heaviness* of the two NPs into account. Heaviness can be interpreted in three ways: the structural complexity of a syntactic structure, its length in words, or its prosodic complexity (Inkelas and Zec 1995). But Wasow (1997) shows that the first two of these, length in words and structural complexity, correlate so tightly that these measures make virtually identical predictions about the effects of heaviness. Arnold et al's (2000) finding was that speakers tended to place heavier NPs last; that is, they selected the Ditransitive more often when the Theme was expressed by a longer NP, and Caused Motion more often when the Recipient was expressed by a longer NP. In other words, speakers are more likely to produce sentences like (2b) over (2a), but (4a) over (4b). This skewing is also known in procedural terms as Heavy NP Shift.

- (2) (a) Mom gave a large, steaming bowl of French onion soup to the boy.
(b) Mom gave the boy a large, steaming bowl of French onion soup.
- (3) Subject + Verb + Direct-Object(Theme) + Indirect-Object(Recipient)
- (4) (a) Maddy tossed the boy the teddy bear I had accidentally pulled the eyes off of but slept with until early adolescence.
(b) Maddy tossed the teddy bear I had accidentally pulled the eyes off of but slept with until early adolescence to the boy.

In both of these cases, the use of a particular combination of syntactic and lexical constructions seems to be probabilistically affected by the phonological length of one of the constructional constituents, either length in syllables or length in words. Similar results have been found for Right Node Raising (Swingle 1993).

An open window

The bread and butter of grammar, grammatical classes and syntactic constructions, are both subject to non-categorical phonological generalizations. These generalizations make up part of human linguistic knowledge, and yet most have not been well integrated into linguistic theory. In Chapter 5, I present a modeling architecture that permits us to capture the types of cross-modal, probabilistic, associative knowledge outlined in this section, and in Chapter 6, we will see that there are neural explanations for this type of knowledge.

In the rest of this chapter, we will see that phonological knowledge is not related probabilistically to other grammatical knowledge alone, but to semantic knowledge as well. That is, while up to now we have seen evidence of the non-categoriality of grammar-specific knowledge, in the following two chapters, we will explore how linguistic knowledge relates to knowledge that interacts with the conceptual system.

3. Sound-meaning correlations

Human language is frequently characterized as a finite system with infinite capacities (Chomsky 1995). The key aspect of linguistic systems that allows this property is the possibility of combining linguistic units with others in novel ways. In fact, there are two components to this process. The first is the ability to see linguistic units as composed of smaller units, *compositionality*. The second is the ability to combine these smaller units in novel ways with each other, *productivity*. Many linguistic theories view this productivity and compositionality as defining features of human language. Syntactic, phonological, and morphological productivity and compositionality have been particularly central in shaping mainstream linguistics of the end of the 20th century. The centrality of these capacities in theories of word-formation is entirely unsurprising; a large part of what language users know about their language is indeed how to combine morphemes to construct novel words.

By itself, though, productivity and compositionality are only part of what humans know about the meanings expressed by words of their language. They are also aware of restrictions on productivity from phonological, syntactic, and semantic sources (Bochner 1993). They have unconscious knowledge of groups of words that share meaning and are phonologically or prosodically similar (Serenio 1994, Kelly 1992). They know subtle

correlations between a person's gender and the phonology of their name (Cassidy et al. 1999) and between the number of syllables in a word and the complexity of the object it describes (Kelly et al. 1990). They also know when pieces of a word bear meaning, whether or not those parts can be productively combined with others (Bybee and Scheibman 1999).

And yet, for many theories of the word-internal relation between sound and meaning, the ability to novelly combine units to produce larger ones defines the domain of study. On Baayen and Lieber's view for example, "Processes of word-formation that are no longer productive [...] are of little or no interest to morphological theory" (Baayen and Lieber 1991:802). Many morphological theories view complex words as compiled in production or decomposed in perception. The processes can be seen as the simple concatenation of morphological units, by so-called Item-and-Arrangement models (e.g. Lieber 1980, Di Sciullo and Williams 1987, Halle and Marantz 1993). They can also be seen as the manipulation of a morphological root through some process or processes, by so-called Item-and-Process models (e.g. Aronoff 1976, Zwicky 1985 Anderson 1992).² What these models share is the view that the morphological capacity is something akin to a metaphorical assembly line, which either compiles a complex word from its morphological parts, or takes apart a complex word to analyze the pieces.

And yet, as we will see below, compositionality and productivity as it is usually understood do not adequately capture what language users know about word-internal form-meaning relations. The rest of this section demonstrates that people extract non-categorical form-meaning generalizations from the words in the lexicon and make use of this knowledge when they process language.

²Word-and-Paradigm models (Hockett 1954) do not make strong use of compositionality, though they do depend on productivity.

The challenge of phonaesthemes

Despite its prevalence, the assembly line metaphor falls apart when faced with non-productive morphological elements (cf., for example, Aronoff 1976 or Baayen and Lieber 1991 on metrics for productivity). The assembly line metaphor provides no insight into morphological units what cannot be combined, or cannot be combined easily with others to form complex words. Even more challenging are sub-parts of words that do not combine with any other material: non-compositional morphological units. Obviously, there is no place in the assembly line metaphor for the knowledge of such units, if they are shown to be psychologically real.

Phonaesthemes (Wallis 1699, Firth 1930) are a particularly illustrative example of non-productive, non-compositional morphological units. Phonaesthemes are sub-morphemic sound-meaning pairings, like the English onset *gl-*. *Gl-* occurs in a statistically excessive number of words with meanings related to ‘VISION’ and ‘LIGHT’ (5a). Another well-documented phonaestheme is the English onset *sn-*, which occurs in a large number of words relating to ‘MOUTH’ and ‘NOSE’ (5b).

- (5) (a) *gl-* LIGHT, VISION *glimmer, glisten, glitter, gleam, glow, glare, glint*, etc.
(b) *sn-* NOSE, MOUTH *snore, snack, snout, snarl, snort, sniff, sneeze*, etc.

Phonaesthemes are identified by their statistical over-representation in the lexicon. For example, consider the distribution of these two phonaesthemes in the Brown corpus (Francis & Kucera 1967). 39% of the word types and a whopping 60% of word tokens starting with *gl-* have definitions that relate to ‘LIGHT’ or ‘VISION’ (Figure 5) in an online

version of Webster’s 7th Collegiate Dictionary. 28% of word types and 19% of word tokens with a *sn-* onset have meanings related to ‘NOSE’ or ‘MOUTH’. I am using here the familiar distinction between word types, the list of word entries that appear in the corpus, and word tokens, the number of actual instances of each of those words. Complete lists of each set of word types can be found in Appendix 1.

	<i>gl-</i>		<i>sn-</i>		<i>sm-</i>		<i>fl-</i>	
	<i>Types</i>	<i>Tokens</i>	<i>Types</i>	<i>Tokens</i>	<i>Types</i>	<i>Tokens</i>	<i>Types</i>	<i>Tokens</i>
<i>LIGHT/</i>	38.7%	59.8%	1.1%	0.3%	0%	0%	10.9%	10.1%
<i>VISION</i>	(48/124)	(363/607)	(1/88)	(1/311)	(0/75)	(0/1042)	(28/256)	(146/1441)
<i>NOSE/</i>	4%	1.4%	28.4%	19.0%	25.3%	27.3%	5.6%	2.7%
<i>MOUTH</i>	(5/124)	(9/607)	(25/88)	(59/311)	(19/75)	(284/1042)	(15/256)	(39/1441)

Figure 5 - Phonaesthemes *gl-*, *sn-*, *sm-*, and *fl-* in the Brown Corpus

These overwhelming statistical distributions are quite contrary to what one might predict if there is no non-combinatorial correlation between the sounds of words and their meanings. That is, on the basis of the assumption that a word’s phonological form is entirely arbitrary, given its semantics, the same proportion of *gl-* and *sn-* words should have meanings related to ‘LIGHT’ or ‘VISION’. And yet, as Figure 5 shows, this is clearly not the case. While more than one third of *gl-* word types relate to ‘LIGHT’ or ‘VISION’, only one percent of *sm-* words do. The reverse is true for *sn-* words; while nearly one third of them relate to ‘NOSE’ or ‘MOUTH’, only four percent of *sm-* words share this semantics. One can also compare the proportion of *gl-* or *sn-* words with these particular meanings with the distribution of *fl-* initial words in Figure 5. While some *fl-* initial words have ‘LIGHT’ or

‘VISION’ meanings, like *flicker* and *fluorescent*, these words are less numerous than are *gl-* words with similar meanings. Note that although all these examples are drawn from English onsets, other parts of syllables like codas in words like English *smash* (Rhodes 1994) and even entire syllables in languages like Indonesian (Blust 1988) have been proposed as phonaesthemes.

Morphological models are attempts to understand distributions of lexical and sub-lexical sound-meaning pairings in a language, which makes phonaesthemes at first blush appear to be perfect fodder for such models. However, phonaesthemes like *gl-* and *sn-* are problematic for assembly line theories of morphology. Words that contain these phonaesthemes also contain a complement which does not constitute a unit. For example, removing the onset phonaesthemes from *glint* and *snarl* yields *-int* and *-arl*. It would be problematic both semantically and structurally to say that *-int* and *-arl* were units which contributed to the words as a whole. For example, we cannot mix and match the phonaesthemes and their complements: *glarl* and *snint* are nonsensical. It is equally obvious that no regular morphological process will change a morphological root *gl* into a surface lexical form *glint*. This state of affairs is much like that of bound roots that combine with only one affix, as in words like the verbs *consort* and *refuse*. The difference is that while the affix *con-* of *consort* occurs in other words with roots that have their own combinatorial lives, phonaesthemes never or rarely do.

There are two ways to secure assembly line models against the difficulties posed by non-productive morphological forms. First, it is possible that phonaesthemes are not counter-evidence to these models since they are facts about the distribution phonological content in the lexicon, and not about individuals’ linguistic knowledge per se. This position was advocated earlier in this century by Saussure (1916) who held that onomatopoeic words

are never organic elements of a linguistic system. Bühler (1933/1969) similarly called that onomatopoeia a reversion, since language has evolved beyond primitive needs and means of self-expression. As I will demonstrate in Section 3 below, however, there is ample evidence that phonaesthemes are in fact used in language processing, and that they play a fundamental role in the lexicon's organization.

Phonaesthemes, viewed as distributions of sub-morphemic sound-meaning pairings in the lexicon, are pervasive in human languages. They have been documented in such diverse languages as English (Wallis 1699, Firth 1930, Marchand 1959, Bolinger 1980), Indonesian (McCune 1983), and other Austronesian languages (Blust 1988), Japanese (Hamano 1998), Ojibwa (Rhodes 1981), and Swedish (Abelin 1999). In general, though, this documentation has been performed on the basis of corpus data alone. Only recently has the psychological reality of phonaesthemes been placed under close experimental scrutiny.

The difficulty for assembly line models in dealing with non-productive units like phonaesthemes is mirrored by a methodological one. The standard linguistic metric for evaluating the psychological reality of a unit is to test its productivity and combinability. For example, in her seminal study of regular English plural marking, Berko (1958) argued for the psychological reality of morphological knowledge on the basis of subjects' generalizations of an existing pattern in the language to new forms, such as *wug*. This metric - productivity - is used in morphological study after morphological study.

But a methodological predisposition towards productivity and compositionality leads to the following quandry. How can we evaluate the psychological reality of phonaesthemes, which are generally not productive or compositional, if the only metric we have are productivity and compositionality themselves? A partial solution to this dilemma arises from the observation that phonaesthemes are in fact partially productive.

In the past three years, three theses (Hutchins 1998 and Magnus 2000 on English and Abelin 1999 on Swedish) have addressed the psychological reality of phonaesthemes from essentially the same perspective: that of neologisms. Magnus' work is representative of all three, so I will use her thesis as a model to demonstrate their methodologies. Magnus uses two types of experiment to test the psychological status of phonaesthemes. Both of these are based upon subjects' recognition and production of neologisms. The first methodology tests the productivity of phonaesthemes by providing subjects with definitions for non-existent words and asks them to invent new words for those definitions. In Magnus' study, subjects tended to invent novel words that made use of phonaesthemes of their language. For example, given the definition "to scrape the black stuff off overdone toast", 27% of subjects invented a word that started with *sk*. The second type of experiment tested the use of phonaesthemes in the perception of novel words. Subjects were presented a non-word and were asked to provide a definition for it. Again, they responded as if they were aware of phonaesthemes. For example, *glon* (which starts with our infamous *gl*-onset) evoked definitions relating to 'LIGHT' in 25% of the cases.

These studies each provide independent evidence for the role of phonaesthemes in the processing of neologisms. Magnus (2000) interprets her results as evidence that individuals are aware of the semantic content of individual phones. Abelin (1999) adds to this the insight that the fewer words there are that share the phonology of a phonaestheme, and the more common that phonaestheme is among the words sharing its phonology, the more likely subjects are to treat neologisms along the lines of that phonaestheme. Hutchins' (1998) results confirm the role of phonaesthemes, even cross-linguistic ones, in the processing of neologisms. Interestingly, they are qualitatively similar to results of neologism

experiments on the Obligatory Contour Principle in Arabic (Frisch and Zawaydeh 2001) and Labial Attraction in Turkish (Zimmer 1969).

These studies make intuitive sense, given what we know about lexical innovation in fiction. Probably the most famous use of literary neologisms are in *The Jabberwocky*, a poem that confounds young Alice in Lewis Carroll's (1897) *Through the Looking Glass*. The first stanza of this poem appears in (6) below.

(6) 'Twas brillig, and the slithy toves
Did gyre and gimble in the wabe:
All mimsy were the borogoves,
And the mome raths outgrabe.

As to what these words mean, Humpty Dumpty is gracious enough to explain the poem to Alice: "Well, 'slithy' means 'lithe and slimy [...]" (Carroll 1897) and so on. This example is particularly illustrative of the problem with neologism studies.

Although they provide valuable insight into how language users interact with new words, neologism studies do not by themselves constitute conclusive evidence for the psychological reality of phonaesthemes. Even in light of these neologism results, one could still hold the position that phonaesthemes are only static, distributional facts about the lexicon, which speakers of a language can access consciously. This is problematic since most all of language processing happens unconsciously. That is, we know that language-users are able to access all sorts of facts about their language upon reflection. People can come up with a word of their language that is spelled with all five vowel letters and 'y' in order, or a

word that has three sets of double letters in a row. (See below for the answers.³) This ability by itself doesn't lead us to conclude, though, that orthographic order of vowel letters in a word is a fundamental principle of implicit cognitive organization. If it were combined with convergent studies confirming the status of this knowledge in unconscious processing, it would. For the same reason, subjects' ability to consciously access distributions of sound-meaning pairings in their language does not imply that those pairings have been extracted from their lexical origins to become part of the subjects' linguistic system.

In fact, different varieties of model could be constructed to account for the conscious cognitive processing that the neologism studies described above demonstrate. For example, subjects could be taking a sample of the lexicon, and finding a set of words that share structure with the neologisms. Or they might be randomly selecting nearest neighbors. Alternatively, they could be looking for cues in the stimulus itself; notice that the word 'scrape' appears in the example for *sk*- provided above. A final possibility is that subjects just might be unconsciously accessing phonaesthemes. Only this last, unconscious processing account supports the hypothesis that phonaesthemes enjoy a place in language processing. But we can't distinguish between any of these solutions on the basis of neologism tasks like the ones described above.

What we need, then, is a way to tap into unconscious language processing. To reiterate, the stakes are high. If we can demonstrate phonaesthemes to have psychological reality, then the assembly line model of morphology will have to be reconsidered; the notion of morphemes as concatenative, productive units will be insufficient for defining the human capacity to extract and use generalizations about the internal structure of words. From a

³Facetiously and bookkeeper, respectively.

broader perspective, phonaesthemes can provide key evidence for or against the modularity and determinacy of phonological and semantic knowledge.

4. An experimental study of phonaesthemes in language processing

From the perspective of their distribution, phonaesthemes are statistically over-represented partial form-meaning pairings. In order to test whether phonaesthemes also have a psychological reality - whether they are internalized and used by speaker-hearers - we will have to answer two related questions. First, does the presence of a phonaestheme in a word affect the processing of that word? If phonaesthemes play a role in language processing, then words they occur in should be processed differently from words that do not. Second, if phonaesthemes do in fact affect lexical processing, is this effect significantly different for phonaesthemes than for subparts of words that correlate some form and some meaning, but not in a statistically significant way in the lexicon? In other words, is there a role for the frequency of a form-meaning pairings in determining processing effects?

The morphological priming methodology (used first by Kempley and Morton 1982, summarized by Drews 1996) provides an excellent starting point for developing a methodology capable of addressing such experimental questions. Morphological priming is the facilitation (speeding up) or inhibition (slowing down) of mental access to a TARGET word on the basis of some other PRIME word, which has been presented previously. A large number of morphological priming studies have found priming patterns that are unique to morphologically related PRIME and TARGET words, and are not shared by morphologically unrelated words (see Feldman 1995). On the basis of this difference between morphological priming effects on the one hand and phonological and semantic

priming effects on the other, it has been argued that morphemes are psychologically real. In the same way, we can hypothesize that if phonaesthemes have some cognitive status, then they should similarly display priming effects that are different from those exhibited by words that do not share phonaesthemes. The study described below aimed to test this possibility.

In this study, 20 native speakers of English, aged 18 to 36 were first presented with a PRIME stimulus, which was an orthographic representation of a word, on a digital computer. It appeared slightly above the center of the screen for 150msec, just enough to be barely perceived by the subject. 300msec later, a second stimulus, the TARGET stimulus, also a typewritten word, was presented in the center of the screen for 1000msec or until the subject responded to it. This inter-stimulus latency of 300msec was chosen because it most effectively pulls apart morphological from non-morphological priming effects (Feldman and Soltano 1999) - at this time delay, there is the least phonological or semantic priming relative to morphological priming. Subjects performed a lexical decision task; they were asked to decide as quickly as possible whether the TARGET was a word of English or not. They were to indicate their decision by pressing one of two keys on the computer keyboard - 'z' for non-words and 'm' for words.

The stimuli were 100 pairs of words and pseudowords, falling into five categories (Figure 6). In the first condition, the PRIME and TARGET shared some phonological feature, always a complex onset, and some semantic feature, such as 'NOSE' or 'LIGHT'. Additionally, a significant proportion of the words in English sharing that onset also shared the semantics, giving pairs like *gleam* and *glitter*. These words were all selected by taking a statistical measure of how many words beginning with a purported onset phonaestheme shared a given semantics. In the second condition, the PRIME and TARGET shared only a phonological feature, always an onset, yielding stimuli such as *class* and *def*. PRIMES and

TARGETs in the third condition shared only some semantics, like *blue* and *green*. In the fourth condition, the pseudo-phonaestheme case, stimuli shared phonology and semantics but were a very small number occurring in the Brown corpus and did so. An example of such a pair is the nouns *flock* and *fleet*, which share the semantics ‘AGGREGATE’. Finally, for comparison, stimuli in the baseline condition did not share an onset or meaning, giving pairs like *play* and *front*.

<i>Condition</i>	<i>Characteristics</i>	<i>Examples</i>
1 Phonaestheme	PRIME and TARGET shared a semantic feature and a phonological onset and were a statistically significant subclass of the lexicon.	<i>glitter:glow</i>
2 Form	PRIME and TARGET shared an onset.	<i>class:clef</i>
3 Meaning	PRIME and TARGET shared a semantic feature.	<i>blue:green</i>
4 Pseudo-Phonaestheme	PRIME and TARGET shared a semantic feature and an onset but were not a statistically significant subclass of the lexicon	<i>flock:fleet</i>
5 Baseline	PRIME and TARGET were unrelated in form and meaning	<i>play:front</i>

Figure 6 - Five test conditions with examples

All categories were matched for average token frequency, since frequency of PRIME and TARGET have been demonstrated to have significant effects on morphological priming (Meunier and Segui 1999). PRIMES of all classes had an average frequency of

approximately 7 in the Brown Corpus and TARGETS 20. All categories were also matched for word length, with an average of approximately 6 characters and 1.1 syllables, each.

Translated into specific hypotheses about the result of this study, the questions at the top of this section appear as follows. First, if the presence of phonaesthemes affects processing, then responses to condition 1 (the phonaestheme case) should be significantly different from those to condition 5 (the baseline, unrelated condition). To be certain that this effect does not derive from priming effects resulting from known phonological (Zwitserslood 1996) and semantic relatedness (Thompson-Schill et al. 1998) between PRIME and TARGET, the responses to class 1 must also be distinct from those to classes 2 and 3. Second, if this priming effect is restricted to statistically over-represented form-meaning pairs, then responses to condition 1 (phonaesthemes) should be significantly different from those to condition 4 (pseudo-phonaesthemes).

Results

The average reaction times per condition fell out as shown in Figure 7. By far, the fastest average reaction time was to the phonaestheme condition, which was 59 msec faster than the baseline (unrelated) condition. That is to say, there was an average facilitation of 59 msec when the PRIME and TARGET shared a phonaestheme, relative to the condition where PRIME and TARGET were semantically and phonologically unrelated. This is an indication of phonaesthetic priming. By comparison, PRIME and TARGET pairs which shared only some meaning but no form were identified only 23msec faster than the baseline condition, the semantic priming effect. By contrast, the form priming effect was in fact inhibitory, causing form-related TARGETS to be identified 3msec more slowly than unrelated

TARGETS. In other words, there was a significantly greater degree of priming between phonaesthemically related words than between words that shared only form or meaning.

But what of the role of statistical over-representation in the lexicon? Remember that the phonaestheme condition elicited reaction times 59 msec faster than the baseline. By comparison, the pseudo-phonaestheme condition, containing words that shared form and meaning but were the only words in the lexicon sharing both, yielded reactions only 7 msec faster than the baseline. This indicates that the extent of the form-meaning pairing's distribution in the language is crucial to its role in processing. Specifically, statistically prevalent form-meaning pairings yield more facilitory priming than do similar words that are statistical loners. The differences between the phonaestheme condition and each other condition were determined all by a two-way ANOVA to yield at worst a significance of $p < 0.02$ - the phonaestheme condition was significantly different from all other conditions.

	<i>Phonaestheme</i>	<i>Form</i>	<i>Meaning</i>	<i>Pseudo-P</i>	<i>Baseline</i>
<i>Average (msec)</i>	606.7	668.2	642.7	658.7	665.3
<i>Priming effect (msec)</i>	59	-3	23	7	0

Figure 7 - Average TARGET reaction times and priming effects by condition

The graphical representation of these results in Figure 8 shows the degree of facilitation engendered by the presence of a phonaestheme in a PRIME-TARGET pair.

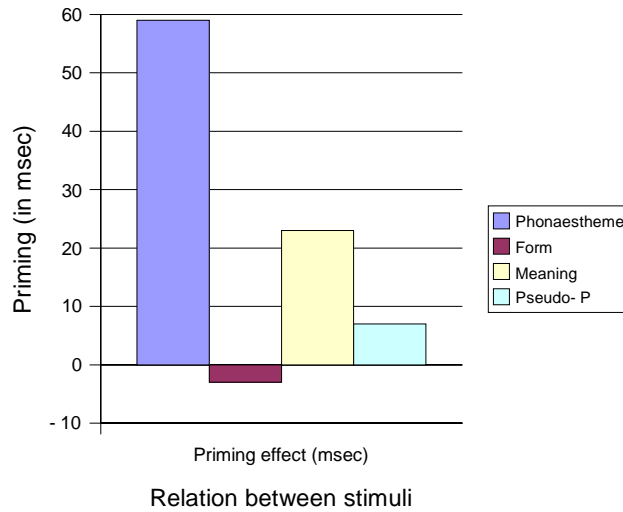


Figure 8 - Priming by relation between stimuli

Summary of results

These results very clearly indicate facilitory priming by both semantically and phonaesthemically related PRIMEs, but not by phonologically related PRIMEs. Moreover, the phonaesthetic priming cannot be accounted for simply on the basis of other sorts of priming. The meaning priming and form priming both yield quantifiably different priming effects from phonaesthetic priming. As a potential complication, it has been proposed in the morphological priming literature that since stimuli that share a morpheme share both form and meaning, the sum of the priming effects from each of these domains might give rise to the actual priming observed with phonesthemes (Feldman and Soltano 1999). However, taken together, form and meaning priming are still distinct from phonaestheme priming; their sum (20mesc) differs significantly from the phonaesthetic priming effect (59mesc). No other metrics aside from summation have been proposed in the literature to deal with combined phonological and semantic effects. Phonaesthetic priming is an entity

unto itself, resulting from the statistical over-representation of a particular sub-morphemic pairing between form and meaning in the lexicon.

These results disconfirm the possibility that language-users access phonaesthetic knowledge only during tasks allowing reflection, such as the neologism tasks described in section 1.2., above. Instead, phonaestheme effects emerge they even while individuals are performing a task that is tightly constrained by time pressure, and is therefore processed unconsciously, as is natural language.

Discussion

We have seen that it is impossible to deny the psychological reality of phonaesthemes. But what of their relation to other morphological knowledge, for examples known affixes?

It seems that in terms of three definitional criteria, phonaesthemes fall fully within the range of normal morphological phenomena. Unlike the most frequent and productive affixes, but like other semi-productive units like *en-* (Baayen and Lieber 1991).

phonaesthemes are only slightly productive. Again unlike prototypical affixes, but like bound morphs like the roots of the verbs *consort* and *refuse*, phonaesthemes cannot be readily combined with other morphological material. Most importantly, as we will see below, phonaesthemes align tightly with canonical morphemes in terms of their priming behavior.

In general, morphologically-related forms have been shown to demonstrate facilitory priming that differs from both semantic priming and phonological priming in terms of its degree and time course. For example, an experiment described by Feldman and Soltano (1999) was nearly identical in format to the one described above, except that the test stimuli, instead of being phonaesthemically related, were morphologically related. Unlike the

phonaestheme stimuli, the morphologically related forms in Feldman and Soltano’s experiment shared a productive suffix, like the regular past tense. Figure 9 lays out the priming effects from Feldman and Soltano’s experiment in comparison with those from the phonaestheme experiment.

<i>Relation Between PRIME and TARGET</i>				
	<i>Morphological</i>	<i>Phonaesthemic</i>	<i>Formal</i>	<i>Semantic</i>
<i>Feldman & Soltano</i>	49	--	-18	34
<i>The present study</i>	--	59	-3	23

Figure 9 - Morphological priming (Feldman and Soltano 1999) and phonaesthemic priming (this chapter) in milliseconds

We cannot directly compare the recognition speeds in these two experiments, because the test conditions differed slightly. We can nevertheless draw a valuable conclusion from this juxtaposition. The phonaesthemic priming detected in the experiment described above is qualitatively similar to evidence for morphological priming. In both cases, special status is accorded to a hypothesized unit of organization on the basis of priming that cannot be accounted for on the basis of form priming and meaning priming alone.

Similar Studies

Further support for the position that probabilistic sound-meaning pairings constitute central, rather than peripheral aspects of human linguistic knowledge comes from the wide range of sound-meaning pairings that have been documented by other researchers.

Research on a number of seemingly unrelated topics has indicated that language users integrate and make use of probabilistic correlations between sound and meaning, even when these relations do not play a productive role in the linguistic system. For example language users have internalized subtle correlations between a person's gender and the phonology of their name. Cassidy et al. (1999) documented these correlations and then presented adults and four-year-old children with words that bore statistically prevalent characteristics of female or male first names and asked them to perform neologism and sentence completion tasks. For example, while male names like *Richard* and *Arthur* tend to have trochaic stress, female names like *Irene* and *Michelle* tend to have iambic stress; male names like *Bob* and *Ted* tend to end in consonants, while female names like *Sue* and *Mary* tend to end in vowels. Cassidy and her colleagues found that subjects tended to process the names along the lines predicted by the phonological cues to gender.

In another study, Kelly et al. (1990) found that adults and children have internalized and make use of the correlation in English between the number of syllables in a word and the complexity of the object it describes. When presented with novel shapes, they were more likely to pair complex shapes with polysyllabic words than they were simpler shapes. Geometric shapes are particularly exemplary of this phenomenon; consider the tendency expressed by the visual complexity and syllabicity of *line*, *square*, and *point* with those of *trapezoid*, *hexagon*, and *parabola*.

Finally, a segment's morphological expressiveness correlates probabilistically with the extent to which speakers are willing to reduce that segment. English has a word-final deletion process which hits the final coronal stops of many words in most dialects. It has consistently been shown that individuals reduce the final *t* or *d* of a word more, the more frequent the word in question is (Labov 1972, Guy 1980, Bybee 2000). But in addition to

this effect, if a final coronal stop has some morphological status, if, for example, it indicates the past tense such as in the words *tossed* and *buzzed*, then it is less likely to be reduced. Just like probabilistic effects from gender and visual complexity, this meaning-bearingness effect provides external support that language users internalize probabilistic form-meaning correlations and make use of them during language processing.

5. Future Directions

Clearly, the present situation begs the question of the precise relationship between morphological and phonaesthetic priming. Only an experiment that combines those two types of test stimuli can answer this question. Also of future interest is the specific role that statistical frequency plays in phonaesthetic priming. Meunier and Segui (1999) have shown that frequency plays a role in morphological priming. We don't yet know what the exact character of the effect of frequency in phonaesthetic priming is. Is it an all-or-nothing feature, or is degree of frequency correlated with degree of phonaesthetic priming? The present study did not include frequency as a graded parameter, but rather as a discrete one; statistically significant phonaesthemes were assigned to the same condition, regardless of how significant they were. We can therefore provide evidence on this question. Another factor that might play a role in phonaesthetic priming is the degree of semantic overlap between PRIME and TARGET. This factor has been shown to play a role in the degree of morphological priming between stimuli, and might also do so for phonaesthemes (Laura Gonnerman, p.c.).

In addressing the psychological reality of phonaesthemes, I have addressed only the degree to which phonaesthetic knowledge has been internalized by language users. Other

questions about phonaesthemes deserve attention, though. For example, the study described above assumes that the potential universality of or motivation for phonaesthemes is orthogonal to their representation. In fact, phonaesthemes have been predominantly studied in terms of their sound-symbolic properties. Various studies have attempted to document cross-linguistic phonaesthemes for the purpose of demonstrating that there is some correlation between the sound structure of words that have phonaesthemes and the meanings they bear, such as the often reported use of high vowels to express small size and low vowels to indicate largeness (e.g. Ohala 1984). The results of these studies are frequently interpreted as disproving the arbitrariness of the sign, e.g. Bolinger (1949). Both the results and this conclusion, though, are hotly contested.

It seems at one level essential to pull the *motivation* for a phonaestheme (the diachronic, conceptual, or developmental reason for its occurrence) apart from its *distribution* (its presence in the lexicon) and its eventual *representation* (the manner in which it is instantiated in an individual mental linguistic system). These are qualitatively different sorts of issues. While the motivation describes ultimate historical or acquisitional causes for the phonaesthemes existence, the distribution is a measure of the phonaestheme's synchronic use, and the representation is part of a model of an individual speaker's linguistic knowledge. Nevertheless, there may in fact be some relationship between the semantic motivation for a phonaestheme and its distribution. More semantically unified or more semantically "basic", however basicness is measured, might tend to be more prevalent. One way to test these hypotheses would be to repeat the experiment presented above with phonaestheme stimuli which varied in terms of their proposed universality.

Whatever direction phonaestheme research moves in, the take-home message of the present study is that like other non-categorical pairings between phonology and semantics,

and between phonology and grammatical features, phonaesthemes have a proven psychological status. How might phonaesthemes be learned, and represented? What is the neural basis for priming, and what does that tell us about the role of associative learning in linguistic models? While we will have to wait until Chapter 5 to see answers to these questions, in the next chapter we will find that phonology is also non-categorically paired with social characteristics of individual speakers.

Chapter 3. Social factors and interactions in French liaison

Outline

1. Introduction
2. Variability in the language of the community
3. Interactions between factors
4. French liaison and the liaison corpus
5. A statistical test of autonomous factors
6. A statistical test of interactions between factors
7. Final note

I don't want to talk grammar. I want to talk like a lady.

Eliza, in Pygmalion, George Bernard Shaw

1. Introduction

Complementing the non-categorical co-variation between phonological patterns and grammar or syntax that we saw in the last chapter are a wide range of non-categorical social effects on phonological form. Much of sociolinguistic research focuses precisely on the role of social factors in conditioning *variation* in the phonological realization of linguistic units (Chambers 1995).

As we will see in this chapter, not only do social factors non-categorically influence phonology, but in fact they sometimes do so in an *interacting* fashion. In other words, some probabilistic factors do not contribute to phonology in the same way in all contexts. Sometimes the degree or even quality of their contributions will be affected by the values of other extraphonological variables that also have effects on phonology.

After presenting the issue of socially correlated variation in the language of the community, as well as the problem of interacting factors, I will move on to demonstrate these phenomena through the statistical analysis of factors influencing French liaison (Tranel 1981) in a large corpus of spoken Swiss French. French liaison is a phonological pattern that is subject as we will see to numerous *autonomous* and *interacting* factors. The work of documenting these factors in this chapter is done in the interest of investigating, in the next chapter, the extent to which language hearers make use of non-categorical effects on French liaison, in particular interaction effects, during language processing.

2. Variability in the language of the community

When Eliza Doolittle, the young, impoverished, cockney flower seller of *My Fair Lady* finds herself wishing for a higher social station, she doesn't go seek the council of a hairstylist or a tailor; she heads to linguist Henry Higgins. Through the course of *My Fair Lady*, we witness Eliza's linguistic transformation away from her vernacular cockney, riddled with glottal stops. With the help of Higgins, who wagers he can transform her in half a year, she learns to produce her glottal fricatives in 'Hartford, Hereford, and Hampshire' and to readjust her diphthongs in 'Rain' and 'Spain'.

Eliza's motivation is the all too painful realization that characteristics of her speech allow others to immediately identify her social status. Her working class speech cues are theatrically highlighted in the movie through repetition and juxtaposition with Higgins' standard pronunciation of the same words. But in real life, language users rely on less explicit presentation of sociolinguistic indicators.

Some of these indicators are categorical, such as Eliza's unwillingness to articulate her [h]s. Most, though, are non-categorical. That is, they are just like the probabilistic relations between sound patterns and extraphonological features we saw in Chapter 2. They provide hints about extraphonological properties, in Eliza's case the speaker's social status, but not hints that even profligate gambler Henry Higgins would be likely to bet a lot of money on. In this section, I will list some of the ways that aspects of phonology can vary as a product of social characteristics of the speaker.

Types of variability

In the end, Eliza finds herself quite adept at producing realistic-sounding standard British English, to the point where Higgins wins his bet and Eliza wins the fancy of the dashing and cultured Freddie. The moral of Eliza's story for present purposes is that those properties of Eliza's speech that tagged her as a flower seller were *mutable*. They could be, and eventually ended up being, altered.

The sociolinguistics literature is rife with studies of socially relevant mutability in phonology and syntax. Linguistic elements whose production varies as a function of social aspects of the speaker are known as social variables. Mutable social variables are the product of social and linguistic exposure and group-identification. As a consequence, they tend to involve differences among individuals, depending upon their social characteristics.

An early example of a mutable phonological variable involves the pronunciation of /ŋ/ in Norwich, England (Trudgill 1972). The suffix 'ing' of *walking* and *laughing* in the received pronunciation includes a mid front vowel and a velar nasal: [ɜŋ]. In Norwich, though, it can surface as alveolar, syllabic [nɪ]. The probability that 'ing' will be reduced to

[ŋ] in Norwich depends on the speaker's social class. As shown in Figure 1, in a casual speech style, middle class speakers on average use the received pronunciation much more frequently than do working class speakers. In this figure, 'L' = 'Lower', 'M' = 'Middle', 'U' = 'Upper', 'W' = 'Working', and 'C' = 'Class'.

	MMC	LMC	UWC	MWC	LWC
Percent of [ɪŋ]s produced	72%	58%	13%	5%	0%

Figure 1 - The average production of 'ing' by middle and working class speakers in Norwich, in fluent speech (from Trudgill 1974)

Montreal French vowels similarly show socially-based variability. Santerre and Milo (1978) demonstrated that age plays a central role in determining the diphthongization of certain vowels. At the end of a stress group, or at the end of a word before a pause, Montreal French vowels can be produced as monophthongs or diphthongs. For example, /ø/ can be diphthongized to [œ^ø]. Diphthongization of /ø/ is more frequent when the speaker is female than male; with women on average producing diphthongized variants around 35% of the time, and men about 22%.

Mutable variables such as Eliza's /h/, Norwich's /ɪŋ/, and Montreal's /ø/ depend upon social properties of the speaker. As a consequence, the different social histories and identities of different speakers give rise to variation among speakers, and these mutable variables are most clearly observed as cross-speaker variation in the language of the community. Mutable variables can also be manipulated for stylistic or other social purposes. If Eliza chose to slip back into her cockney ways when talking to her father, than anyone observing her would be facing a case of speaker-internal variation; sometimes she selects one

variant, other times the other, possibly with other social correlates, like the identity of her interlocutor. Social action is among language's main functions. Language is a tool, used for establishing or testing social status or social identity, constructing or destroying social connections, and carving out the limits of group membership. When we study mutable social variables, we investigate the impact that this social dimension of language has on the details of linguistic structure.

By contrast with mutable variables, which depend on the social environment in which a speaker learns a language, immutable variables depend on inherent, biological characteristics of individual language users. For example, the shape and size of the vocal tract bear a direct influence on the speech signal. A larger larynx yields a lower fundamental frequency. Individual differences in the shape of the vocal tract give rise to speaker-specific differences in the frequency of F3. Biological sex is a central factor that contributes to determining characteristics of the vocal tract, including nose and larynx size; human males tend to have larger larynges and noses than their female counterparts. Variability in these immutable properties is observed predominantly at the level of the production of a community. Individuals will vary in terms of their innate biological properties, and these differences give rise to variability, once again, across speakers.

Note that even presumably immutable properties can be modified. Sometimes the modification is quite convincing, too. Most viewers were convinced they were watching a female performer seduce the male protagonist, Fergus, in 1992's *The Crying Game* until presented with prima facie evidence that the seducer was in fact biologically male. But while we are somewhat taken aback when an Eliza Doolittle, dressed elegantly for a day at the races, begins to swear like chimney sweep, we are utterly confounded upon realizing that we

have been led to miscategorize biological gender by a high fundamental frequency, breathiness, and other prototypical female voice characteristics.

As we have seen, both mutable and immutable social factors can influence phonological variables non-categorically. In the next section, I'll underline the relation between social variability in the language of the community and the linguistic system of the individual.

Non-categoriality in individual production

The line of inquiry we are following through this thesis involves the linguistic system of an individual language user. What is the cognitive structure of the human linguistic system? What does it have to do with the rest of cognition? With the human body? What information do individuals access when they perceive language? We have seen thus far that phonology correlates non-categorically with other domains of knowledge, and that individuals make use of these correlations in perception. Now that we have established that social properties also covary with phonological ones, we can ask to what extent these correlations play a part in individual linguistic systems.

On a purely theoretical plane, there are two potential paths by which social variability as discussed in the previous section might enter into the cognitive system of an individual language user. In the first, the variability displayed by a community is played out at the level of individual production. This is demonstrated de facto by inter-speaker variation, as seen above, and expanded on in this section. In the second path, individual language users make use of social variability during language perception – making social assumptions about speakers on the basis of their phonology, and making phonological assumptions on the basis

of believed social aspects of the speaker. The potential role of social variability in perception is explored in the following section.

In fact, there are at least two ways in which the production of a linguistic community can be reflected in an individual's speech patterns. The first is the most straightforward; individual speech might be deterministic, and social variability a speech community-level phenomenon. We might study some social variable and discover, for example, that middle working class speakers in Norwich produce /*ŋ/ as [nɔ] 5 percent of the time. If this is the case, then it could be that what that 5 percent represents is the percent of working class speakers who pronounce /*ŋ/ as [nɔ] all the time. This sort of variation is known as *inter-speaker variation*, because only when we look across speakers is there any degree of non-categoriality. If social variation had only this sort of impact on individual speech, then we might be justified in dismissing the social patterns as irrelevant to individual grammar.

But another possibility is that individual percentages could reflect class percentages. For example, in the Norwich case, each middle working class speaker could produce /*ŋ/ as [nɔ] about 5 percent of the time. This variable production by individuals might also depend on other aspects of the speech context, such as the style, register, or other aspects of the context. This scenario is known as *intra-speaker variation* since one can observe non-categoriality in a single speaker. Although I've depicted this reflection of community usage patterns in the individual as a separate possibility from categorical behavior, hybrid cases are also possible; there might be both categorical and reflective speakers in a given speech community.

Unfortunately, Trudgill (1972) does not tell us which scenario obtains in the Norwich /*ŋ/ data. However, the second, intra-speaker variation scenario actually does play out in numerous cases. Fasold (1978), for example, cites a number of studies that show just

this. He cites studies on final stop deletion, in black peer groups in Harlem (Labov et al. 1968), in Philadelphia (Guy 1977), in Detroit (Wolfram 1973), and in Appalachian English (Wolfram and Christian 1976). In addition, he points to studies of postvocalic /r/ deletion in Appalachian English (Wolfram and Christian 1976) and the case of Montreal French *que* deletion (Sankoff 1974), both of which show the same trend: individual speakers follow the patterns of their social groups, when a large database of their speech is collected.

In other words, there's evidence that the probabilities that characterize the use of a variable phonological rule by a language community are reflected in individual production. But when a phonological variable varies along immutable social dimensions, the effects of different social categories are less likely to be represented in an individual's speaking patterns. Individual speakers can modify their language for social purposes, in ways that are generally studied under the rubric of register or style. But this does not mean that the full range of socially correlated variation is reflected in intra-speaker variation. Aside from those individuals whose personal interests lead to trans-gender identification, most members of a community retain the same gender identity throughout their lives. Although individuals are able to vary sex-driven speech characteristics to some extent, intra-speaker variability is always restricted by biological and social role constraints. Much the same is true for class affiliation in many societies. Despite our best efforts to resist it, age, too, is fundamentally immutable. (Granted, age changes over time, but it's not intentionally mutable.) The linguistic reflections of these social categories are even more likely, therefore, to enter into individual cognition at the level of language perception than in individual language production.

Non-categoriality in individual perception

While we have seen that social variability correlates in some cases with individual speech variability, we have yet to establish what effect if any social variation has on individual language perception. We saw in the previous chapter that both grammatical and semantic non-categorical correlates of phonological patterns are internalized by speakers when they perceive speech. In order to understand how these correlations are learned and represented, we must first establish what domains of knowledge can enter into these non-categorical pairings. The rest of this chapter and the following chapter are dedicated to demonstrating that language hearers pick up on social correlates of phonology, which implies that even extralinguistic factors can correlate non-categorically with phonological variables.

In fact, there is relatively little research in this area. One of the main reasons for this is the notion of normalization that is prevalent among linguists. As Johnson and Mullennix (1997) describe it, the conventional view of language processing in the face of variability looks as follows:

“[...] canonical linguistic representations are derived from the speech signal. In these accounts, after the system makes a response to variation in the signal and derives the canonical representation, information about nonlinguistic variation is discarded.”

(Johnson and Mullennix:1)

This prevalent notion of the role of the social correlates of linguistic variation places most of the processing burden on the extraction of an invariant linguistic representation..

A number of recent studies, however, indicate that speaker variability is in fact available to language users during speech perception. Goldinger (1997), Nygaard and Pisoni

(1998), Bradlow et al. (1999), and Goldinger et al. (1999) have all established, using different experimental methods, that a great deal of phonetic detail associated with speaker identity is encoded in episodic (short-term) memory during language processing.

Just one example of this type of research should suffice. Bradlow et al. (1999) presented subjects with words, produced by a number of different speakers, and also varying in rate and amplitude. They asked subjects to determine whether they had previously heard words they were presented with. Subjects could respond that the word was old or new. Across different experimental setups, Bradlow and colleagues found that if a word had been uttered previously by the same speaker, it was more likely to be identified as ‘old’ than if it had been previously uttered by a different speaker. By comparison, amplitude did not have an effect on identification - just because a word had been previously heard at the same amplitude, that didn’t make subjects any more likely to recognize that they had already heard it. Exactly what speakers-specific characteristics subjects were picking up isn’t clear, but the study does imply that some aspect of speaker identity is kept in memory along with the lexical token.

These studies and others like them have led to an alternate view of “normalization” in speech processing. So-called *exemplar* models (Johnson 1997 and Pisoni 1997) place more weight on complex and explicit representations of perceptual inputs and less on the processes that arrive at those representations on the basis of the speech signal. Give two examples of such models. Exemplar models of speech processing view speech perception not as the extraction of invariant linguistic representations from an inherently noisy signal, but rather as the internal encoding of a linguistic input in significant phonetic detail. The argument for exemplar based models of perception is based on three major points. Acoustic inputs are highly variable, and depend on speaker-specific and group-specific properties.

They are also multidimensional, with some dimensions co-varying, and others standing independent. Finally, the extralinguistic details of the signal are informative for linguistic processing, and information about a given dimension of the signal relevant to other dimensions.

Research like the work cited above that supports exemplar models thus demonstrates that the variable, multidimensional phonetic detail of the signal that is available for linguistic processing is put to use by hearers.

The challenge

While we have seen in this section evidence that detailed phonetic information, indicating social aspects of the speaker, plays a role in speech perception, we have yet to establish its relation to phonology. The patterning of speech sounds might be completely autonomous from the social detail that individual language perceivers seem to memorize and recruit when understanding language (as was demonstrated in the previous sub-section). It might also be independent of socially conditioned variability in speech production.

In investigating syntactic, morphological, and semantic correlates of phonological cues in Chapter 2, we saw that language users have internalized these correlations and make use of them during unconscious language processing. No studies that I am aware of take up the questions of whether and how social variability affecting phonology is implemented in language perception.

The challenge we find ourselves with at present is to find a case of socially varying phonology in the speech of a language community, and to determine the extent of its internalization in individuals' language processing behavior. If this challenge can be met, it

will establish the place of probabilistic socio-phonological correlations in human language processing, alongside phono-syntactic, phono-morphological and phono-semantic ones. Section 4 is dedicated to describing one phonological variable with myriad probabilistic correlates. French liaison is a phonological phenomenon with phonological, morphological, syntactic, and social correlates. We will see in a large corpus study in Sections 5 through 7 of this chapter that many of these correlates are significant in the language of the community. We will first run across a complicating factor, though; some of these correlates are not independent of each other. In the next section, I will describe the problem of interacting factors on phonological variables. In the next chapter, two perception experiments will test the psychological reality of the autonomous and interacting factors, including social ones, and will establish the place of social factors in phonological processing.

3. Interactions between factors

In the previous section, I surveyed evidence that social factors correlate non-categorically with phonological patterns, and that this information is available in the speech signal and speech context for perceivers to take advantage of. In the present section, I demonstrate that social and other factors make more complex contributions to phonology than the situation depicted up the present indicates. At the heart of this complexity are interactions between contributing factors. In this section, we will be introduced to French liaison, a phonological process which is subject to effects from a large number of social and other constraining factors, as well as, I argue, interactions between some of these factors.

In the language of a community, phonological variability correlates with variation in social characteristics of a speaker, as established in the preceding chapter. But these

correlations are sometimes mediated by other factors, which modify their effects. In such cases, we say that factors *interact*.

When there is no interaction between factors, we say that they are *autonomous*. To make this distinction clear, let's consider a classic study by Labov (1966) on New York English. His research in New York was an effort to identify the social determinants of the variable production of /r/ (along with three other linguistic variables). Labov hypothesized that /r/'s realization as [r] or \emptyset is not due to pure free variation, as Hubbell (1950) and others had suggested. He sought out social factors that might condition the realization of /r/: *independent* factors.

In his study, Labov conducted interviews with New Yorkers of varying socioeconomic status, which he grouped into three categories: Lower, Working, and Middle class. But aside from class, Labov also investigated the contribution of another contributing factor: the style of speech a speaker was engaged in. This variable he classified into five distinct types, in order of decreasing formality: reading minimal pairs (MP), reading isolated words from a list (WL), reading a passage (RP), conversation with an interviewer (IS), and conversing casually with a peer (CS). So when studying the effects of these two independent social variables on the occurrence of /r/, Labov had a total of 5 style conditions times 3 class conditions, making a total of 15 conditions.

The total percentages of possible /r/s which were articulated is reproduced in Figure 2, below. We can clearly determine two trends from this presentation. The first involves class. In every speech style, lower class speakers produced fewer [r]s than working class speakers, who themselves produced fewer [r]s than middle class speakers. This can be easily verified by simply comparing the values in each row. From left to right, the percentages of produced [r]s increase, as socioeconomic class increases. Second, in all classes, decreasingly

formal speech style correlates directly with decreasing production of [r]s. That is, in the figure, as we move down the chart in each column, we find a constant decrease in the production of [r].

Style	Class		
	LC	WC	MC
MP	49.5%	55%	70%
WL	23.5%	35%	55.5%
RP	14.5%	21%	29%
IS	10.5%	12.5%	25%
CS	2.5%	4%	12.5%

Figure 2 – Tabular representation of percent of /r/ production, as a function of socioeconomic class and speech style (Labov 1966)

Thus, these two factors, socioeconomic class and speech style, each contribute to the realization of /r/ in New York English. Moreover, for all intents and purposes, their contributions are autonomous. In other words, no matter what the social class of the speaker, increasing formality will on average yield an increase in [r]s. In the same vein, no matter the speech style, relatively higher ranked socioeconomic classes yield on average higher percentages of [r] production. This autonomy of the two variables stands in contrast to variables that interact.

To exemplify interaction between factors on phonology, we need look no further than another classic sociolinguistic study, this one conducted by Milroy and Milroy (1978). In their work, the Milroys studied a number of different variables, including a number of

vowels, and, most relevant to the point of interaction between variables, the voiced interdental fricative /ð/. In Belfast English, /ð/ can be elided intervocally. Words like *mother* and *brother* can be realized with or without a voiced interdental fricative in their center. The Milroys found this elision to have social correlates, particularly, the sex and age of the speaker. For present purposes, we will simply consider two age groups, 18-25 and 40-55, and two sexes, giving us a total of four conditions.

But Figure 3, which summarizes the average production percentage for each of the four conditions, shows a slight divergence from the autonomy of the factors in Figure 2. We first notice that, as in the case of New York /r/, there is a consistent difference between the two clusters. No matter what their age, males produce fewer [ð]s than do females. However, and this is the crux of the difference, it is not the case that gender is irrelevant to age differences. Rather, the effect of age depends on the gender of the speakers in question. Specifically, for males, increasing age correlates slightly with increasing use of [ð]. That is, older males produce slightly more [ð]s than do younger males. However, the reverse trend characterizes females, where increasing age in fact correlates with decreasing production of [ð]. In other words, older females produce significantly fewer [ð]s.

To know what effect a group's age has on /ð/ production, it is imperative to know their sex. Otherwise, the magnitude and even the directionality of the effect might be falsely predicted. Sex and age are thus said to interact in their influence on the dependent variable of /ð/ production.

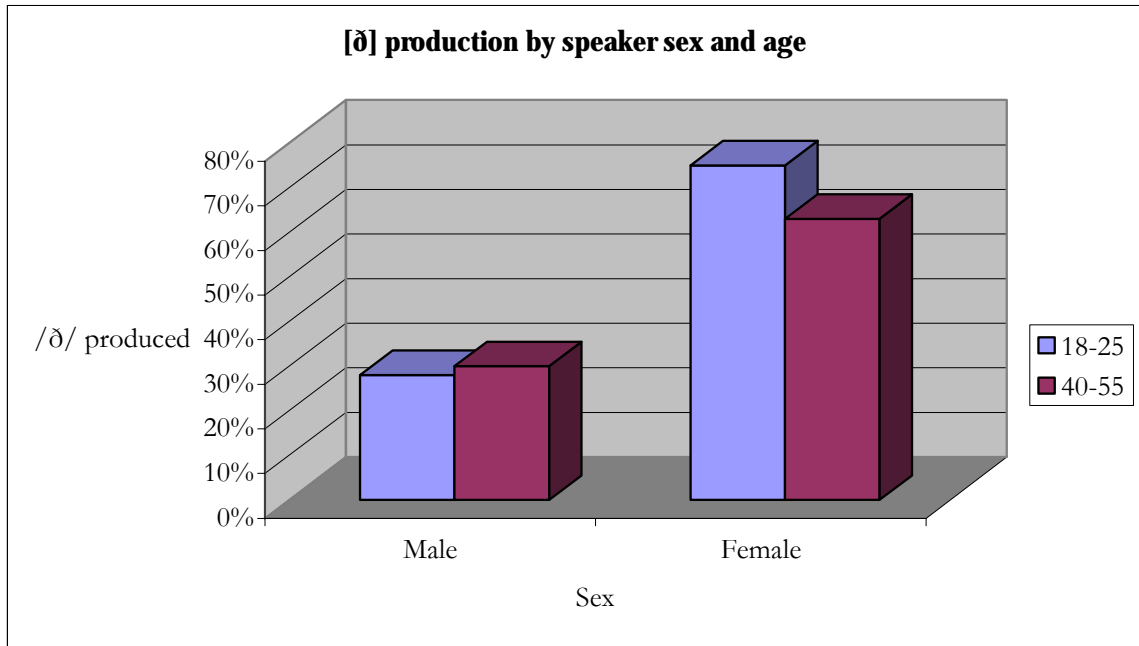


Figure 3 - /ð/ production as a function of speaker sex and age

Now that the distinction between interaction and autonomy of independent factors has been defined, we are in a position to ask what the relevance of this classification has. The question is significantly more intriguing when one considers that there is little if any discussion in the sociolinguistic literature of interacting factors (although Kay (1978) and De Jong (1989) are salient counter-examples). This question will be addressed in significantly more detail in Chapter 5, in which I discuss how one might go about trying to model interactions between factors influencing phonology. But briefly, as we will see in the next section, linear models of socially correlated variability, such as variable rule models based on VARBRUL, are unable to deal with interactions between variables, and thus in effect exclude them from consideration. This means first that models of such variation are required to be slightly modified (in no conceptually dramatic way) to include interaction effects. Second, when factors from different domains, such as social, semantic, syntactic and phonological knowledge, interact, they have dramatic ramifications for theories of the interfaces between

these domains of knowledge. In particular, this means that to make a prediction about the distribution of a particular variable, a number of different sorts of information must be brought to bear simultaneously in order to evaluate the effects of any of the factors. But let's not put the theoretical cart before the empirical horse. First we need to take a more detailed look at an extreme case of factor interaction on a large scale, the case of French liaison.

4. French liaison and the liaison corpus

French liaison is the variable production of a set of word-final consonants. Not all final consonants in French are liaison consonants - liaison consonants are lexically or morphologically specified. That is, liaison consonants can be the final consonants of words, like *jamais* 'never', or morphemes, like the plural *-s*. These liaison segments are produced or not, depending on number of factors. The most determining of these factors is the quality of the following segment, that is, the first segment of the following word. Liaison consonants are most likely to be produced when followed by a vowel, and are almost never produced when followed by a consonant, or utterance-finally (Tranel 1981). For example, (1) demonstrates the differential behavior of word final liaison consonant *-s* when followed by a vowel (1a) or a consonant (1b), or when it's utterance-final (1a and b). The notation used is to underline the relevant liaison consonants. A liaison consonant likely to be pronounced is marked with an uppercase letter, while one more likely to go unproduced takes a lowercase letter.

- (1) a. leS ananas 'the pineapples'
b. les bananes 'the bananas'

On the basis of data like those presented in (1), liaison is often classified as a type of external sandhi, or word-boundary juncture rule. External sandhi has been uncovered in numerous languages, most saliently in Classical Sanskrit (Whitney 1891). In Sanskrit, word-final segments are subject to assimilation with the first segment of the next word, as demonstrated by the examples of *tat* ‘that’ in (2), from Kessler (1994).

- (2) a. *tat indriyam* > *tad indriyam* ‘that sense’
 b. *tat manas* > *tan manah* ‘that mind’
 c. *tat cetas* > *tac cetah* ‘that intellect’

Just as the Sanskrit allophony in (2) appears to be conditioned by word-initial segments, so French liaison seems, given just the facts in (1) to be a simple, phonological word-boundary phenomenon. In fact, just like Sanskrit sandhi, French liaison appears to be phonologically natural. Sanskrit’s assimilation can be argued to serve an articulatory purpose. In the same way, a number of phonological frameworks, notably Optimality Theory (Prince and Smolensky 1993, Kager 1999), assume an iterated CV structure to be phonetically optimal. Liaison consonants are most often produced before word-initial vowels. When resyllabified with the following syllable, liaison consonants serve not only to reduce the coda of the preceding syllable, their syllable of origin, but also to add a consonant to the beginning of the following syllable. In both these ways, liaison production makes French syllables more frequently take a more CV-like structure.

But liaison, and perhaps also Sanskrit external sandhi (Rice 1990, Kessler 1994), is not as simple as it first appears. While liaison consonants followed by consonants are

exceedingly rarely produced, those which are followed by a vowel (as in (1a)) are less predictable. The realization of pre-vocalic liaison consonants is variable. If all we have to work with is the phonological nature of the following segment, the best we can do is to assign some probability to the realization of a liaison consonant. Before a vowel, as we will see in section 5 below, that probability is something like 0.5.

The following segment, though, is not the only factor that influences liaison. A number of other factors bear on the production of these consonants. Among these is the grammatical class of the liaison word. While nominal modifiers such as determiners and adjectives favor liaison (3a), most nouns disprefer it (2b).

- (3) a. *deS Anglais* ‘some English (people)’
 b. *dès anglais* ‘English dice’

The use of ‘prefer’ and ‘disprefer’ above is not accidental. Statements such as “liaison always occurs between a modifier and a noun it modifies” or “liaison is forbidden between a noun and a modifying adjective” while perhaps prescriptively useful, should be made with great caution, as the corpus study later in this chapter will show. The most accurate quantitative characterization we can make of the data is in terms of the probability that a liaison consonant will be produced given all its influencing factors that we can measure.

Before we look to the liaison literature for an indication of the extent and number of influences on liaison, an outstanding issue deserves treatment. For our present purposes, the production of liaison consonants is binary. That is, a liaison consonant can either be produced or not. Several caveats are required at this juncture. There are two ways in which liaison consonant production is in fact non-binary. The first regards the syllable assignment

of the consonant. When produced, liaison consonants are often, but not always, resyllabified such that they occur in the onset of the following word (Encrevé 1988). So just because we know that a liaison consonant was produced, that doesn't mean we also know what its syllabic placement was. Second, the quality of a liaison consonant can vary. Liaison consonants can in fact have one of six different identities: [z], [n], [t], [k], [•], and [p], and sometimes these identities are confused, both by aphasics (Kilani-Schoch 1983) and by normals (Leon 1984, Desrochers 1994, and Laroche-Bouvÿ 1994). As an example of identity confusion, the very common French *trop* 'too much' is occasionally heard with a liaison consonant other than the prescribed and orthographically suggested [p] liaison. The semantically primed [z] liaison, which occurs in the majority of plural nouns, adjectives, and determiners, and is exemplified in (3) above, is most likely to take its place. Whether this identity switching is due to reanalysis or speech errors, we can't predict with certainty which liaison consonant will be produced in a given context.

The first complication, that of resyllabification, will not be treated below, except indirectly, through the effect of pauses and punctuation on liaison. This is done partly to make the present task manageable, and partly because Encrevé's results show resyllabification to be more or less orthogonal to liaison. Resyllabification is a process that occurs throughout the language. Thus, although it is essential to some purely phonological accounts of liaison, such as in an OT framework, resyllabification can hopefully be methodologically separated from the process of liaison.

The second complication, the issue of liaison segment identity, raises a host of frustratingly central questions, surrounding the morphological and semantic status of liaison. Although these questions will be addressed below, the present study yielded very little identity variation, only two tokens in the corpus were instances of "mistaken" identity:

reanalysis or speech errors. This was likely due to the fact that the corpus I discuss below is constructed from read speech, where orthographic indicators can serve to override natural linguistic tendencies. Therefore, since the present work contributes no new empirical evidence on quality variation, identity will only be referred to tangentially as evidence for the semantic content of certain liaison consonants.

Proposed influences on liaison

The literature on French liaison is extremely large and fragmented. One might in fact prefer to say that there are two or even three literatures on liaison: an older one, focusing on the language of the individual, and which is broken into a phonological and a syntactic branch, and another, more recent one, targeting the language of the community.

Studies of the first sort ask questions about the representational status of liaison. Is a liaison consonant a unit with morphological status or is simply a phonological one? Is it added to a bare stem in contexts in which it appears, or deleted from full forms in those contexts in which it does not occur? Does its occurrence have a functional explanation, or is the occurrence of a liaison consonant a purely formal aspect of Francophone linguistic knowledge? If liaison is a formal operation, to what domain does it owe its distribution: phonology or syntax? What are the syntactic configurations which respectively allow, require, and disallow liaison?

Two lines of generative research attempt to explain liaison in different ways. The phonological view sees liaison as a phonological process, the insertion or deletion of a consonant, where the conditioning factors are purely phonological. It tries to explain why a liaison consonant would be produced or not on the basis of purportedly universal or

functional criteria. Studies of this type are generally unconcerned with the details of whether other factors cross-cut phonological ones. As such, they are strictly limited in their explanatory scope.

The alternative generative view, the syntactic one, follows the traditional prescriptive view that liaison contexts can be strictly categorized by their effect on liaison. On this view, there are three types of liaison contexts (Delattre 1966).

1. Contexts in which liaison is *obligatory*, such as, for example, between a plural determiner and a noun it modifies.
2. Contexts in which liaison is *prohibited*, such as before an h-aspiré word
3. Contexts in which liaison is *optional*, such as between a verb and whatever follows it

Generative syntactic models differ in what they see as the defining properties of these contexts. The traditional view, that liaison occurs within a ‘rhythmical unit’ (Grammont 1938) yielded in the 1970s to a syntactic definition of its context. After an attempt to define the domain of liaison in purely syntactic terms (Selkirk 1974), Selkirk moved on to develop an extremely influential syntactic/prosodic account (Selkirk 1978). On this view, the syntactic structure of an utterance is mapped to a prosodic structure, which is composed of prosodic categories like the syllable, foot, and prosodic word, and which is then subject to prosodic rules. Selkirk (1978) argues that the phonological phrase is the domain of obligatory liaison. What’s a phonological phrase? As defined by Nespors and Vogel (1982:228-229), it is as described in (4).

- (4) Join into a phonological phrase a lexical head (X) with all the items on its

nonrecursive side within the maximal projection, and with any other nonlexical item on the same side (such as prepositions, complementizers, conjunctions, copulas).

More recently, Selkirk (1984) has proposed that it is not the phonological phrase, but rather silent *demibeats* that determine the obligatoriness of liaison. Demibeats are positions on the metrical grid of a sentence, and the number of silent demibeats between any two words is determined by a silent demibeat addition rule. The details of this rule are unimportant for our present purposes. What is important is that the metrical structure of a sentence is believed in this theory to be derived directly from its syntactic structure.

Moving back to the pure syntax of liaison contexts, Kaisse (1985) presents an alternative view to those of Selkirk, and focuses only on liaison in informal speech. Here, Kaisse presents liaison as an external sandhi rule, and argues that in order for liaison to be obligatory, the second word must c-command the first. Again, this is a purely syntactic definition of obligatory liaison.

Despite the differences in what these models see as essential to obligatory or prohibited liaison, they all agree that the main focus of studies of liaison ought to be on defining necessary and sufficient conditions to distinguish the obligatory cases (and sometimes the prohibited ones) from the others. And the only criteria considered are syntactic ones. In general, the detail of what factors weigh on the distribution of liaison consonants in optional contexts is of no interest. Of equally little interest is empirical validation that the three-way schism is actually valid, or, if it is, whether tokens are distributed through it as prescribed. Finally, models such as these have little if anything to say about what happens with the domain of optional liaison, where any syntactic rules that might have any validity would by virtue of the optionality of the context, apply non-

categorically.

In fact, in spoken French, liaison is obligatory in only a small handful of contexts. Encrevé (1983) and Booij and De Jong (1987) both conducted large corpus studies, in which they investigated the contexts in which, and degree to which, liaison was actually used as though obligatory in spoken French. The studies both indicate that liaison is obligatory only in a restricted subset of those syntactic contexts described in the prescriptive and generative literature. Only after determiners and between a personal pronoun and a verb or between a verb and a pronoun are liaisons always produced. An additional group of obligatory liaisons occur inside of fixed expressions, such as *de temps-en temps* ‘from time to time’, *états-unis* ‘United States’, and *pyrenées-orientales* ‘eastern pyrenees’ (Bianchi 1981).

By comparison, contexts where liaison production is optional are plentiful. Moreover, rather than simply being the product of a randomly applied optional rule, liaison use in optional contexts is sensitive to number of linguistic and extra-linguistic factors. The restrictedness of obligatory liaison, paired with the systematicity of influences on optional liaison suggest that obligatory liaison simply falls at one end of a continuum of liaison use, rather than constituting an entirely separate entity.

By contrast with individualist, generativist investigations, studies addressing the behavior of liaison as it is produced by a linguistic community focus unsurprisingly on the social correlates of liaison production. That is, they ask questions like the following. Are there effects of gender, age, and socioeconomic status on the use of liaison in a community? What is the directionality of these effects? For example, are men or women more likely to use liaison? What is the magnitude of these social effects?

Studies in this line of work have also focused some attention on the linguistic correlates of liaison use in optional liaison environments. This perspective has yielded

statistical measures of effects from these two domains. The first includes purely linguistic or semantic constraints, constraints purportedly at work in the individual during language use. The second includes purely social constraints, which are claimed to structure the intra-cultural differences in the use of liaison. We will shortly see a cursory survey of a large number of the factors of each of these classes, starting with linguistic factors and ending with social factors. Along with this survey will be descriptions of how a large corpus, which included all these factors, was constructed. First, though, let's briefly look at the global architecture of the French Liaison Corpus.

Construction of the French Liaison Corpus

Our overarching goal in studying French Liaison is to investigate the extent to which individuals have internalized non-categorical correlations between phonology and social aspects of the speaker. The perception experiment described in Chapter 4 tests this knowledge. But in order to evaluate the internalization of probabilistic patterns, we need first to establish what those patterns are in the language of the community – the ambient language from which correlations are extracted. To study variation in a language community, in particular with the exceedingly large number of factors identified in the literature (and enumerated later in this section), a large, well-balanced, and controlled corpus is of the utmost importance (Biber et al. 1998).

To this end, I based my study on the IDIAP French Polyphone corpus (Chollet et al. 1996). The French Polyphone corpus consists of speech from several thousand French-speaking inhabitants of Switzerland. The speech is of two types. The majority of the corpus is composed of read speech, intended to maximize the range of phonological contexts

recorded. A much smaller portion consists of spontaneous human to computer speech. All recordings were made over the telephone, meaning that their resolution is somewhat reduced and that there is often background noise. All data is in two forms – the acoustic signal and a human transcription.

I chose to work exclusively on the read speech, for two reasons. First, significantly more data per subject was available in that form. Second, since linguistic interactions between humans and computers remain extremely infrequent, natural human-to-computer speech is extremely variable in its register and style. Including this speech would have added additional factors to my analysis, factors which, it will be shown below, play a role in liaison. Since it would be very difficult to evaluate which style or register a particular speaker chose spontaneously to use in a non-circular manner, I chose to restrict analysis of the corpus to read speech.

Tagging

I randomly selected 200 speakers from this corpus, of whom half were men and half women. Of the 200 speakers I selected, I pruned away those who self-reported as native speakers of a language other than French, since L2 French speakers display quite different liaison behavior from natives. This segregation left me with a total of 173 speakers. Of these speakers, 90 were male and 83 female.

I found all tokens of liaison consonants. Some of this was automated - a small set of frequent words, such as *les* ‘the’ and *dans* ‘in’ was automatically identified. But not all final consonants are liaison consonants. For example, while the final “s” in *films* ‘strings’ is a liaison consonant, pronounceable as either /z/ or \emptyset , the terminal “s” in the homograph *fils* ‘son’ is

always pronounced, as /s/. Because of cases like this, much of the tagging had to be done manually.

Once I had established all the liaison consonants in the corpus, I turned to influencing factors on liaison. First, although the literature claims strongly that liaison production before a consonant is strictly impossible, I wanted to make sure not to exclude any potential variation. To this end, I marked up 5% of the liaison cases for whether they were followed by a consonant or a vowel, and considered all the consonant cases. In this sample, of 250 liaison consonants preceding consonants, exactly none were produced. So for the rest of my analysis, I considered only liaison consonants preceding vowels.

On the basis of the literature discussed in section 4, above, each potential liaison was analyzed along 19 different dimensions. Some of these factors were coded automatically, others completely by hand.

What follows is a description of all of the factors pulled from the liaison literature that were included in the corpus, along with a characterization of how each factor was coded. A synopsis can be found in Figure 7.

Phonological identity

As mentioned above, there are six possible identities of liaison consonants. Three studies have examined the effect of liaison consonants' identity on their production, with dramatically different results.

Ashby's (1981) study, described below, ranked liaison consonants by their effect on liaison. He found [n] to most strongly favor liaison, followed by [t] and [z]. He did not have large enough numbers of other segments to make conclusions about them. But Ashby's ranking is contradicted by that of Encrevé (1983), who, in a large corpus composed of politicians' radio and television performances, found [t] to be the most liaison-inducing consonant, followed by [z] and [•], with [n] being word-specifically obligatory or prohibited. Finally, Malecot (1979) found in a large corpus of educated, middle-class Parisian speech that [•] was more likely to be produced than [z], which was itself more often articulated than [t]. Unfortunately, Malecot had too few tokens of [•] for their production data to be reliable. These authors' divergent results are summarized below in Figure 4.

[n]	[•]	[t]	[z]	Source
61%	NA	37%	28%	Ashby 1981
NA	11%	72%	40%	Encrevé 1983
NA	94%	52%	61%	Malecot 1979

Figure 4 – Liaison frequency as a product of liaison consonant identity

1. Liaison orthography

The orthographic realization of the liaison segment was extracted from the French Liaison Corpus automatically, using a Perl script (Schwartz and Christiansen 1998). It was encoded directly as a, c, d, e, g, l, n, p, r, s, t, u, x, or z. The reason vowels were included in this list was to mark “morphological” *t* insertion between verbs and following clitic subject pronouns. You will likely have noticed the seemingly aberrant presence of l as a liaison consonant. Although it is not considered a liaison consonant, I wanted to test the hypothesis that word-final *l* is also subject to liaison-like behavior.

2. Liaison phonology

Similarly, the phonological realization (or potential realization) of each liaison consonant was derived directly from its orthography, which happens to be possible because liaison consonants have consistent phone-grapheme mappings. I also listened to every token and found only three cases where the predicted phonological realization of a segment and its actual realization differed. These were not included in further statistical analyses. The classes which were included were k, l, n, p, r, t, and z.

Preceding segment

It's not the nature of the liaison consonant that bears a phonological influence on liaison.

The preceding segment also plays a role. And just like other cases of final consonant variability (such as t/d-deletion), a preceding vowel raises the likelihood that the liaison consonant will be produced, while a preceding consonant lowers it (Ashby 1981, Morin and

Kaye 1982).

It has also been suggested that when preceded by two consonants, rather than just one (that is, when the liaison word has the form CVCCC), a liaison consonant is even less likely to be produced (Delattre 1966). For example, words like *armen* [a• m̃] ‘arm (pl. 3rd pres.)’ and *parten* [pa• t̃] ‘leave (pl. 3rd pres.)’ should drop liaison consonants more frequently than words like *veulen* [vœl̃] ‘want (pl. 3rd pres.)’ and *coulent* [kult̃] ‘flow (pl. 3rd pres.)’. But this claim has been contested by Kovac (1979) and Agren (1973).

Finally, when a semivowel precedes a liaison consonant, its effect is ambiguous. Ashby (1981) reports that although liaison is more frequent after semivowels than after consonants, his statistical analysis shows that given all the other factors weighing on liaison production, semivowels actually have the reverse effect. They decrease the likelihood of liaison, relative to consonants.

3. Preceding segment

The preceding segment had to be coded manually, since outside liaison, French is just like in English, in that it’s quite a challenge to determine a French word’s phonology from its spelling. For example, in *parten* [pa• t̃] ‘leave (pl. 3rd pres.)’, the ‘en’ is silent, while in *men* [mã] ‘lies (v)’, it isn’t. Each segment preceding a liaison consonant was manually classified as a vowel, semivowel, or consonant.

Following segment

It's generally believed that liaison preceding consonants is always disallowed, while it may be permitted before vowels. But what of semivowels, the mutts of the CV world? Words starting with semivowels abound in French, and this class includes those words listed in Figure 5 below (from Watbled 1991). Since many are recent borrowings into French, these might display aberrant behaviors with respect to liaison.

/j/	<i>hiérarchie</i> 'hierarchy', <i>hiéroglyphe</i> 'hieroglyphe', <i>hyène</i> 'hyena', <i>iode</i> 'iodine', <i>iodé</i> 'iodized', <i>yacht</i> 'yacht', <i>yaourt</i> 'yogurt', <i>yeux</i> 'eyes', <i>yoga</i> 'yoga', <i>yougoslave</i> 'yougoslavian', <i>yo-yo</i> 'yo-yo'
/● /	<i>huile</i> 'oil', <i>huiler</i> 'oil (v)', <i>huissier</i> 'usher', <i>huit</i> 'eight', <i>huitième</i> 'eighth', <i>huître</i> 'oyster'
/w/	<i>oie</i> 'goose', <i>oiseau</i> 'bird', <i>oisif</i> 'inactive', <i>oisiveté</i> 'inactivity', <i>ouate</i> 'cotton wool', <i>ouest</i> 'west', <i>oui</i> 'yes', <i>ouistiti</i> 'marmoset', <i>watt</i> 'watt', <i>week-end</i> 'week-end', <i>western</i> 'western (movie)', <i>whisky</i> 'whisky'

Figure 5 – Some French semivowel-initial words

In fact, this question has been little-studied. While statistically unsubstantiated proposals abound, such as Watbled's (1991) suggestion that "foreign" words block liaison, there are to my knowledge no large scale studies that have investigated the effects of semivowels on liaison.

In addition to the complicated behavior of semivowels, there is one major exception to the generalization that vowel-initial words can allow liaison in a preceding word. The trouble is caused by a class of vowel-initial words, most of which are spelled with an

unpronounced initial ‘h’: words known as *h-aspiré*, or ‘aspirated-h’, words. (As to the name for these words, many of them had initial [h] before being borrowed into French, although there is little evidence that they were ever actually produced with an initial aspirate in French.) As shown in (5c), h-aspiré words like *hiboux* ‘owls’ defy liaison, unlike other vowel-initial words (5b), but just like consonant-initial words (5c).

- (5) *leS autres* ‘the others’
 les hiboux ‘the owls’
 les fermiers ‘the farmers’

H-aspiré words are plentiful in French - among them are those listed in (6) below. Just about every aspect of h-aspiré seems to vary across speakers: which words belong to this class, whether they disallow liaison in all contexts, and if not, what the more permissive contexts are (Good 1998). What is consistent, though, is that h-aspiré words tend to allow less liaison than other vowel-initial words.

- (6) *hache* ‘axe’, *haie* ‘hedge’, *haine* ‘hatred’, *halte* ‘stop’, *hameau* ‘hamlet’, *hamster* ‘hamster’,
 hanche ‘hip’, *handicap* ‘handicap’, *hangar* ‘hangar’, *hareng* ‘herring’, *haricot* ‘green bean’,
 hasard ‘chance’, *hausse* ‘raise’, *haut* ‘high’, *héros* ‘hero’, *hibou* ‘owl’, *Hollande* ‘Holland’,
 homard ‘lobster’, *honte* ‘shame’, *hurler* ‘howl’

4. Next orthography

The first segment of the word following a liaison consonant can be either a vowel or a semivowel. But among the vowel-initial words that block liaison is the class of h-

aspiré words. Since the list of words with h-aspiré differs from speaker to speaker, I decided to simply code all h-initial words as having a potential h-aspiré, rather than decide on the basis of prescriptive judgments which words belonged to this group for each speaker. For the purpose of the analysis of *l*-final words, I also included the possibility that a following word could begin with a consonant, to establish the total frequency with which *l*s are produced. The following segment was automatically extracted from the database and the grouped into one of the four possible values: vowel, semivowel, h, and consonant.

Pauses

The notion that liaison occurs within a prosodically defined unit is an old one (e.g. Delattre 1966), and there is some quantitative substantiation of this claim. Agren (1973) found liaison to be significantly less prevalent when a pause separated the liaison word and following word. As few as 5% of liaison consonants were produced before pauses in his corpus.

5. Pause

Duez (1985) has shown that pauses in French of 180-250 msec across contexts are usually not interpreted as pauses. So the lower cutoff for pauses was set at 200 msec. Just in case the length of pause also contributed to liaison, I also incorporated the length of the pause into the coding for this feature, which had four possible values, no pause (or a pause of less than 200msec), short pause (200-400msec), medium pause (400-600msec) and long pause (greater than 600msec). I manually measured the delay between the offset of the liaison word and the onset of the following word,

using ESPS Waves. There is also a possibility that “pause” could be expressed through the prolongation of a syllable in French, but I could not detect any such lengthening through simple observation.

6. Punctuation

Since the speaker in the corpus were reading text, punctuation may have also played a role in determining where pauses were placed. I therefore automated the extraction of cases where punctuation intervened between the two words, as the texts subjects were to read were also provided. These were subsequently segmented into minor segmentations (,) and major ones (;, :, and .). We’ll see in the statistical analysis in Section 5 below that pause and punctuation overlap to some degree, but not entirely.

Morphological status

In studies of t/d deletion in English (see Bybee 2000 for an overview), it is well established that when a final coronal stop bears some morphological status, it is subject to less deletion than when it bears none. While the same hypothesis seems not to have been forwarded generally for the presence of liaison, there is nevertheless some evidence that meaningfulness has an impact on liaison. For example, in adjective + noun sequences, when the adjective’s final segment expresses a plural in /z/, liaison is more frequent than when the noun is singular (Morin and Kaye 1982). This is tantamount to saying that the plural morpheme surpasses non-expressive final consonants in their tendency to liaise.

Aside from adjectival number, liaison consonants can signify verbal inflection, as in (7), or nominal number, as in (8).

- (7) a. *parT avec* ‘(he/she) leaves with’
 b. *pars avec* ‘(I, you) leave with’
- (8) a. *eleveS anciens* ‘former students’
 b. *eleve ancien* ‘former student’

French verbs are inflected for number, person, and tense with suffixes. The final consonants of these suffixes are all liaison consonants, and can be ‘s’ [z], ‘z’ [z], ‘x’ [z], ‘t’ [t], or ‘d’ [t]. Suffixes that end in [t] usually mark the third person and occasionally the second person singular, while those ending in [z] usually mark the first or second person.

Any effect found due to morphological status in the adjectival case cited by Morin and Kaye (1982) or in the cases in (7) and (8) could be explained by morpheme-specific omission frequencies, or could be due to functional pressure. Indeed, meaning-bearing liaison units serve any of a number of functions, outlined by Garrote (1994), following Malmberg (1969), and summarized in Figure 6. They can serve the purpose of lexical opposition, distinguishing words that would otherwise be homophonous. they can serve a morphological opposition function, distinguishing between morphological variants of a lexeme. Finally, they can serve a syntactic opposition function, distinguishing between different syntactic arrangements.

Opposition type	Examples of near-homonyms
1 Lexical opposition	<i>les hauteurs</i> ‘the heights (an h-aspiré word)’ vs. <i>les auteurs</i> ‘the authors’
2 Morphological opposition	<i>Il était la</i> ‘he was there’ vs. <i>ils étaient la</i> ‘they were there’
3 Syntactic opposition	<i>un savant anglais</i> ‘an English savant’ vs. <i>un savant Anglais</i> ‘a knowing Englishman’

Figure 6 - Types of functional oppositions served by liaison consonants (from Garrote 1994)

Looking more specifically now at liaison consonants’ possible manifestations, we see that the expression of person influences liaison asymmetrically. Morin and Kaye (1982) found more liaison with verbal 3rd person [t] than with 1st person [z]. The relation between this effect and overall frequency of these consonants remains, however, unclear.

To these observations that number and person can influence liaison use, Martinet (1988) adds anecdotal evidence that their effect is at least partly functional. She suggests that when the meaning borne by a semantically charged liaison segment is expressed elsewhere in the sentence, this decreases the probability of the liaison’s expression.

7. Plural

In manually coding tokens for their expression of number, I selected all those final /z/s and /t/s which had some morphological status and placed them in a single category. But a small set of words (80 tokens in the corpus) ends in one of these two liaison consonants yet do not include, on most linguistic analyses, a plural

morpheme. Examples of such pseudo-plurals occur in words like *deux* ‘two’ and *trois* ‘three’. These were placed into a separate class. Along with no plural and regular plural, this yielded a total of three values for this factor.

8. Person

Liaison consonants can mark any one of three persons on French verbs, as seen in (7) above. Tokens were manually classified as bearing no person information, or instantiating the first, second, and third person.

Syntax

We’ve already seen the important role that syntax plays in determining whether a liaison is obligatory, if such a classification is valid. According to Booij and De Jong (1987), liaison is obligatory: (1) after determiners, (2) between a personal pronoun and a verb, or (3) between a verb and a pronoun. But a whole lot of other syntax effects can be felt as non-categorical influences on liaison use.

Morin and Kaye (1982), for example, show that between an inflected head and its complement, liaison is significantly more frequent when that head is a verb than when it’s a noun. So even in a restricted syntactic environment, grammatical class effects liaison use. By the same token Booij and De Jong’s corpus and Ashby’s corpus demonstrate that after clitics and determiners, liaison is extremely frequent, but not so after other modifiers.

9. Liaison word class

The grammatical class of the liaison word was determined by means of an automated tagging program, TreeTagger, described in Schmid (1994). This tagger consistently exceeds 95% accuracy in matching the hand-labelled data in the Penn Treebank. Although the tagger provides more detail as to the grammatical intricacies of the words in the corpus, including verbal mode and the like, I only included grammatical class in a broader sense. Values for this variable were adjective, adverb, conjunction, determiner, interrogative, noun, proper noun, preposition, pronoun, verb, and abbreviation. Of these, conjunction and interrogative were not numerous to provide significant results and were discarded.

10. Next word class

The TreeTagger program automatically assigned grammatical category classifications to the words following liaison consonants. These ranged over conjunction, determiner, interrogative, noun, proper noun, preposition, pronoun, verb, and abbreviation.

Frequency

Lexical frequency is a key determinant in a number of phonological processes, including t/d deletion in English (see the summary in Bybee 2000). In t/d deletion, more frequent words are more commonly subject to deletion. Bybee explains this in terms of the progressive automation of the articulatory gestures responsible for the coronals, arguing that the more frequent a word is, the more it is to be eroded through routinization, and the more

entrenched the eroded version of the word will become. Another plausible explanation could take the informativeness of final stops into account. The more frequent a word, the higher a listener's expectation that that word will be produced. In other words, more frequent words might need less phonetic detail to be identified, and speakers might thus take advantage of listener's expectations, by producing them less clearly.

But with liaison, the reverse trend seems to emerge, although indirectly. Agren (1973) demonstrated that before more frequent words, liaison consonants are more likely to be produced, and Booij and De Jong (1988) showed that frequent word-combinations also led to greater liaison use. Confounding the issue of frequency is the existence of a number of fixed expressions which unequivocally require liaison, expressions like *cas échéant*, *fait accompli*, and *accent aigu* (from Malecot 1979), some of which are frequent, others of which are less so. These frequency results leave open the question of what effect if any a frequent liaison word will have on liaison production, although they indicate a role for word-pair frequency and following word frequency.

11. Liaison word frequency

In order to determine the frequencies of the several thousand words in the French Liaison Corpus, I needed an even larger corpus of French, from which to extract lexical frequencies. In practice, there is only one large, publicly accessible corpus of French, available from the Project for American and French Research on the Treasury of the French Language, or ARTFL (<http://humanities.uchicago.edu/ARTFL.html>). The ARTFL corpus is composed entirely of written texts, almost 2000 of them at present. There is a total of 115 million words, comprised of around 439,000 word-forms. I again automated this

process using a Perl script which produced the number of total occurrences of each liaison word form in the entire ARTFL corpus. These ranged from 0 for *Azerbaijan* to 1,903,168 for *les* ‘the (pl.)’. These were compared on a token-by-token basis, without reference to lexemes, which is appropriate since ARTFL does not include lemma representations. The tokens turned out to be more or less equally distributed among five frequency classes, 0-10³, 10³-10⁴, 10⁴-10⁵, 10⁵-10⁶, and greater than 10⁶. This distribution is indicative of a decrease in token number as frequency increases – there are many fewer very frequent words (10⁶) than fairly frequent words (10⁵-10⁶), and so on.

12. Next word frequency

The following word’s frequency was coded in the same manner as the liaison word’s frequency, through a search of the ARTFL corpus and subsequent segmentation into five categories: 0-10³, 10³-10⁴, 10⁴-10⁵, 10⁵-10⁶, and greater than 10⁶. Word-pair frequency proved intractable to extract from the ARTFL database due to the coding therein, and was therefore not included. We will see, however, the combined effects of the frequency of the liaison and next word in Section 5, below.

Word length

Word length, and relative word length in particular, cross-cut all of the factors we have already encountered. First, it has been widely observed (by Encreve 1988 and others) that greater liaison numbers arise where a liaison word is monosyllabic than where it is polysyllabic. In the same vein, Ashby’s (1981) study demonstrated an effect of relatively

shorter liaison word length on liaison production. He showed that in cases where the liaison word and following word were the same length or where the following word was longer than the liaison word, liaison was more likely than when the liaison word was longer than the following word.

13. Liaison word length

The number of letters in the liaison word was automatically counted and the results were grouped into four classes: 1-3 characters, 4-6 characters, 7-9 characters, and 10- characters.

14. Next word length

The following word's length was coded in the same fashion as was that of the liaison word. Possible values were again 1-3 characters, 4-6 characters, 7-9 characters, and 10- characters.

15 and 16. Liaison word syllables and next word syllables

Since orthographic length is a coarse metric, I also manually counted the number of syllables in both the liaison word and the following word, all of which had between one and six syllables.

Social factors

The sociolinguistic literature identifies five social variables that correlate with liaison production in the language of the community. These are speaker age, socioeconomic status,

and sex, as well the speaker's dialect and speech style.

A number of authors agree that a speaker's age and socioeconomic status have an impact on how speakers use liaison. Two corpus studies have addressed the question of the social correlates of liaison, but they differed in their data collection methods. Booij and De Jong (1987) based their analysis on an unpublished manuscript co-authored by De Jong (De Jong et al. 1981), in which 38 native French speakers from Tours, France were asked to read aloud a text containing a number of different liaison contexts. All of these were so-called obligatory liaison contexts. That is, these were liaisons that were assumed to be produced in all cases. Ashby (1981), on the other hand, studied the use of liaison by 16 informants, also from Tours, in an interview situation. His data came from interview speech, rather than read speech. From these interviews, Ashby collected only those liaison contexts in which liaison was applied variably. He ignored contexts in which liaison was either always or never applied. Although their data differed in origin, Ashby 1981 and Booij and De Jong 1987 both found that with increasing age comes a greater likelihood that a speaker will produce liaison consonants. They didn't, however, allow us to determine if the language is changing, or if these effects are purely due the chronological age of speakers. The same two corpus studies uncovered a correlation between higher socioeconomic status and increased use of liaison consonants.

Additional confirmation of the role of age in the use of liaison comes from Morel's (1994) experimental study of differences between adult's and children's use of liaison. She observed that children of kindergarten age tend to retain liaison consonants heard in one context when they produce the same word in another context more than adults do. This lexical consistency could be a reliable correlate of extremely young age of the speaker.

However, when it comes to the role of sex in liaison use, there is not yet any clear

consensus. Booij and De Jong found women to be consistently more likely to articulate a liaison consonant. From a sociolinguistic perspective, this is an entirely unsurprising state of affairs. In variable segment after variable segment, women are found to be more linguistically conservative than men are (Trudgill 1983). This trend is often explained in terms of women's ostensibly superior social awareness (Trudgill 1972) or linguistic facility (Chambers 1995). However, in Ashby's study which only measures liaison in variable contexts, men tended to produce more liaison consonants than their age- and status-matched female peers. So although these two studies converge on a common conclusion – that sex matters to liaison – they diverge in exactly how it does.

Speech style, too, seems to condition liaison use. De Jong et al. (1981) elicited data of different levels of formality through two different tasks: a word list task (very formal) and a passage reading task (formal). In case after case, they found liaison production to be much higher in the word-list condition than in the reading condition, indicating that the more attention speakers pay to their speech, the more likely they are to articulate their liaison consonants. In the same vein, Ashby separated his interviews into two halves, and measured the degree of liaison use in each. His statistical analysis showed that early interview placement tended to push liaison consonants to be produced, while later interview placement led to more liaison omission. Once again, it seems that as speakers begin to pay less attention to their speech, that is, as their speech becomes more informal, liaison use drops off.

A final social factor that influences liaison use is dialect. I call this a social factor because it belongs not to the realm of the linguistic structure proper, but is part of the social identity of the speaker. The world's 13th most spoken language, French is spoken widely in a dozen countries. Aside from the 50 million Francophones in France, large populations are

also to be found in Canada (6 million), Switzerland (1.3 million), and Belgium (4 million), not to mention those speakers throughout former French and Belgian colonies throughout Africa and East Asia (Grimes 2000). Each geographical locus of French has particularities that are proper to it, and some of these touch liaison. For example, liaison interacts with the regular Canadian French processes of assibilation and vowel lowering (Walker 1980).

Even within France, there is substantial regional dialectal variation. The French of Brittany or Burgundy can be easily distinguished from that of the Parisian region. In terms of liaison, Marseille in particular displays some peculiarities. A number of h-initial words that are reported to be h-aspiré words in Standard French, and are used as such across France, attract liaison in Marseille, rather than rejecting it. Watbled (1991) lists a number of such words, including *hamster* ‘hamster’, *handicap* ‘handicap’, *haricot* ‘green bean’, *hasard* ‘chance’, *hauteur* ‘height’, *héros* ‘hero’, *Hollande* ‘Holland’, and *hurler* ‘howl’.

17. Sex

One of the questions each informant was asked to respond to was their sex. Only two speakers did not answer this question. Both had unambiguously male voices, and were therefore included in the database as male speakers.

18. Age

Informants provided the year of their birth in responding to the questionnaire. Recordings were performed in 1995-6, and subjects fell more or less evenly into three groups corresponding to rough age classes. Young speakers, those 25 years old or younger, were born after 1970. Older speakers, those over 45 years of age, were born before 1950. All others were classified as middle aged.

19. Education

The educational system in French-speaking Switzerland, much like elsewhere in Europe, provides an early academic track bifurcation. High school-aged students can find themselves in a regular high school, in which case they usually anticipate some sort of tertiary academic schooling, or in a technical high school, where more emphasis is placed on technical skills to be applied in a profession requiring such abilities. Informants were asked to specify the highest level of education they had obtained. Those who reported no secondary education were grouped into the primary class. Those who stated they had obtained some post-secondary schooling were grouped into the tertiary group. All other responses were classified as secondary education.

Summary

We have run through a long list of potential probabilistic influences on liaison (summarized in Figure 7), including some interactions between factors. All of this has been done in the broader interest of answering the questions: What does the individual know about the correlates of variation when perceiving language?

What remains can be separated into two main tasks. First, we need to document variation in the production of the community, a goal which is addressed through the statistical analysis of these factors in the French Liaison Corpus, described in the next two sections. We then need to measure the extent of individual use of this variation in perception. This second task will be tackled through two perception experiments, reported on in Chapter 4.

#	Feature	Description	Values
Dep	Valence	Whether the consonant was produced	yes, no
1.	Liaison orthography	Orthography of liaison segment	a, c, d, e, g, l, n, p, r, s, t, u, x, z
2.	Liaison phonology	Phonological realization of liaison	k, l, n, p, r, t, z
3.	Preceding Segment	Type of the preceding segment	consonant, vowel, semivowel
4.	Next orthography	Orthography of the following segment	vowel, semivowel, h, consonant
5.	Pause	Length of pause between the words	none, short (0-200msec), medium (200-400msec), long (>400msec)
6.	Punctuation	Punctuation between two words	none, ‘,’ ‘,’ ‘,’ ‘,’
7.	Plural	Plural marking of the liaison	none, regular plural marking, interpretable as plural marking
8.	Person	The person expressed by the liaison	none, 1st, 2nd, 3rd
9.	Liaison Word Class	Grammatical class of the liaison word	adj, adv, con, det, int, nom, npr, pre, pro, ver, abr
10.	Next Word Class	Grammatical class of the following word	Adj, adv, con, det, int, nom, npr, pre, pro, ver, abr

11.	Liaison Word Frequency	Tokens of the liaison word in ARTFL	0-10 ³ , 10 ³ -10 ⁴ , 10 ⁴ -10 ⁵ , 10 ⁵ -10 ⁶ , 10 ⁶ -
12.	Next Word Frequency	Tokens of the next word in ARTFL	0-10 ³ , 10 ³ -10 ⁴ , 10 ⁴ -10 ⁵ , 10 ⁵ -10 ⁶ , 10 ⁶ -
13.	Liaison word length	Orthographic length of the liaison word	1-3 characters, 4-6 characters, 7-9 characters, 10- characters
14.	Next word length	Orthographic length of the next word	1-3 characters, 4-6 characters, 7-9 characters, 10- characters
15.	Liaison word syllables	Number of syllables in liaison word	1-6
16.	Next word syllables	Number of syllables in following word	1-6
17.	Sex	Gender of the speaker	female, male
18.	Age	Speaker's year of birth	1924-1949, 1950-1969, 1970-
19.	Education	Educational class of the speaker	no info, lower (primary), middle (secondary), higher (tertiary)

Figure 7 – Features tagged in the French Liaison Corpus

5. A test of autonomous factors

In this section, I'll address the role of independent factors on liaison in the language of the community, from a statistical perspective. The main question to be answered is the following:

- Which factors enumerated above are confirmed in this corpus to have statistical significance?

Analysis

In the past twenty years, linguistic variables have been principally studied using an analytical tool known variably as VARBRUL (Sankoff 1987) or GOLDVARB (Rand and Sankoff ms.). These computer programs implement, among other data collection and preparation tools, a statistical method for analyzing the significance and degree of an arbitrarily large number of factors on a single linguistic variable.

The statistical method in question is the construction of a particular type of regression model. A regression model consists of an equation, which is a mathematical description of the effects of certain independent factors on a dependent variable. In the simplest case, there is only one influencing and one influenced factor. Let's take the example of the relation between the length of a liaison word and the production of a liaison consonant. As the graph in Figure 8 shows, there is a direct relationship between the frequency class of the liaison word and the probability that liaison consonants (Cs) will be produced. (If tokens were no grouped into these classes, the curve would have a flatter slope.)

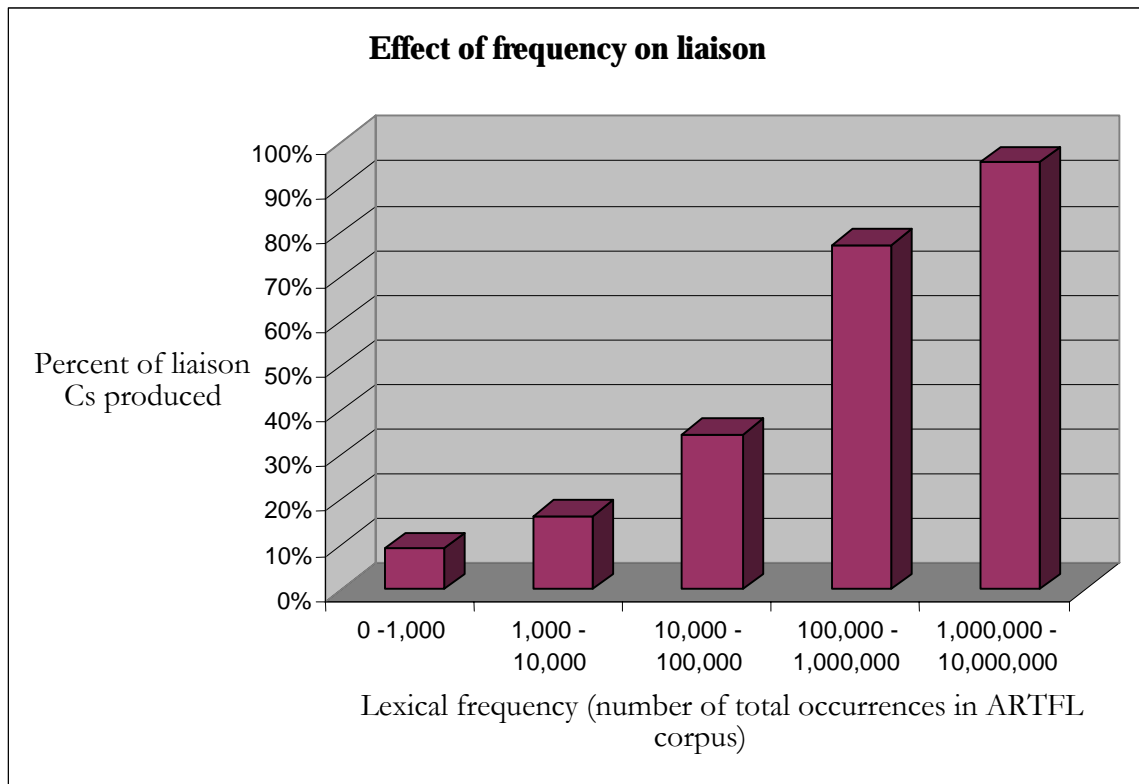


Figure 8 – Liaison production by liaison word frequency

If our task is to explain liaison frequency, the y-axis of Figure 8, and if frequency is the only explaining factor we are considering, then we can build a model which predicts the likelihood of liaison as a function of the length of the liaison word. One type of model would simply list the various length categories and the associated liaison probabilities. We could say that words that have a frequency of 0 to 1,000 occurrences in ARTFL, have a 0.07 probability of being produced. Such a model, though, would be sorely lacking in several respects. First, it would not be articulated in a language that is transferable to other domains, like the effect of frequency of an English consonant-final word and its rate of final consonant deletion. Since the model is specified in terms of two specific variables and their

values, one would be hard pressed to apply it directly to another similar problem without a whole lot of interpretation. Second, there is a very clear insight that it fails to capture. This is the direct relation between how frequent a word is and how likely its liaison consonant is to be produced. A simple statement of the liaison probabilities per condition is not in a form that can capture this generalization.

So instead of an exact description of the co-distribution of the two variables, what we need to develop is a more general characterization, one which captures the trend in question and also allows comparison with other, similar results. The statistically prescribed method involves the construction of a formula describing a regression line or regression curve, which maximally captures the correlation between the two variables.

Different relationships can hold between variables. In some cases, the relationship can be *linear*. That is, a straight line best captures the relationship. If we are trying to predict the value of the dependent variable, we can represent it as y in an equation like the one in (9), which captures such a linear model. Here, x is the value of the dependent factor, A is a coefficient, and B is a constant.

$$(9) \quad y = Ax + B$$

When a number of different independent variables influence the dependent variable, we need to formulate an equation which includes them. A simple example of such a model would predict the value of dependent variable y on the basis of independent variables x and z in an equation like (10). Here, A and B are coefficients, while C is a constant.

$$(10) \quad y = Ax + Bz + C$$

Given now that we've defined our dependent and independent variables, and that we have a proposed infrastructure for describing the relation between those variables, we still need two devices. The first is a method to determine the best-fitting model for the data. The second is a way to evaluate the significance of models and parts of models derived through this method.

Linear regression analysis provides two principal ways to arrive at a best model. Each of these involve taking a starting model and adding or subtracting factors that progressively improve the model. In step-forward methods, the starting model includes only a constant, such that there is no effect of the independent variables on the dependent one. Progressively, the term (such as Ax or the constant C) that best fulfills some selection criterion is added to the model. A very useful selection criterion, and the one used in the present study, is the *likelihood* of the model, that is, the probability of the observed results, given the model. In each step, the term is added which would increase a function (actually -2 times the log) of the model's likelihood to the greatest statistically significant degree.

The alternative selection method, step-backwards analysis, starts with a model including all the possible terms given all the factors included in the potential model. So if we start with 19 potential factors on liaison, all 19 will be included in the starting model. Unlike step-forward selection, which progressively adds terms, step-backward selection discards the terms whose loss will lead to the smallest statistically insignificant loss of Log Likelihood. That is, it throws away the most disposable terms until discarding any would create a significant decrease in the Log Likelihood of the model.

When possible, both methods should be used in order to triangulate on a single solution. Unfortunately, when too many possible terms are available, step-backward analysis becomes both too computationally taxing to be used practically, and too indiscriminate in its

retention of terms. Thus, step-backward selection will often yield a model with a large number of totally statistically insignificant terms. In these cases, only step-forward selection is possible.

Once either of these methods has arrived at a most likely model, that model and its terms must be evaluated. The model can be very transparently evaluated by the percentage of tokens that it correctly classifies, its significance and log likelihood. Each term included in the model, aside from its de facto relevance due to its inclusion in the most likely model, is assigned a significance measure, a value for p . P can range in value from 0 to 1, with smaller values indicative of more significant terms. The usual cut-off in the social sciences .05, which indicates that there is a 1/20 chance there is no relationship between the model or term and the data; that the empirical trends described by the term are purely random fluctuations.

Before moving on to a discussion of the limitations of the VARBRUL system, there is one further complication. In studies using VARBRUL and GOLDVARB to build regression models, as well as in the present work, dependent variables are always dichotomous – they have two possible values. But notice that linear equations generate continuous values for the dependent variable. In the case of liaison, we don't want to say that a certain set of factors conspire to make a liaison consonant be produced to a degree of 0.54; liaison consonants are either articulated or they aren't. In the construction of a general model of the influences of independent factors on a dependent one, this binarity wreaks all sort of statistical havoc with a model like linear regression that assumes its output to be continuous.

For this reason, VARBRUL models make use of a special type of linear regression known as logistic regression. Logistic regression uses the same methods as linear regression, but adds to it a filter that modifies the predicted value for the dependent variable. This filter

is called a logit transformation. The logit transformation Y of a probability p of an event is the logarithm of the ratio between the probability that the event occurs and the probability that the event does not occur, or $Y = \log(p/(1-p))$. This produces a *sigmoidal* or S-shaped curve as the output of the regression function, which more closely approximates the on-off or present-absent dichotomy needed to describe liaison, /*/ deletion, or /r/ deletion. The transformation from a linear to a sigmoidal function is demonstrated graphically in Figure 9.

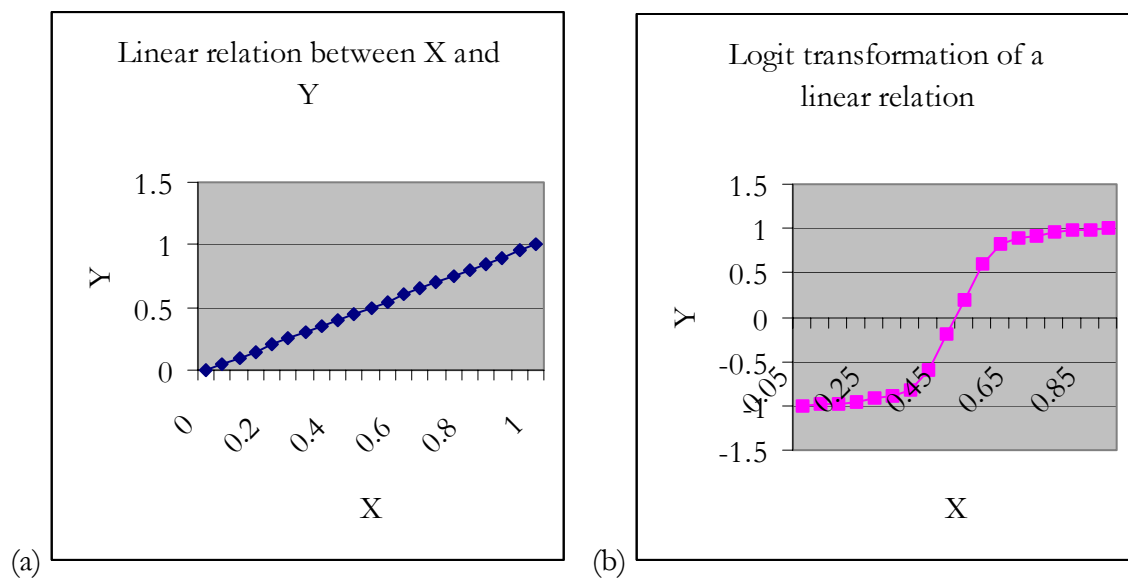


Figure 9 – Logit transformation of (a) a linear relationship into (b) a sigmoidal one

In sum, then, what VARBRUL uses to construct mathematical models of the effects of independent factors on dichotomous dependent variables is logistic regression. But the logistic regression included in VARBRUL, and consequently in all statistical studies that make use of VARBRUL is limited to independent factors which do not interact. Looking back at the form of the functions that can be produced by a logistic regression analysis, there is no place in an equation like the one in (11) for x and z to display interactions. Rather, in

equations like this one, \mathbf{x} and \mathbf{z} are necessarily independent; the contribution of the term $A\mathbf{x}$ will be identical, no matter what the value of \mathbf{z} is, and vice versa. Whether interactions effects were excluded for theoretical or practical reasons, the current state of affairs does not allow the possibility of testing the significance of interaction effects in VARBRUL.

(11) a. $y = A\mathbf{x} + B\mathbf{z} + C$
b. $y = A\mathbf{x} + B\mathbf{z} + C\mathbf{xz} + D$

And yet, there is a natural way to extend the statistical tools used in VARBRUL to cover interactions between factors as well. Logistic regression outside of VARBRUL can include not only terms describing variables multiplied by some constant, but also terms including the product of two independent variables and a constant. If we want to test the significance of interaction effects, such as the interaction between age and gender in the Milroy's study of */s/* deletion, these effects can be represented as in the interaction term $C\mathbf{xz}$ in (11b). This interaction term will have a value that is a function of the values of the two variables \mathbf{x} and \mathbf{z} and the constant C assigned to this term. Interaction terms are treated exactly like other terms in a full-scale logistic regression. They can be added in a step-forward selection or dropped from a step-backward selection. They are also assigned significance measures.

Because interactions between factors are of interest at present, full-scale logistic regression, including interaction terms, will form the basis for the statistical analysis of the French Liaison Corpus.

Now that the current study's empirical basis and analysis method have been outlined, we can move on to the crux of this chapter, the statistical analysis of factors on liaison. You

will recall that there are a number of outstanding questions that need answering, all of which revolve around a single query: *What factors probabilistically effect liaison use in the language of the community?* In the remainder of this section, we will test only autonomous effects, and in the following section, we will test interaction effects.

Details of the test of autonomous factors

In a first statistical test, I ran a step-forward logistic regression on all of the 19 factors enumerated in section 1 above, without including the possibility of interactions between them. The effects of these factors were tested only as they independently influenced liaison production.

The regression model that resulted from the selection process was able to correctly predict 88.3% of the data. That is, of 2559 tokens, 299 were predicted by the model to have the wrong liaison valence. This is an average to high degree of reliability for a problem with a similarly large problem space.

Of the 19 factors originally included, the selection procedure selected 12 for inclusion in the model on the basis of statistical increases of model log likelihood, as discussed above. These were those listed in Figure 10 below. Also included are some statistics. The leftmost number is the value of their Wald statistic for that variable, which is the square of the ratio of the coefficient to its standard error. This metric of significance is less important for present purposes than the measure of p, to be discussed shortly. In the next column is the number of degrees of freedom the variable had, followed by its significance (or p value), and its R value, which is a measure of the partial correlation between the particular independent variable and the dependent variable.

----- Variables in the Equation -----				
Variable	Wald	df	Sig (p)	R
AGE	12.8011	2	.0017	.0499
LIAIS-FREQ	58.1429	4	.0000	.1190
LIAIS-GRAM	113.2900	9	.0000	.1641
LIAIS-ORTH	17.6644	12	.1263	.0000
LIAIS-SYL	32.5244	5	.0000	.0798
NEXT-GRAM	78.1082	10	.0000	.1281
NEXT-ORTH	54.2204	4	.0000	.1143
PAUSE	9.1311	3	.0276	.0297
PERSON	8.1899	3	.0422	.0249
PLURAL	17.1285	2	.0002	.0609
PREC-SEG	12.8800	2	.0016	.0501
PUNCT	35.1698	2	.0000	.0939
Constant	.0301	1	.8623	

Figure 10 – Results of a step-forward logistic regression analysis of only independent factors.

We can clearly see from these results that aside from LIAS-ORTH, the spelling of the liaison consonant and the constant, all the factors included are significant with a $p < .05$. Why the liaison consonant's orthography was included despite its lack of significance is clear from the selection criteria. When a factor would significantly increase the power of the model, it was added to it. When LIAS-ORTH was included in the model, the model's reported its predictions to have improved by a full 1%, an improvement significant to a degree of $p < .001$ (not shown here). Other factors included were, from the top, the speaker's age, frequency of the liaison word, grammatical class of the liaison word, orthography of the liaison consonant, syllabicity of the liaison word, grammatical class of the next word, orthography of the next segment, length of pause, person and plural expressed by the liaison segment, the identity of the preceding segment, and the punctuation intervening

between the two words. A constant was also included in this model, which served to skew the output of the regression slightly to one side regardless of the values of any of the factors. Let's look now in more detail at the effects of these factors.

Liaison, following and preceding segments

Aside from LIAIS-ORTH, there are two other orthographic factors included in the above model: PUNCT (punctuation) and NEXT-ORTH (the next segment's orthography).

Interestingly, aside from PAUSE, none of these factors' phonological equivalents, LIAIS-PHON, PAUSE, or NEXT-PHON are included in the model. What we can conclude from this is that the orthographic versions better explain the liaison distribution than their phonological peers. This is at first surprising, until we remember that the corpus is composed entirely of read speech. Thus, spelling pronunciation might overrule natural phonological patterns.

In a broad sense, each of these orthographic-phonological pairs should have the same sort of effect on liaison. We would expect pauses to be predicted by commas, and for following semivowels to have similar effects whether spelled or spoken. In order to put this to the test, I ran another step-forward logistic regression, this time excluding the orthographic elements. This would test whether when the orthographic factors weren't busy explaining the data away, their phonological partners were significant. The regression test turned up negative – neither NEXT-PHON or LIAIS-PHON were included in this model.

The spelling of the liaison segment was in fact highly informative. Figure 11 shows the percentages of liaison segments articulated when they had each of the four most common phonological forms. Compare this with the effects of liaison segment orthography,

in Figure 12. While the phonological identity of the liaison consonants only really allows a major division between /r/ and the other phonological forms, as shown in Figure 11, the orthographic variation internal to liaison phones is highly informative. Liaisons in /z/ spelled with “z” are quite less likely than those spelled with “x”, for example. The significance of both orthography and frequency, as we will see below, implies that the relatively high production of liaison spelled in ‘x’ cannot be solely explained in terms of the potentially greater frequency of ‘x’ as a nominal plural marker or verbal first or second person marker.

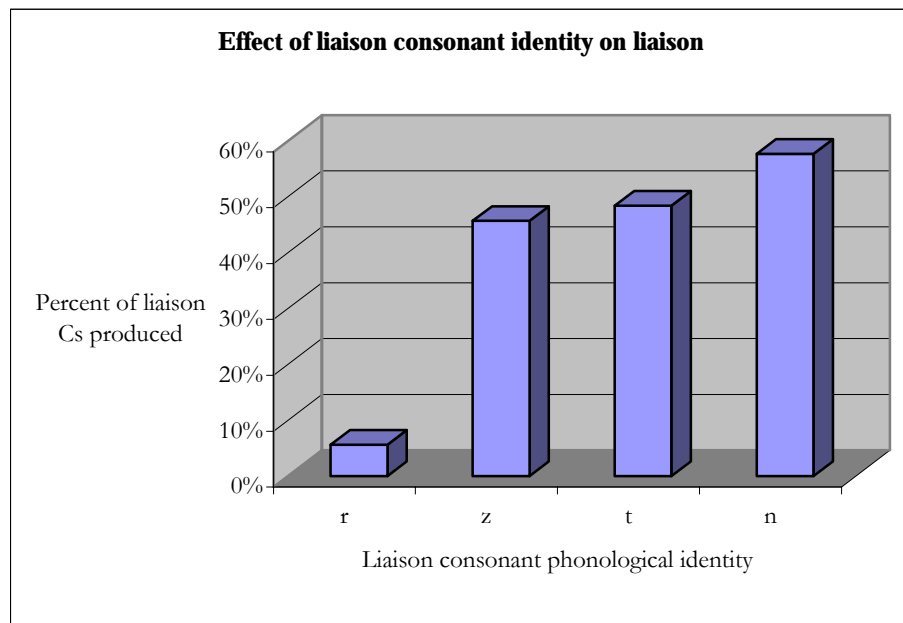


Figure 11 – Percentage of liaison consonants produced as a product of liaison segment identity

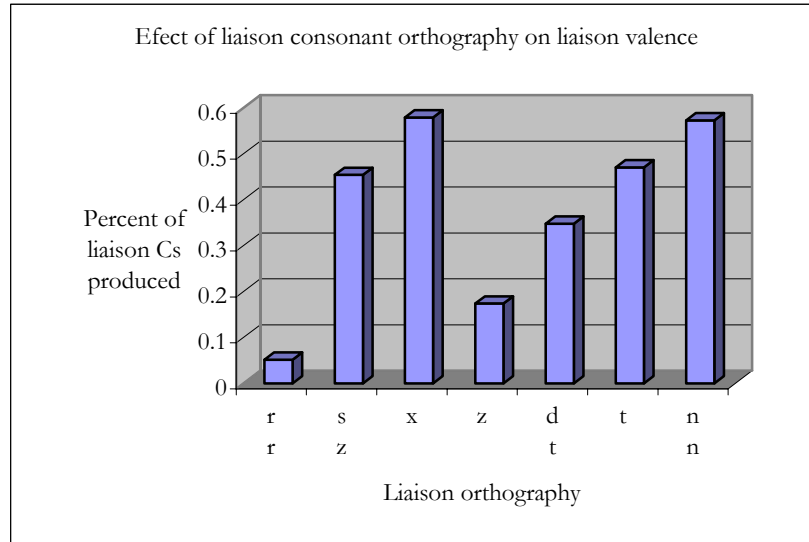


Figure 12 – Percentage of liaison consonants produced as a product of liaison segment orthography

In a similar turn of events, the orthography, rather than the phonology, of the following segment is most significant to liaison valence (Figure 13). But the effect of orthography is quite the opposite of what we might reasonably expect. Rather than decreasing the likelihood of liaison, words starting in “y” increase it. Could this indicate that semivowels in French are more vocalic than vowels? Upon further inspection, the great majority of words starting with “y” in the corpus (86%) were the word *y* ‘there’. This word, pronounced /i/, provided a very fertile liaison context. This fact, although interesting from a grammatical perspective, since *y* is a clitic pronoun and not a lexical head, leaves the semivowel question open. The second most frequent word beginning in orthographic “y”, *yeux* ‘eyes’ provoked liaison in all seven of its instances, in all cases following a modifying determiner or adjective. But these numbers are too small for us to conclude anything substantial about the role of semivowels.

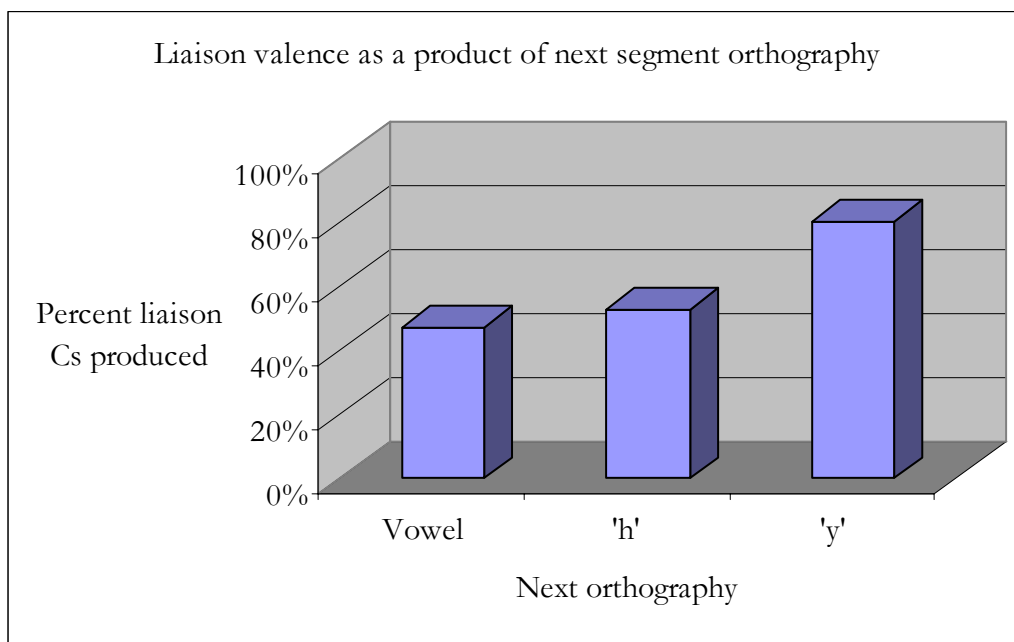


Figure 13 – Percentage of liaison consonants produced as a product of next segment orthography

On the other hand, the “h” results are worthy of some investigation. Just like “y”-initial words, words with initial “h” provoked liaison in preceding words more often than did vowel-initial words. Now it just so happens that of the words with initial “h” in the corpus, listed in (12) below, only *hasards* ‘chances’, *hausse* ‘increase’, *haut* ‘high (m)’, *haute* ‘high (f)’, *hauts* ‘high (m pl)’, and *home* ‘home’ are h-aspiré tokens. Of these, only *hasard* gave rise to a produced liaison consonant in its single occurrence, while the others, each of which was produced once, except *haute*, which had two occurrences in the corpus, defied liaison. Liaison was slightly more prevalent in the rest of the h-initial words than it was for vowel-initial words. From these data, we can draw the weak conclusion that h-aspiré is not used exactly as prescribed, because of the token of *hasard* that repressed liaison. The general

attraction of liaison consonants to non-h-aspiré h-initial words has, as far as I know, never been previously documented, and deserves further study.

(12) Words with initial 'h' in the liaison corpus

habitants 'inhabitants', *habitats* 'habitats', *hasards* 'chances', *hausse* 'increase', *haut* 'high (m)', *haute* 'high (f)', *hauts* 'high (m pl)', *hectares* 'acres', *hectometers* 'hectometer', *hélas* 'alas', *hermétique* 'hermetic', *heures* 'hours', *heureuse* 'happy', *hexagonale* 'hexagonal', *hiérarchisées* 'hierarchical (pl)', *histoires* 'stories', *home* 'home', *homage* 'homage', *homes* 'men', *honneurs* 'honors', *horreurs* 'horrors', *hors* 'outside', *hospitalier* 'hospital (adj.)', *huit* 'eight', *huitante* 'eighty', *humain* 'human', *humains* 'humans', *hyènes* 'hyenas'

The discussion above about the relevance of the preceding segment to liaison use cited two studies, both of which showed that a preceding vowel increases the likelihood that a liaison consonant will be produced, while a preceding consonant diminishes it. The presence of the PREC-SEG factor in the regression model demonstrated above provides further confirmation of the role of the preceding segment. In fact, the disparity was quite significant – liaison consonants following consonants were produced only 23% of the time, while those following a vowel were articulated in 54% of their possible utterances. The question that was left open above, the effect of semivowels in pre-liaison position, will remain unanswered, as there were too few data points in the current corpus for any statistical merit to be accorded them.

Pause and Punctuation

Turning now to prosody, both PAUSE and PUNCT were included in the regression model, and both met significance criteria. As can be seen from Figure 14, the presence of a pause greater than 200msec correlated with significantly fewer liaison consonants being produced. A similar effect is produced by punctuation (Figure 15). When there is no punctuation intervening between the liaison and following word, liaison is more than six times more likely than when a comma intervenes. No liaisons were produced crossing a period, semicolon, or colon. The fact that both of these variables were included in the regression model indicates that they explain slightly different data; otherwise the more powerful one would explain away the substance of the less powerful one, leaving it out of the running.

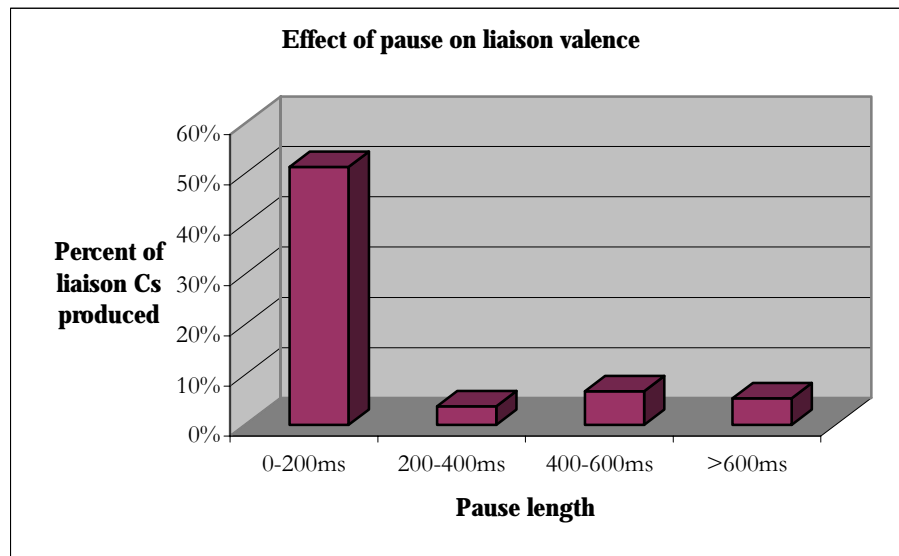


Figure 14 – Percentage of liaison consonants produced as a product of pause

We can test the degree of overlap between the two factors statistically. The correlation coefficient for a set of data, where each data point has a value for each of two

variables is the degree to which the values of the two variables correlate. Values for correlation coefficients range from -1 to 1 , with negative values being negative correlations. A positive value indicates a positive correlation, such as the one between age and ear-lobe size. A value of 1 or -1 means an exact correlation between the variables, and 0 means no correlation. The correlation coefficient between PAUSE and PUNCT, considering only whether there was comma or not, and whether there was a pause or not, was 0.43 . That is, there was a positive correlation of very slight magnitude.

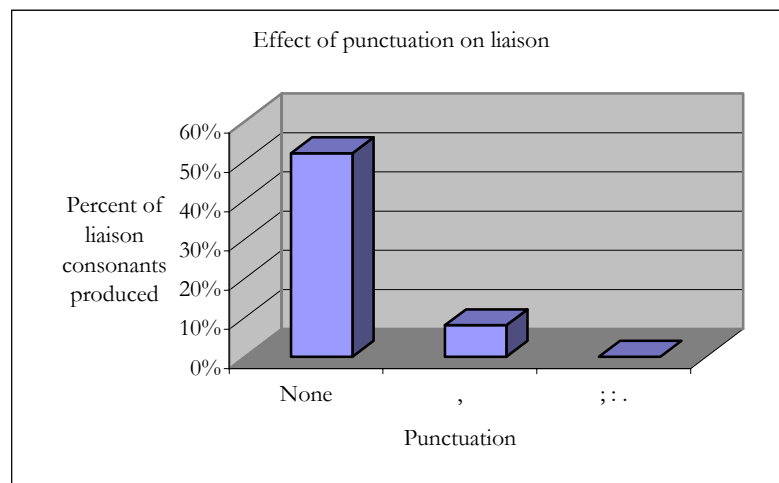


Figure 15 – Percentage of liaison consonants produced as a product of punctuation

Possibly, the divergence in the variables lies in the wider range of possible locations of pauses. Speech is peppered by pauses that are not marked by punctuation. Some delimit prosodic or syntactic units. Others are disfluencies, due to processing, articulatory, or other constraints on real-time speech production. Still others can be used for contrastive emphasis. But none of these groups include comma-delimited pauses. If this is true, it suggests that the classic view of liaison as occurring within a prosodically defined breath unit has some validity. Turning to the corpus, we find that of 2359 comma-less tokens, 67 of these had a

pause greater than 200msec (Figure 16). And of these 67 comma-less pauses, only 7 (10%) has articulated liaison consonants. So there were a number of pauses without commas, and only rarely were liaison consonants produced in these environments.

	None	>200msec
None	1216/2292 (53%)	15/111(10%)
Comma	15/111(14%)	0/76 (0%)

Figure 16 – Liaison valence as a function of punctuation and pause

Similarly, since the corpus was composed of read speech, the presence of a comma could have suppressed liaison consonants, even where no pause was articulated. Just anecdotally, I can attest after listening to a number of tokens with commas and no liaison that many of these had no pause intervening between the two words or lengthening of the first word. Numerically (in Figure 16), of the 187 tokens with commas intervening between liaison word and following word, 111 had no pause, and of these cases, only 15 (14%) was their liaison consonants produced. This suggests that the presence of a comma repressed liaison use, even when no pause was actually articulated. The independent effects of pauses and punctuation on liaison valence suggest that prosodic fluency and orthographic marking independently suppress liaison.

Morphological status

Whether or not a liaison consonant bears a meaning, and what meaning it bears, are significant factors in determining whether that consonant will be produced. Both PLURAL and PERSON were included in the regression model. Let's first look at the effect of person.

Our first observation is that when a liaison consonant expresses a person marking on a verb, that affects the probability that it will be articulated. But the influence is not uniform. Although we might predict that meaningful segments would be more likely to be produced, it's only third person /t/, and not first or second person /z/ that gains from its meaningfulness. This is shown in Figure 17, where the percent of liaison segments produced in the first and second person is much lower than the percent produced in the third person condition. In fact, being a first or second person morpheme actually decreases the likelihood that a segment will be produced to a significant degree.

This fact can't be explained simply in terms of the difference in phonological identity between these two segments. You might remember from Figure 11 above that /t/ liaisons were barely more frequent than /z/ liaisons. But /z/ liaisons expressing the first and second person are almost always spelled with "s", occasionally with "x", and never with "z", for example in *pais* 'leave (sg. 1st or 2nd pres.)' or *veux* 'want (sg. 1st or 2nd pres.)'. Thus, looking now at Figure 17, we recognize that first and second person /z/ liaison might even be favored over third person /t/ liaisons on the simple basis of their spelling. In other words, just as Morin and Kaye (1982) suggested, third personal verbal inflection attracts more liaison production than does first (or second) person inflection. And this morphological effect can't have a purely functional explanation, since first and second person morphemes see less liaison production than forms expressing no person.

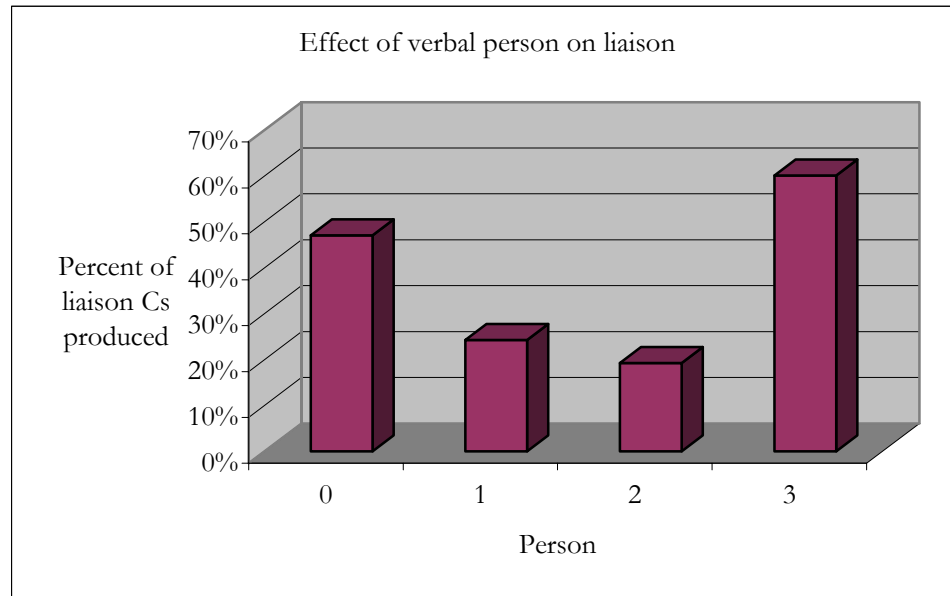


Figure 17 – Percentage of liaison consonants produced as a product of person

The number expressed by a liaison consonant also has an effect on its pronunciation, although a little less transparently than that of verbal person. In the entire corpus, words having a final liaison consonant that expresses the plural are in fact pronounced somewhat less frequently than consonants not expressing the plural (Figure 18). But pseudo-plurals, words with a plural semantics and a final, non-morphological /z/, give rise to liaison production much more frequently than either of these first two classes. This factor was included in the regression analysis, with an extremely strong significance rating, so it must be capturing some vital generalization that cannot be better explained by other factors.

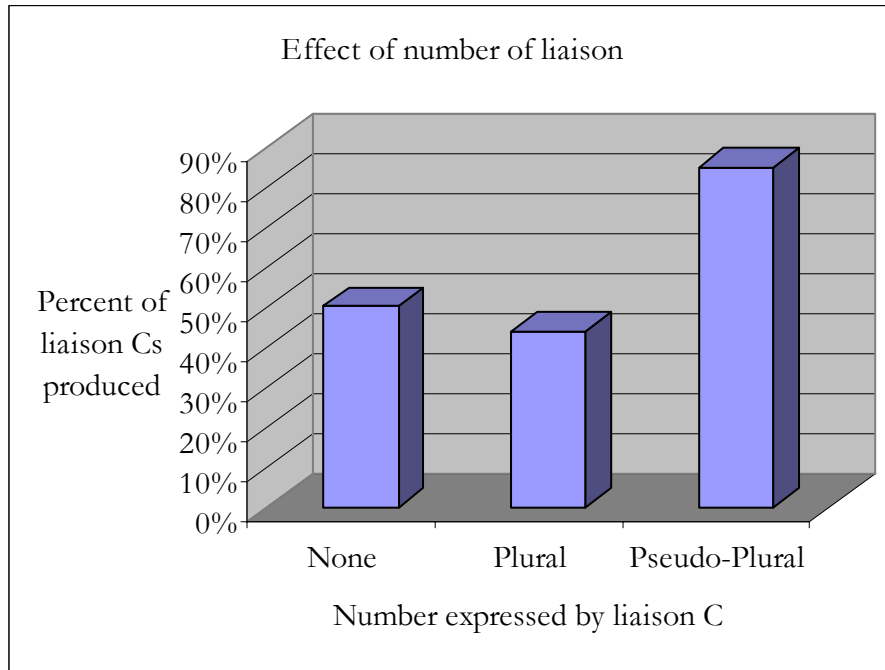


Figure 18 – Percentage of liaison consonants produced as a product of number

When we look at the words that were coded as having final pseudo-plurals, we find that they are not particularly numerous, totaling only 82 tokens. They are mostly frequent, short words with final /z/ or /t/. As such they would already be expected to be quite attractive to liaison. And yet, because of the inclusion of this factor in the regression analysis, we can conclude that this set of words happens to inspire liaison more than would otherwise be predicted on the basis of the other factors. In other words, the inclusion of this feature may be due to the lexically specific bias for these words in particular to be produced with their final liaison consonants intact.

Syntax

The syntactic classes of the both the liaison word and the following word were deemed sufficiently informative by the step-forward analysis process to be included in the regression model. LIAIS-GRAM and NEXT-GRAM both proved extremely significant. In fact, they were the second and the third variables added to the model. This means that, after the frequency of the liaison word was incorporated into the model, they were the most efficient predictors of the data.

From the very first glance at Figure 19, which summarizes the effect of grammatical class on liaison, we can tell what a crucial role grammatical class plays. A first, compelling observation is that content words and function words seem to cluster among themselves. The leftmost half of the figure includes all and only content categories: nouns, proper nouns, verbs, and adjectives. The right half is the exclusive domain of function classes: conjunctions, preposition, pronouns, and determiners. This trend even goes so far as to place adverbs in the very center. Adverbs are of somewhat ambiguous content/function status. While adverbs can be formed more or less productively in French from adjectives by adding the suffix *-ment* ‘-ly’, there are also subgroups of adverbs whose membership is quite restricted. Examples are adverbial quantifiers, like *tant* ‘so much’, *ainsi* ‘thusly’, and *trop* ‘too much’, a group to which no new members can be readily added, or conjunctions like *mais* ‘but’ and *sinon* ‘otherwise’.

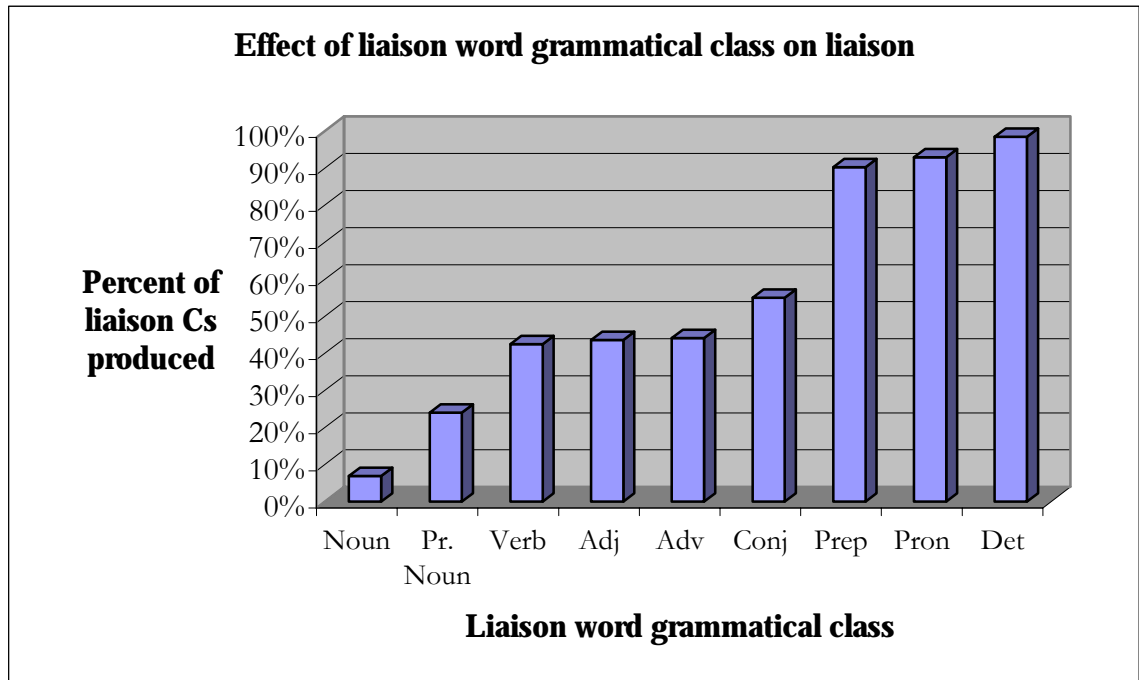


Figure 19 – Percentage of liaison consonants produced as a product of liaison word grammatical class

This schism in the liaison induction of function and content words constitutes yet another probabilistic pairing between phonological and syntactic knowledge, as argued for in Chapter 2. In a similar fashion, within the form and content categories, different grammatical classes tended to evoke liaison to different degrees. For example, verbs were almost eight times more likely to have a final liaison consonant produced than nouns were. Divisions like this among grammatical classes serve to reinforce the argument in Chapter 2 that morphosyntactic properties are relevant to phonological ones.

You'll remember that according to Booij and De Jong (1987), liaison is obligatory after determiners. In Figure 19, we see that determiners do not quite always evoke liaison. We might draw from this that there are exceptions to Booij and De Jong's generalization. In fact, there are only two cases of liaison absence on a determiner, in the sequences *un home* 'a

retirement home’ and *un à* ‘one to’. In the first case, the following word, a recent borrowing from English and therefore an h-aspiré word, defies liaison. That is, it seems that in this case the determiner’s desire for a liaison consonant is trumped by the following word’s rejection of it. In the second case, the word *un* ‘one’ was misclassified by the part-of-speech tagger as the identically spelled article *un* ‘an’.

So determiners, aside from exceptional cases, require liaison. With other function words, however, the story is quite different. Pretty much every frequent preposition and pronoun in the corpus has at least one token without liaison. Conjunctions, the least liaison-attracting function class, consisted mostly of *mais* ‘but’ and *quand* ‘when’, and had liaison production about half of the time. Notice that this still gave them more articulated liaison consonants on average than any of the content categories. This is most likely the case because the exceedingly frequent conjunction *et* ‘and’ was not included as a potential liaison word. Despite its final orthographic “t”, *et* does not have a final consonant in any context. If *et* had been included, it would have sharply driven down the liaison frequency for conjunctions.

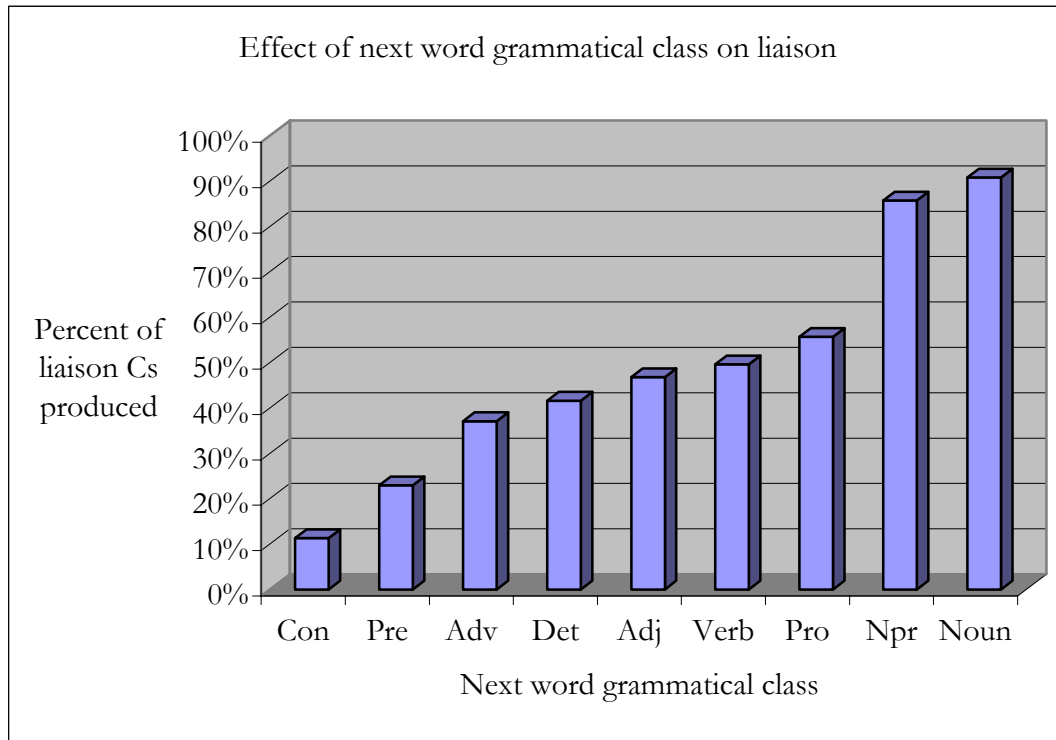


Figure 20 – Percentage of liaison consonants produced as a product of next word grammatical class

Turning now to the grammatical class of the following word (Figure 20), we once again discover a split between function and content words, with two differences. The first is that pronouns seem to have drifted into the jurisdiction of content words. This aberrant situation is due to one particular construction, which unequivocally requires liaison. This is the subject-verb inversion construction, exemplified in (13b) and (13d) below. In interrogative sentences, the verb can come before a subject pronoun. In these cases, which are subject to other syntactic constraints not relevant to the present discussion, the final liaison consonant of the verb is always produced. When no final liaison consonant appears on the verb, as in (13d), an epenthetic /t/ is added, and is produced in all cases. This subject-verb inversion accounts for about 40% of the instances of pronouns as words

following liaison consonant, and is thus responsible for the strange placement of pronouns among verbs and nouns. This epenthetic /t/ seems to be a morpho-syntactically governed liaison-phenomenon.

- (13) a. *Ils veulent des haricots.* ‘They want green beans.’
b. *VeulenT-ils des haricots.* ‘Do they want green beans?’
c. *Il parle de quoi?* ‘What’s he talking about?’
d. *De quoi parle-T-il* ‘About what is he talking?’

The second obvious distinction between the effect of preceding words and following words is the directionality of the effect. A liaison word evokes the production of more liaison consonants when it is a function word, while a following word promotes less liaison if it’s function word. We’ll test whether there is any interaction between these generalizations through statistical tests of interactions between factors in Section 6 below. By doing so, we’ll also be able to investigate the role of head-complement structure and specific pairings of liaison words and following words.

Word length

While Encrevé (1983) has established that monosyllabic liaison words are more likely to evoke liaison than are polysyllabic words, the relative effects of polysyllabic words of different lengths have yet to be assessed. The inclusion of the factor LIAIS-SYL in the regression model indicates that liaison word length is significant to liaison production. First, as Figure 21 shows, monosyllabic liaison words were much more likely to give rise to liaison

production than were polysyllabic words, confirming Encrevé's result. But in addition, among polysyllabic words, disyllabic words see the most liaison production. Liaison word length seems to be not simply a binary variable, but at least a ternary one.

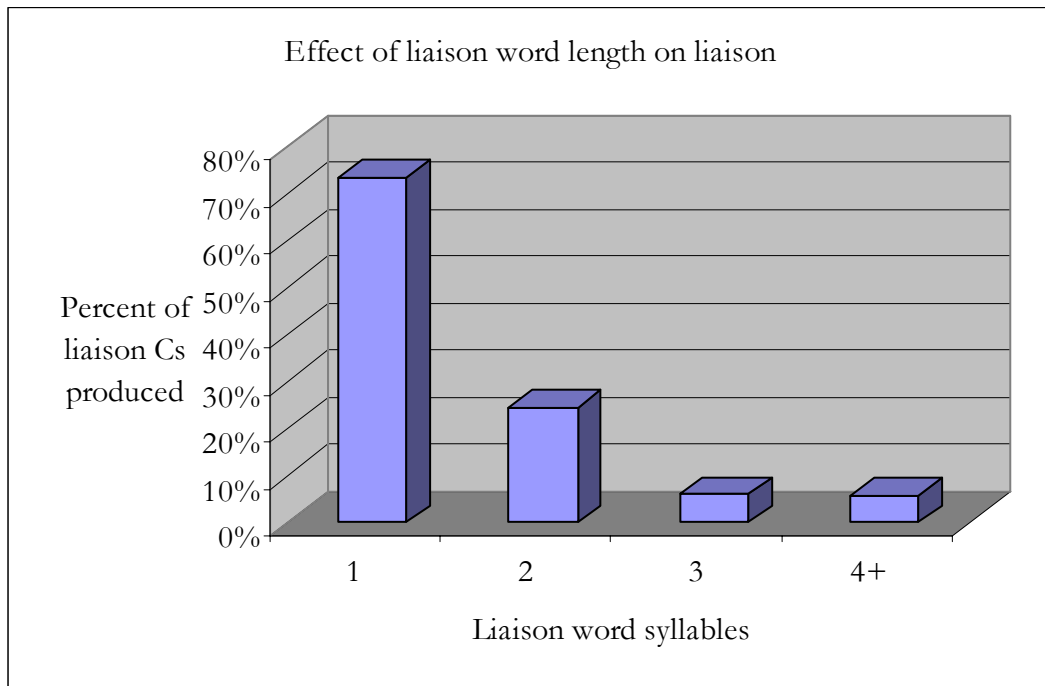


Figure 21 – Percent of liaison consonants produced as a function of number of liaison word syllables

The non-inclusion of the next word's length in the regression model denies its independent contribution to liaison. But relative length of the two words might very well prove significant. We'll have to wait to see if there was an interaction between LIAIS-SYL and NEXT-SYL to know about relative length.

Social factors

The final factor included in this model is the social factor AGE, confirming Ashby's (1981) and Booij and De Jong's (1987) finding that age correlates with liaison use. Just as in those studies, older speakers used liaison more than younger speakers in the current corpus, as shown in Figure 22. Interestingly, EDUCATION and SEX are excluded.

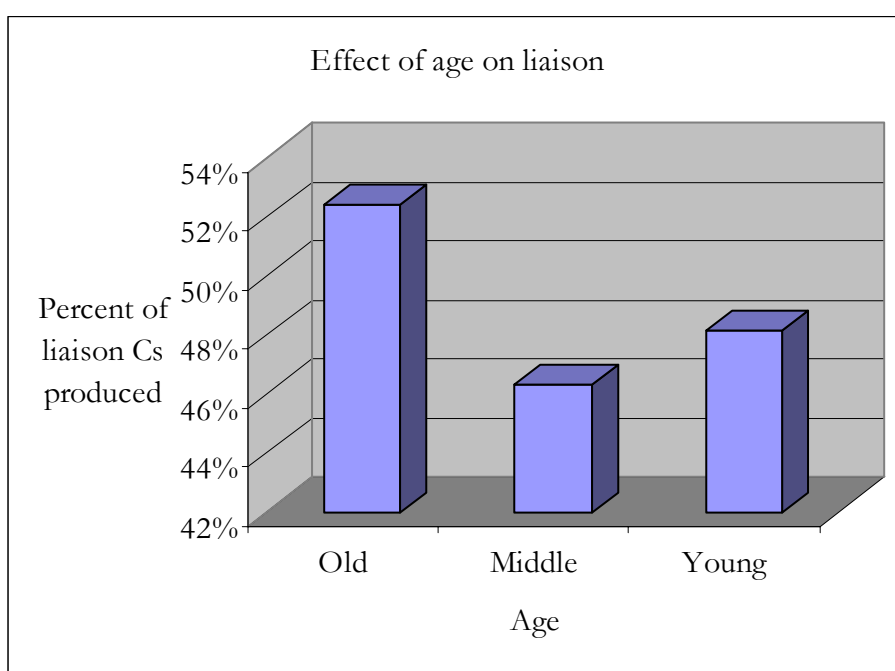


Figure 22 – Percent of liaison consonants produced as a function of speaker age

There was also a slight tendency for more educated speakers to produce liaison consonants more often than their less educated peers (Figure 23) and for women to produce more liaison consonants than men, 50% versus 48%. But these tendencies did not meet statistical significance criteria, and so can only be taken as anecdotal.

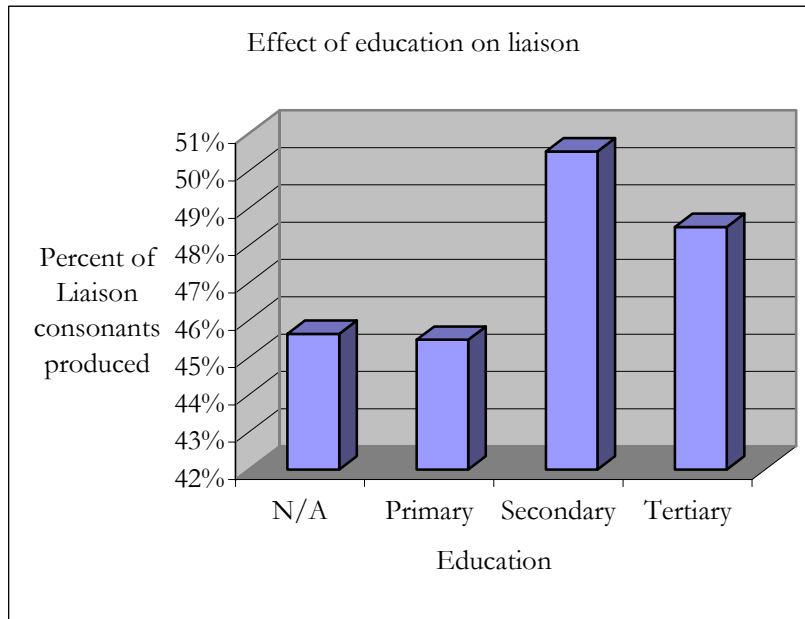


Figure 23 – Percent of liaison consonants produced by education

Frequency

We’ve already had a chance to observe the effect of liaison word frequency on liaison – this was the case study of a simple direct effect of an independent variable shown in Figure 8, above. Not only was liaison word frequency highly significant, it was the first factor included in the model, and explained a total of 83% of the data, all by itself. This confirms the hypothesis that increased frequency of a liaison word increases the likelihood of liaison being produced, a trend precisely opposite to what happens in the realm of t/d deletion, for example, where increased frequency correlates with increased deletion.

The fundamental criterion for liaison seems to be frequency. Frequency is what ties several of the major factors already identified together. Determiners and other function words are much more frequent than content words, and produce more liaison. /t/, /z/, and /n/ are more frequent than /k/ and /r/, and thus produce more liaison. Shorter words are

more frequent than longer words, and thus afford the production of liaison consonants. Of course, the results of the regression analysis presented above show that frequency can't be the only game in town. The other factors were included in the model for a reason; they explain some of the data that frequency alone cannot. But we've excluded from the present analysis another factor, which might be as informative as liaison word frequency; the frequency of the pair of words. Reliable word-pair frequencies, as mentioned above, were not available.

Notice that NEXT-FREQ, the measure of the frequency of the following segment, was not included in the model. This goes against Agren's (1973) finding that liaison consonants were more likely to be produced when preceding a frequent word. But there may yet be a role for the next word's frequency in combination with liaison word frequency. That is, there may well be an interaction between LIAIS-FREQ and NEXT-FREQ. For example, frequent liaison words might be more likely to evoke liaison before more frequent following words than before less frequent ones. This would not only demonstrate the relevance of the next word's frequency, but would also constitute circumstantial evidence for the word pair frequency hypothesis, described above.

Summary

This catalogue of effects on liaison, phonological and otherwise, should serve as convincing evidence that phonological and other factors contribute probabilistically to the realization of liaison. In order to predict whether a liaison consonant will be pronounced, phonological details like the segment's identity and the identity of the preceding segment need to be taken into account. Syntactic details like the grammatical class assignments of the liaison word and

the following word, morphological ones like what person and number a liaison expresses, and social ones like the age of the speaker need to be addressed as well. The same results presented above reinforce previous findings, while others conflict with previously reported results and yet others contribute evidence about new factors affecting liaison.

6. A test of interactions between factors

In this section, selected interactions between variables described above are added to the list of potential terms to be incorporated into a regression model by a step-forward regression model. The main question to be asked is the following:

- Are there statistically significant interactions between factors? If so, are these restricted to phonology- morphology-, or syntax-internal interactions? Or are there cross-modal interactions as well? If there are cross-modal interactions, are there any between linguistic knowledge and extralinguistic knowledge, such as, for example, between social variables and syntactic ones?

Methods

Logistic regression models, as explained above, can include interactions between factors.

There are two ways to proceed from the autonomous factor model in Section 5 in testing whether any interactions were significant in their contributions to the realization of liaison.

The first takes the regression model developed in Section 5 as a given, base model, to which additional factors are added, again using a step-forward selection method. Let's let this

approach be known as the *autonomous base* solution. The second approach is to start our modelling effort over, including in the list of potential terms those autonomous factors that were shown to be significant above, alongside those interactions that we're interested in testing. This larger pool of factors can then be subjected to the same modelling efforts as the autonomous ones previously were. This will be called the *re-pooling* solution.

These two methods do not yield identical products, because in some cases, interaction terms may be more powerful explainers of the data than autonomous terms are. The step-forward process selects the most succinct set of factors that are capable of explaining the data. Because of this effect, the re-pooling solution might cause some autonomous variables that would otherwise be classified as significant to be excluded.

The two methods additionally require different assumptions and yield different conclusions. The autonomous base model makes the simplifying assumption that autonomous effects are primary, and interactions secondary. That is, it claims that on top of autonomous effects, which have been well-documented and have been the object of intense empirical and theoretical study, there may additionally be interactions between factors. It goes without saying then, that starting with an analysis based exclusively on autonomous factors, and then adding interactions to that analysis, allows us to investigate whether in addition to the well-known autonomous factors, there are also interaction effects.

The re-pooling approach makes less stringent assumptions and allows more general conclusions. Rather than requiring the primacy of autonomous over interactive effects, the re-pooling solution assumes only that autonomous and interactive terms might have effects on the dependent variable. And on the basis of a model developed through this approach, we can conclude what the most effective model of the data in question is, given all the given factors and interactions between those factors. Unfortunately, re-pooling also can erode the

effects of individual factors, because interaction terms are more powerful than independent ones. Since I assume that autonomous effects are primary, I will use an autonomous-base method.

In the following section, I will present the results of an autonomous-base method, testing a list of interaction terms that I generated on the basis of previous literature and the tendencies I picked out from the literature.

As presented in Section 3, interactions between social variables abound in sociolinguistic data, although they are rarely analyzed statistically, due to the prevalence of VARBRUL as the statistical workhorse in sociolinguistics research.. I included the three possible terms representing potential social interactions, age by gender, age by education, and education by gender.

I hypothesized that the relative conservatism by older speakers might also lead to a difference in how the different age groups used liaison consonants with different functional significances. Thus, older speaker might be more likely to use liaison across the board, while younger speakers might use it only when it has some semantic value. They might also be more likely to use liaison in particular grammatical contexts. I thus also included the interaction between liaison word grammatical class and speaker age.

The interaction between age and number was thus included. Since spelling of the liaison consonant was a significant factor in the autonomous model, I imagined that education might interact with this factor. After all, greater familiarity with orthography or just with written language, might lead to increased spelling pronunciation by more educated speakers. This effect translates into the interaction between education and liaison orthography.

As discussed above, liaison word frequency is an extremely significant factor in determining liaison production. It has also been suggested that the frequency of the following word is relevant. Finally, although it was not established here, the frequency of the word pair has also been claimed to constitute a major determinant of liaison. We can then ask whether the word-pair effect is simply the result of an interaction between liaison word frequency and following word frequency. Perhaps more frequent liaison words evoke more liaison before other frequent words than they do elsewhere. Is there a significant interaction between liaison word frequency and next word frequency?

The same question can be asked for syllabicity. It has been claimed that when liaison words are shorter than following words, the climate encourages liaison. We can use the interaction between liaison word length and next word length (both in syllables) to evaluate the extent to which this effect is felt in the present corpus.

Finally, much of the generative literature on liaison deems the syntactic relation between the liaison word and the following word to be the most crucial factor in determining liaison obligatoriness (e.g. Selkirk 1974, Selkirk 1978, Selkirk 1984, and Kaisse 1985). Presumably, then, liaison might be explained in a large part by the grammatical classes of the liaison word and the following word, in addition to the interaction between these factors. Of course, not all of syntactic structure is captured by part-of-speech classes, but grammatical class alone may serve to explain a large part of the liaison data. A summary of these interactions appears in (6).

- (6)
- | | |
|------------------------|------------------------|
| AGE * EDUC | AGE * GEN |
| AGE * PLURAL | EDUC * LIAIS-ORTH |
| EDUC * GEN | LIAIS-FREQ * NEXT-FREQ |
| LIAIS-GRAM * NEXT-GRAM | LIAIS-SYL * NEXT-SYL |

The autonomous base model

I performed a logistic regression analysis for each of the interactions in (6), each of which started with the model described in Section 5 as the base model. Given this model, a step-forward selection procedure was run on the interactions in question. Four of the resulting models included, in addition to the 13 autonomous factors constituting the base model, an between factors (Figure 24).

----- Variables in the Equation -----				
Variable	Wald	df	Sig	R
AGE * EDUC	17.5384	6	.0075	.0618
----- Variables in the Equation -----				
Variable	Wald	df	Sig	R
AGE * LIAIS-GRAM	10.8066	16	.1821	.0000
----- Variables in the Equation -----				
Variable	Wald	df	Sig	R
LIAIS-FREQ * NEXT-FREQ	28.1362	16	.0305	.0000
----- Variables in the Equation -----				
Variable	Wald	df	Sig	R
LIAIS-GRAM * NEXT-GRAM	63.4919	46	.0445	.0000

Figure 24 – Partial results of four step-forward logistic regression analyses, starting with the model in Figure 10 and adding interactions between factors.

These additions to the model led to a statistically significant increase in the model's predictive power, from 88.3% accuracy to 90.8% accuracy, as well as in several other metrics: log likelihood and goodness of fit. Let's take a closer look now at the interactions that were included in this model.

The first of the interaction terms, age by education (Figure 25), involves differential effects of age, depending on another other social property of the speaker. Young speakers use more or less liaison than older speakers, depending on their degree of education. Old speakers with a primary education omit more liaison consonants than their younger counterparts with the same education (the leftmost cluster of columns in Figure 25). But when old speakers have a secondary or tertiary education, they articulate more liaison consonants than younger, equally educated speakers, on average.

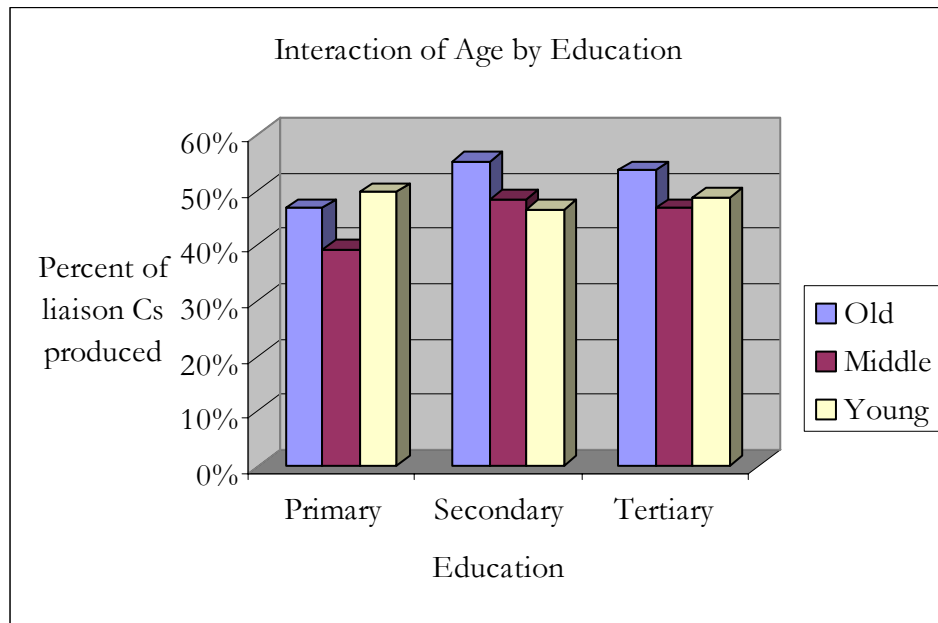


Figure 25 – Liaison valence as a product of speaker age and education

This effect may in fact be due to the simple detail that many members of the youngest age group, those younger than 25 at the time of the elicitation, were not old enough to have entered into higher education. That is, it might not be the amount or level of education that a speaker has had that's influencing liaison use here, but rather the educational track that speaker is on or their educational capacity. This would explain why many younger, less educated speakers would use more liaison than their older counterparts; they are bound to be more educated or capable of being so, but simply have not yet reached the appropriate age. This would raise the liaison use among less educated, younger speakers relative to older, less educated speakers.

This interaction involves only social variables, making it similar in type to the interaction between age and sex in Belfast /ð/, discussed in Section 3. But two other interaction was shown by the regression analysis to be significant involved interactions between social and linguistic variables.

The interplay of age and liaison word grammatical class is depicted in Figure 26. Here, it's plain that young speakers omit more liaison consonants when the liaison word is a verb than when it's an adverb. By comparison, middle-aged and older speakers exhibit little difference in their treatment of these two classes. This trend might indicate a dropping off of the morphological uses of liaison and an increase in its lexicalized uses, as with adverbs.

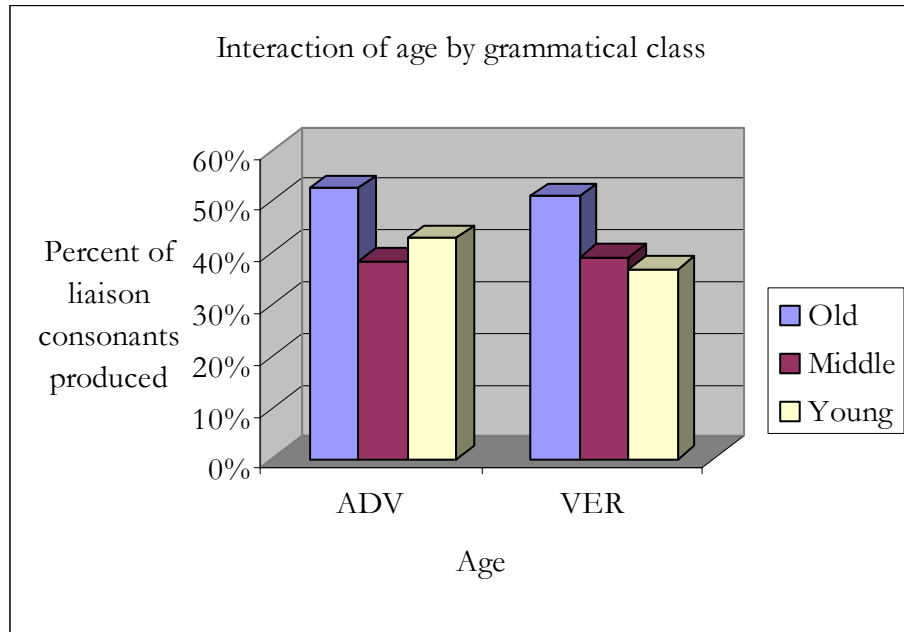


Figure 26 – Liaison valence as a product of liaison word grammatical class and speaker age

The penultimate interaction selected by the regression model is the one between liaison word grammatical class and following word grammatical class. The graphical depiction in Figure 27 departs from the bar graph format of the previous figures as the line graph version of this particular interaction was more interpretable than its bar graph counterpart. Here, the two independent variables are presented in the order of their independent contributions to liaison use. Thus, on the X-axis, syntactic categories of the liaison word decrease in the likelihood liaison will be produced from left to right. By the same token, the next word's grammatical class is listed in the legend in decreasing order of the percent of liaison consonants produced. You will notice that many of the lines are broken. Where there were fewer than 5 tokens for the intersection of a liaison word class and next word class, such as a liaison conjunction and a following noun, no node was placed on the chart, meaning that no line could be drawn through that point.

Aside from the general, downwards left to right trend of the graph, corresponding to the autonomous effect of LIAIS-GRAM, there are also all kinds local interactions between values of the two variables. If there were no interaction, all the lines would be parallel. As it stands, certain combinations of the two variables give radically different liaison results than their near neighbors. For example, when the liaison word is a pronoun, liaison is about half as likely if the following word is an adjective than when it's a noun, pronoun, or verb. As liaison words, determiners always take liaison if they precede an adjective, but not so for pronouns, prepositions, or adverbs when they're liaison words. Before adjectives, they take liaison closer to half of the time.

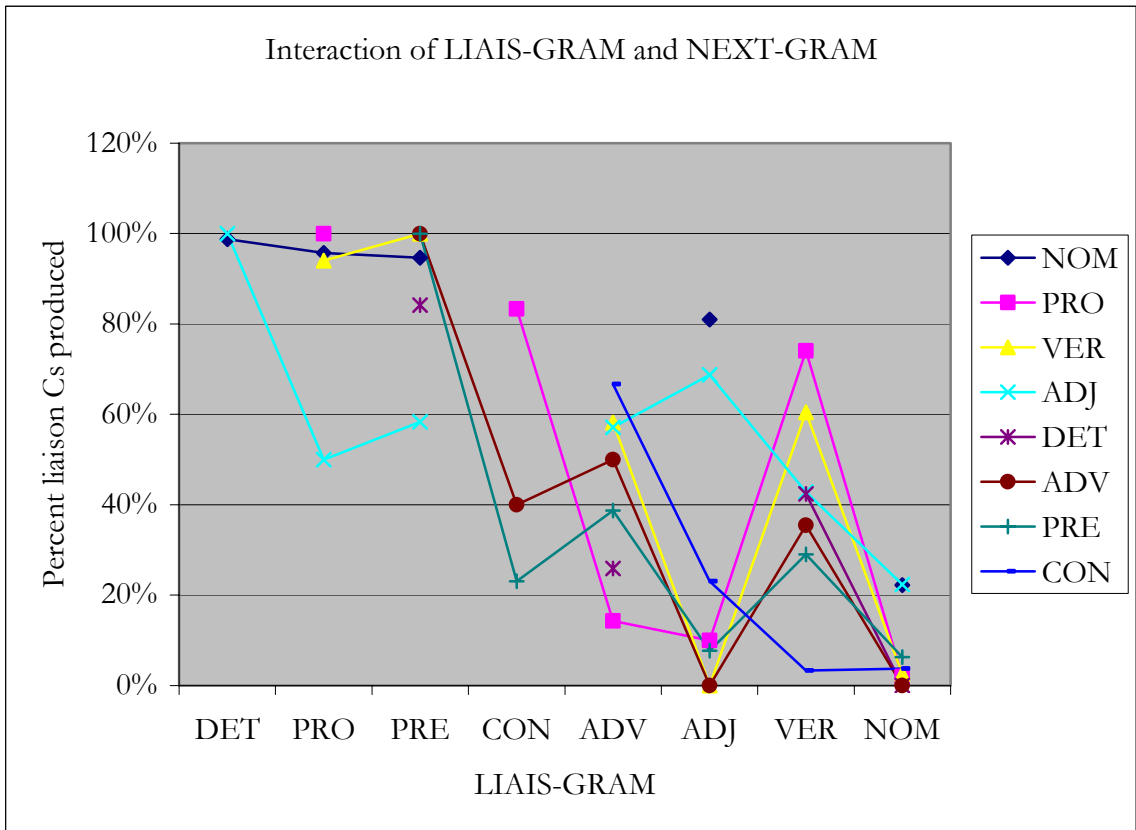


Figure 27 – Liaison valence as a product of liaison word grammatical class and next word grammatical class

The final interaction of note holds between the frequency of the liaison word and that of the following word. Liaison word frequency is the strongest autonomous predictor of liaison production, and in Figure 28, we can plainly see an increase in liaison from left to right, as liaison word frequency increases. However, not all the lines, representing different next word frequencies, are equal. When the liaison word is very frequent or very infrequent, all next word frequencies cluster around 95% and 15% respectively. But when the liaison word has a middle-ranged frequency, more frequent next words give globally less liaison than less frequent ones.

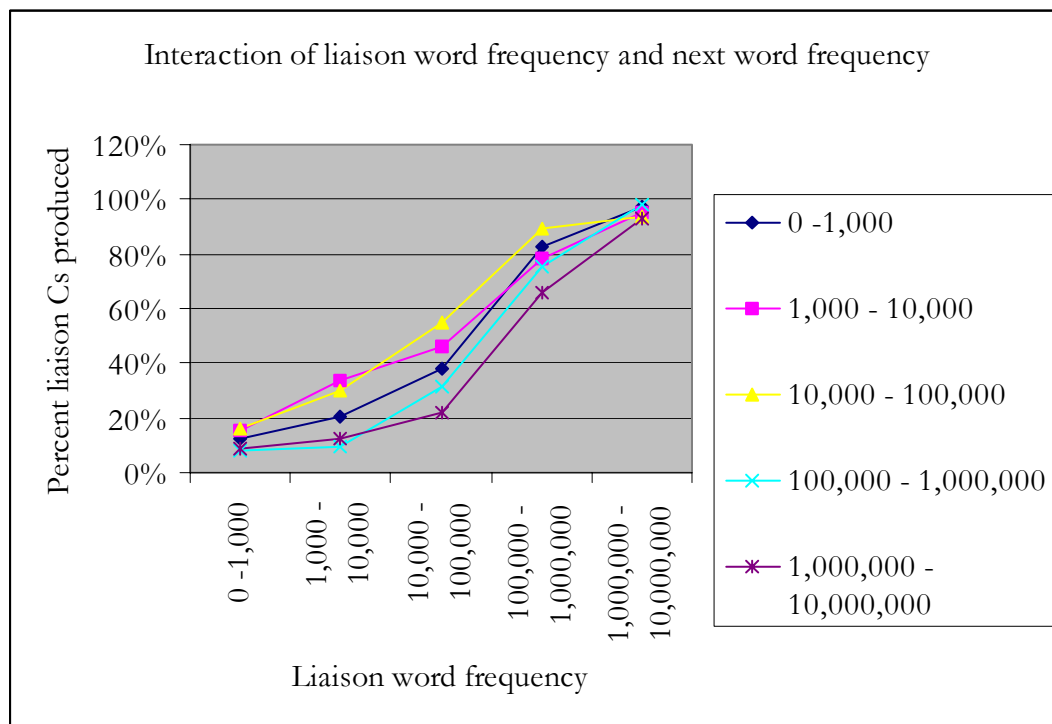


Figure 28 – Liaison valence as a product of liaison word frequency and next word frequency

This finding goes to support the hypothesis that word-pair frequency is relevant to liaison. We might suspect that instead of deriving from the relative frequency of two words,

the probability of liaison derives mainly from the liaison word, in addition to a subtle contribution on the part of the next word's frequency. If this explains middle-frequency word – pairs' behavior, then we should expect that in cases where the liaison word's frequency does not strictly constrain liaison, that more frequent next words should evoke more liaison. However, what we see in the interaction between LIAIS-FREQ and NEXT-FREQ is the reverse trend; with middle-frequency liaison words, the frequency of the following word bears an inverse relation to the percent of liaison consonants produced.

7. Final note

Frequency of a liaison word is the best predictor of whether a liaison consonant will be deleted or not. But frequency could play an even more fundamental role in the distribution of liaison segments than this. It may have a role in determining what expectation individual language users have about the language of their community.

Frequent pairings of liaison with other features of an utterance might be internalized by hearers. If older speakers utter more liaison consonants, perhaps individual hearers are aware of this fact. Perhaps they even come to depend on it, expecting that an older speaker will utter more liaison consonants. By making this assumption, a hearer will be able to establish more reliable prior beliefs about the phonological forms that will come out a speaker's mouth, facilitating the understanding task. If associations between liaison and phonological or extra-phonological features of a string are robust, that is, statistically significant, perhaps speakers have internalized these correlations as well. Perhaps they respond differently to utterances in their language when their expectations about pairings between liaison and other features are met than when they aren't. Finally, when factors

interact in their effects on liaison, perhaps hearers pick up on the trends and use them for the same purposes.

These hypotheses are tested in an experimental fashion in the following chapter.

Chapter 4. Individual perception of variability

Outline

1. Introduction
2. Experiment 1: Autonomous factors
3. Experiment 2: Interactions between factors
4. Discussion

You know how it is when you go to be the subject of a psychology experiment, and nobody else shows up, and you think maybe that's part of the experiment? I'm like that all the time.

Steven Wright

1. Introduction

In the language of the community, liaison use correlates with a range of phonological, prosodic, syntactic, morphological, and social factors. To what extent do individual language users make use of these correlations during language processing? To what extent do they make use of knowledge of interactions between these factors? If they do make use of this knowledge, what is the relationship between the degree of the correlations in the speech of the community and their effect in perception?

This chapter uses the liaison corpus as the basis for a linguistic perception task, that can formulate answers to these questions. The idea is to test the perceptual status of the correlations that emerged in Chapter 3, by taking speech tokens from the corpus and playing them to subjects, soliciting some sort of quick response. The responses the subjects give and how long it takes them to give them are indications of those subjects' linguistic processing.

We then can statistically test the subjects' responses in a manner similar to the analysis of the

original liaison corpus. This gives us insight not only into the factors impinging on the subjects' linguistic processing. It also allows us to assess the relation between the distributional facts and the processing effects of those factors.

Two experiments tested autonomous and interacting factors using different methods.

2. Experiment 1: Autonomous factors

Background

The statistical analysis of the FLC described above demonstrated that a number of factors probabilistically influence the production of a liaison consonant. Among these are liaison identity, liaison frequency, and liaison length. You will remember from the previous chapter that liaisons in *r* are more infrequently produced than those in *l*, and that increasing frequency raises the probability of liaison production, while increasing length lowers it. These three factors were selected to test the following experimental questions:

1. Is the processing of liaison words affected by the production or omission of a liaison consonant?
2. Is the processing of liaison words affected by the various factors that influence liaison production in the language of the community?

Method

In order to elicit quick responses to speech stimuli, the experimental setup was the following cross-modal matching task. Subjects were first presented with a written PRIME stimulus (for 600msec, in the middle of a digital computer screen), which was composed of two words. Immediately afterward, they heard a spoken TARGET stimulus. The subjects' task was to decide, as quickly as possible, whether the written words stimulus and the spoken stimulus were “the same” or “different”.

For all of the test stimuli (described below), the preceding, written PRIME forms were simply written representations the two words, and the following, spoken TARGET forms were spoken versions of those same words. For all of the filler stimuli, the preceding written PRIME form differed in either the first or the second word from the spoken TARGET form. In this way, all the targeted experimental stimuli were in the “same” condition. All of the filler stimuli were different from their written forms, and thereby served to balance the data set. All test and filler stimuli had potential liaisons, and half of each had their liaison consonants produced.

63 subjects were recruited from the campus community of the University of Lausanne and were compensated with 10 Swiss Francs, at the time of experimentation equal to about \$6, for their participation in this and the second experiment. They ranged in age from 19 to 38 and all self-reported as native speakers of Swiss French. Their regional origins were spread throughout French-speaking Switzerland, but most were born in the Vaud canton, where Lausanne is located. Subjects were told that they were participating in a perception experiment, and, after detailed instructions on the experiment process, were instructed to respond as quickly as possible with their responses to the stimuli. Before the

actual recorded part of the experiment began, they engaged in a training session, made up of twelve stimuli very similar to the test and filler stimuli. Both this experiment and Experiment 2, covered in the next section, lasted no more than ten minutes.

Stimuli

In constructing the test stimuli, 48 cases of liaison words along with the word following them were selected from the FLC. They were grouped into three classes.

In the first, pairs of words were selected according to the frequency of the liaison word. The stimuli were grouped into five frequency categories, as in Figure 1 below. A total of 20 tokens were selected. These didn't fall evenly into the five classes since there were not enough cases of liaison not being produced in extremely frequent liaison words.

	1-10 ³	10 ³ -10 ⁴	10 ⁴ -10 ⁵	10 ⁵ -10 ⁶	10 ⁶ -10 ⁷
Liaison	2	3	2	2	1
No liaison	2	3	2	2	1

Figure 1 - Numbers of tokens in each of five frequency conditions and two liaison valence conditions

To control for morphological status and liaison word grammatical class, all the tokens in the frequency-varied group had a liaison word which was an inflected verb. This also made it easier to find following words that shared grammatical classes. All tokens were also inflected for the third person. Tokens in the *liaison* and *no liaison* conditions were balanced for other factors: speaker age, liaison word length, following word length, following word grammatical

class, following segment class, following word frequency, pause presence, and preceding segment class. In other words, the liaison and no liaison groups were evenly balanced for each of these factors.

In the second class of stimuli, whether not the liaison consonant was produced, the *liaison valence*, was cross-cut by the identity of the liaison consonant. The first word of each stimulus pair ended with a liaison consonant which was either *t* or *r*, and the liaison could be produced or not. This yielded four conditions for the 16 tokens to fall into (Figure 2).

	Liaison identity	
	/t/	/r/
Liaison	4	4
No liaison	4	4

Figure 2 - Number of tokens, categorized by liaison identity and liaison valence

Just as in the liaison word frequency stimuli, these liaison identity test stimuli were matched for a number of features. These were liaison word length, following pause, preceding segment class, liaison word grammatical class, morphological status, following word length, and speaker age.

In the third and final group of stimuli, liaison valence was varied, as was the length of the liaison word. Only monosyllabic and disyllabic words were included, so combined with the two valence conditions, these yielded four combined conditions (Figure 3).

Liaison word syllables		
	1	2
Liaison	3	3
No liaison	3	3

Figure 3 - Number of tokens, categorized by liaison word syllables and liaison valence

Controls were effected for this group on the length, grammatical class, and frequency of the liaison word, as well as the speaker’s gender and age, and the identity and morphological status of the liaison consonant.

These three groups of stimuli yielded a total of 48 test stimuli. Along with these, 48 filler stimuli, which were matched pairwise with the test stimuli for frequency and length, were selected.

I manually extracted all the test and filler stimuli from the audio files in the FLC. Care was taken to select as much of the stimulus as was possible, without including material from the surrounding context.

Results

The experiment’s setup afforded only one reasonable metric for assessing the processing differences that might result from relatively more likely or less likely combinations of liaison application and other aspects of a token. This metric is response classification, the measure of how frequently subjects selected “same” or “different”. Unfortunately, reaction time, the time subjects took to respond to the stimulus, was not informative, since tokens varied significantly in their length.

For the cross-modal matching task, subjects were asked to identify whether or not the words that appeared on the screen were the same as the words they subsequently heard. Aside from one subject, about half of whose responses were on a key with an undesignated value, all subjects responded only by pressing a button indicating identity or one indicated difference. This task was under strict time pressure - subjects were instructed to respond as quickly as possible. On average, the time between the end of the sound token and the subject's response was a short 632 milliseconds.

Liaison identity

We'll first consider the identity of a liaison consonant and its potential effect on word processing. You will remember from the discussion in Chapter 3 that liaison consonant identity played a central role in determining the degree to which liaison consonants were produced. In particular, the two consonants considered in this experiment, /t/ and /r/, contrasted sharply in their degree of liaisability. As shown in Figure 4, in only 6% of its occurrences was liaison /r/ produced, while /t/ was articulated in a hearty 48% of its potential appearances.

Liaison consonant identity				
	/r/	/z/	/t/	/n/
% produced	6%	46%	48%	58%

Figure 4 - Liaison valence as a product of liaison identity

The question at hand, then, is the following: Are subjects more likely to respond that a spoken utterance matches its written form if a liaison /t/ is produced than if a liaison /r/ is produced, relative to their unproduced counterparts? As we can see from the results in Figure 5, the answer is yes. First, let's consider responses to /t/. Here, subjects more often responded 'same' when the /t/ was produced than when it wasn't. By contrast, /r/ was more likely to give rise to a 'same' response when unproduced than when produced.

	Consonant	
	/r/	/t/
Liaison	75%	93%
No liaison	91%	83%

Figure 5- Percent of 'same' responses as a product of liaison production and liaison identity

But we also need to establish that this effect is statistically significant. Luckily, we can use logistic regression, the precise statistical tool described in Chapter 3. Remember that logistic regression involved the construction of a regression model for a binary dependent variable as a function of a number of independent variables and interactions between those variables. In the present case, we would like to establish the degree to which we can be certain that there is a significant interaction between consonant identity and liaison valence. That is, the null hypothesis we need to disprove is that any effects of liaison valence and consonant identity are entirely autonomous - nothing about being an /r/ has an effect on how liaison production or non-production impacts subjects' judgments.

In both a forward stepwise and a backward stepwise selection test, the interaction between consonant identity and liaison valence was included in the regression model and

also yielded a significance less than 0.01. In other words, this interaction effect was highly significant. Interestingly, so was the autonomous effect of consonant identity, but not that of valence. So although valence made no autonomous contribution to subjects' responses, it did when taken into consideration along with the identity of the liaison consonant.

Liaison word frequency

In the liaison corpus, a direct relationship was shown between the frequency of a liaison word and the likelihood that its liaison consonant would be produced. As shown in Figure 6, every frequency class yields a higher liaison production average than its less frequent neighbors. Did subjects' categorization of test tokens reflect an expectation of this state of affairs? More precisely, were tokens with more frequent liaison words more likely to be categorized as the same as their written counterparts when their liaison consonant was produced, relative to tokens with less frequent liaison words?

	Liaison word frequency				
	0-10 ³	10 ³ -10 ⁴	10 ⁴ -10 ⁵	10 ⁵ -10 ⁶	10 ⁶ -
% produced	9%	16%	35%	77%	96%

Figure 6 - Liaison valence as a product of liaison word frequency

Figure 7 presents the percentages of 'same' responses as a function of the frequency of the liaison word and the valence of liaison. It is less easy to pick out the interaction between these factors visually than in the previous case, because there are more values for the frequency variable. Grossly speaking, though, we see that for the least frequent liaison

words, the two columns on the left, the percent of ‘same’ responses is greater if the liaison consonant is not produced than if it is produced. However, the most frequent liaison words, in the two columns on the right, yielded the highest percentage of ‘same’ responses when liaison was produced. The middle-range-frequency words, the center column, gave rise to approximately equal responses whether or not the liaison consonant was produced. In other words, there appears to be an interaction at work between liaison word frequency and liaison valence.

	Liaison word frequency				
	1-10 ³	10 ³ -10 ⁴	10 ⁴ -10 ⁵	10 ⁵ -10 ⁶	10 ⁶ -10 ⁷
Liaison	74%	89%	98%	97%	95%
No Liaison	86%	93%	96%	61%	90%

Figure 7 - Percent of ‘same’ responses as a product of liaison valence and liaison word frequency

Step-forward and step-backward logistic regression tests confirmed these results. In both, the interaction of these two variables was included in the final model, with a significance smaller than 0.05. That is, the interaction was significant. The same was true of the autonomous effects of liaison valence and liaison word frequency. There was one blip in the experimental results, a particularly low percentage of “same” responses when no liaison was produced in the second-most frequent words. I don’t have any explanation for this aberrance.

Liaison word length

The third factor on liaison that was tested for was the length in syllables of the liaison word. Just as in the previous two cases, we are looking for an interaction between liaison valence and liaison word length as they effect subjects' 'same'/'different' responses. In the liaison corpus, shorter liaison words yielded more liaison than longer words (Figure 8), so if the pattern shown in the previous two examples continues, subjects should have responded 'same' more frequently when the liaison word was short and liaison was produced than when it was not produced, as compared with longer words.

	Liaison word syllables			
	1	2	3	4+
% produced	74%	24%	6%	6%

Figure 8 - Liaison valence as a product of liaison word length in syllables

Figure 9 shows the percentages of 'same' responses as a product of these two factors. Once again, a clear interaction effect is visible. With monosyllabic words, liaison production led to greater 'same' categorization, while disyllabic words showed the reverse trend. Tokens with disyllabic liaison words were most categorized as the same as their written counterparts when a liaison consonant was not produced.

	Syllables	
	1	2
Liaison	94%	79%
No liaison	87%	94%

Figure 9 - Percent of 'same' responses as a product of liaison valence and liaison word length

Once again, these results were submitted to statistical evaluation. And once again, both a forward stepwise and a backwards stepwise model-building algorithm included the interaction terms as a factor, with a significance smaller than 0.01. Also included in both was the autonomous factor valence, with a significance smaller than 0.05.

In response to the experimental questions designated above, three conclusions emerge from the results reported above.

1. Native French speakers process liaison words differently when the liaison consonant is produced than when it isn't. In condition after condition, liaison valence appeared as a significant factor in subjects' match responses. This suggests that liaison valence is a psychologically relevant linguistic phenomenon.
2. The processing of liaison words is also affected by the various factors that influence liaison. This is indicated by the significance of interactions between liaison valence and liaison identity, liaison word frequency, and liaison word length.

In the next experiment, I addressed the psychological status of social variability effects on liaison, as well as interactions between these and other factors.

3. Experiment 2: Interactions between factors

In the previous experiment, it was demonstrated that variables that correlate with liaison production in the language of the community are translated by the individual perceiver into processing expectations. Now we'll move on to social correlates of liaison and interactions between these and other factors, picking out one particularly salient interaction.

Background

The FLC shows autonomous social factors on the production of liaison in liaison words, as well as interactions between these and other factors. One of the more significant of these interactions is the one that crosses the age of the speaker with the identity of the liaison consonant.

The methodology used in Experiment 1 was quite effective in picking out the degree to which perceivers access liaison-external factors when processing language. But this method would prove less effective were it to be applied to interactions between factors on liaison. In Experiment 1, we were testing the effects of interactions between liaison valence and other factors on responses to a forced-choice task. But if we were to use the same method for investigating effects of interactions between two factors on liaison valence, then the experimental setup would need to test for the significance of a three-way interaction. That is, the test would be whether the three-way interaction between liaison valence, speaker age, and grammatical class was a significant factor in determining a subject's sameness response.

Such an experimental setup would push up the number of data points required to reach significance levels by a factor of three, because now instead of four conditions, we

would be testing twelve conditions (four times the three age conditions). Two other methodological problems would emerge from using the same methodology as in Experiment 1. First, sociolinguistic studies of the effects of age on linguistic variables tend to split speaker age into at least three groups, approximating those selected in the corpus study described in Chapter 3. However, the same-different task used in Experiment 1 only allows for a two-way distinction. Thus, if a S responds that a token and its prime are different, we can't conclude which age group that S classified the speaker as belonging to. Second, listeners are relatively adept s at discerning speaker age from relatively little data. However, there's little reason to think that they are able to instantaneously determine a speaker's age from only two successive words, especially since some tokens were as short as 200 msec. In fact, as I'll show below, subjects did a particularly poor job of guessing speakers' ages.

These methodological complications call out for a different approach to data collection for the present task. Following Bates et al. (1996), I decided solve this problem by taking one of the potentially complicating independent variables, speaker age, and using it as the dependent variable. By changing the task so that subjects were now asked not to make similarity judgments, but rather to determine the age of the speaker, I could directly access their age judgments. This solved the three problems discussed above.

1. Only liaison and grammatical class and their interactions were relevant independent factors, so only four conditions needed to be tested (although we'll see below that this assumption was not quite accurate).
2. There could now be three possible values for the dependent variable, age, just as there were three values for age in the liaison corpus study.
3. Our analysis could gain direct access to the subjects' age judgments for each token, as this was their very response.

Adopting this approach allowed the following questions to be addressed.

1. When making social judgments about speakers, do listeners make use of the speakers' production of liaison consonants?
2. Are these judgments contingent upon the interactions between the production of the liaison consonant, and other variables, like the grammatical class of the liaison word?

Stimuli

36 liaison word and following word pairs were selected from the FLC. These stimuli varied along three factors - the age of the speaker, the grammatical class of the liaison word, and whether or not the liaison consonant was produced. The 36 stimuli were distributed among the resulting eight conditions as follows (Figure 10).

	1924-1949		1950-1969		1970-Present	
	ADV/CON	VERB	ADV/CON	VERB	ADV/CON	VERB
Liaison	3	3	6	5	0	2
No liaison	1	1	6	7	2	2

Figure 10 - Number of tokens, categorized by liaison valence, liaison word grammatical class, and speaker age.

36 filler stimuli were roughly matched with the test stimuli for length and frequency. Stimuli were extracted from the FLC as discussed for experiment 1, above.

Method

The same 63 subjects performed an age-identification task, in which they heard a spoken stimulus, which was selected from either the test or the control stimuli. Their had to decide, as quickly as possible, whether the speaker of the stimulus was “young”, “middle”, or “old”.

Results

Just like responses to the cross-modal matching task, responses to a forced-choice age-judgment task can provide a window into cognitive processing. Under extreme time pressure, subjects make semi-automatic judgments, which cannot be due to conscious processing. On average, subjects responded 835 msec after the end of the stimulus. It’s not surprising that the reaction time to this task was somewhat longer than that to the cross-modal matching task. In the cross-modal matching experiment, for all of the test stimuli,

both of the words subjects heard were primed by a preceding visual stimulus. But there was no priming for the age-judgment task.

In Chapter 3, we surveyed interaction effects in the liaison use of the French-speaking community. The interaction of age by grammatical class, the interaction tested in this experiment, was shown to be subtle but significant. As seen in Figure 11, the proportion of liaison consonants produced in adverbs to those produced in verbs is greater for young speakers than for middle-aged and older speakers. So if subjects have internalized this tendency, then they should judge speakers to be relatively younger if they produce liaison with adverbs rather than not, relative to verbs produced with liaison versus those without.

		Liaison word grammatical class	
		Adverb	Verb
Speaker age	Old	52%	51%
	Middle	38%	39%
	Young	43%	37%

Figure 11 - Liaison valence as a product of grammatical class and speaker age

Before we get to the results, one additional complication needs to be clarified. As shown in Figure 10 above, the experimental setup actually included speakers of each of the three different age classes. This means that there might be some effect of the speaker's actual age on the perceived age, as reported by subjects. In fact, there was a strong effect of age, as Figure 12 attests (with this main effect of speaker age being highly significant, >0.01 , in all the logistic regression test it was placed in). But it's clear from these results that subjects had a great deal of trouble in consciously identifying speaker age. Young speakers were identified

much less frequently as ‘young’ than were their middle-aged counterparts (about one third as often). Old speakers were much less frequently categorized as such than were young speakers (about half as often).

Perceived Age	Actual Age		
	Young	Middle	Old
‘Young’	12%	36%	12%
‘Middle’	34%	51%	61%
‘Old’	53%	13%	27%

Figure 12- Perceived age by actual age

So actual age will only serve to complicate the analysis of the interaction between liaison valence and liaison word grammatical class on age judgments. In the discussion that follows, I will first discuss the statistical analysis of all of the results, including the complicating factors of age, and then I will move on to a more detailed analysis of the effects of the two test factors in a single age class, the middle-age.

Given that the speaker’s age has a significant effect on subjects’ judgments of age, we need to include age as a factor in any analysis of the other two factors that we conduct. In other words, we need to test the significance of not only the interaction of liaison word grammatical class and valence as it affects age judgment. We also need to test for the significance of the three-way interaction between grammatical class, valence, and speaker age, as it affects perceived age. Three-way interactions are extremely difficult to pick out visually from a table or graph, since they involve differences between relations between

relations. The Figure in 13 is characteristic of the level of complexity such visual representations are apt to take on.

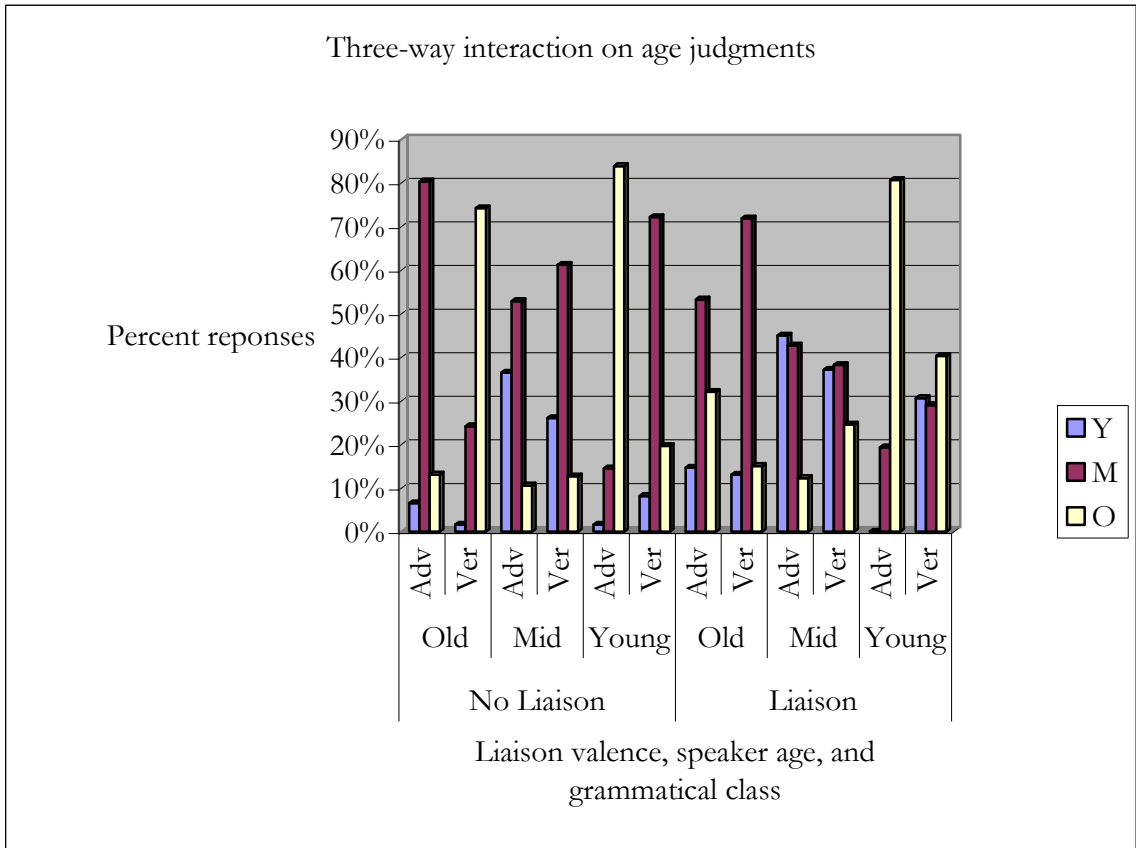


Figure 13 - Interaction of speaker age by liaison valence by liaison word grammatical class

Much more helpful is to test whether this interaction is truly a statistically significant one, or whether it's just a random fluctuation of the data. Unfortunately, logistic regression is no longer exactly appropriate for the analysis at hand. The dependent variable we are now considering is not, as required by logistic regression, a binary variable, but rather a ternary one, with three age categories as response possibilities. One solution to ternary dependent variables was adopted by Cedergreen (1973) in an analysis of Panamanian Spanish /s/,

which can be realized in three ways, [s], [h], and \emptyset . There, the multivalued dependent variable was separated into two dichotomous variables, one relating [s] to its spirantized cousin [h] and another deleting [h]. In similar fashion, we can separate the dependent variable in the present work into two judgment processes, both as deviations from the ‘middle’ response. ‘Middle’ is not only conceptually the default response and the one requiring the least commitment on the part of subjects, but it is also the mode of subjects’ responses, as well as the age category they responded most accurately to.

I performed logistic regression on the age judgment results in two sessions, separating them into ‘middle’ and ‘old’ responses on the one hand, and ‘middle’ and ‘young’ responses on the other. Let’s first look at the comparison of ‘middle’ and ‘old’ responses. Here, in both forward and backward selection procedures, the three-way interaction of speaker age by liaison valence by liaison word grammatical class was included, and had a significance rating less than 0.01. Age, grammatical class, and the interaction of age by grammatical class were also included in the model by both methods, where they all had a significance less than 0.01.

The results of the ‘young’ versus ‘middle’ analysis were slightly different. The backwards stepwise function included in the model the three-way interaction of speaker age by liaison valence by liaison word grammatical class, as well as all those included in the ‘old’ and ‘middle’ model, with the addition of the interaction of valence by grammatical class. But the forward stepwise procedure swapped valence for grammatical class and did not include the three-way interaction in question.

In sum, relative to ‘middle’ responses, ‘old’ responses were governed, among other factors, by a three-way interaction between speaker age, liaison word grammatical class, and liaison valence, as we might have guessed from the speaker age effects discussed above. It

was not, however, subject to an interaction effect between grammatical class and valence, as we originally predicted. We can be slightly less certain about the status of the three-way interaction as concerns deviations in judgment from ‘middle’ to ‘young’. In this case, one statistical test accords the three-way interaction a strong significance, while the other leaves it out. But in both cases, the interaction of interaction by grammatical class was significant.

Since the actual age of speakers seems to have been confounding all of these analyses, both autonomously and in interactions with the other factors, it seems essential to control for speaker age by taking a single age slice and investigating subjects’ responses to those tokens. Middle aged speakers serve as the best baseline case for the reasons discussed above (numerical prevalence, perceived normalcy, and most accurate identification).

And in fact, when we consider only the middle aged speakers, precisely the interaction we expected emerges (Figure 14).

	Valence			
	Liaison		No Liaison	
	Grammatical Class		Grammatical Class	
	Adverb	Verb	Adverb	Verb
Young	45%	37%	36%	26%
Middle	43%	38%	53%	61%
Old	12%	25%	11%	13%

Figure 14- Age judgments for middle-aged speakers as a function of liaison valence and grammatical class

The proportion of ‘young’ judgments to ‘old’ judgments in adverbs where liaison is produced is much greater than the proportion of ‘young’ judgments to ‘old’ judgments in verbs where liaison is produced. Similarly, adverbs with produced liaison consonants garnered many more ‘young’ judgments than did their counterparts with liaison consonants that went unproduced.

This interaction effect is probably easier to pick out visually. Figure 15 shows the age judgments to middle-aged speakers as function of liaison valence and liaison word grammatical class. Compare ‘young’ and ‘old’ responses in the liaison produced columns, the two clusters on the right. In verbs with produced liaison consonants (the rightmost column), ‘young’ and ‘old’ responses are much closer than they are in adverbs with produced liaison consonants. This trend is true to a significantly lesser degree in adverbs and verbs without liaison consonants produced (the two right-most column clusters). There, while verbs without liaison garner fewer ‘young’ responses than their adverb counterparts, they also gain slightly in the realm of ‘old’ responses. It’s the difference in the degree of the difference between the liaison and no-liaison responses that constitutes a potential interaction effect.

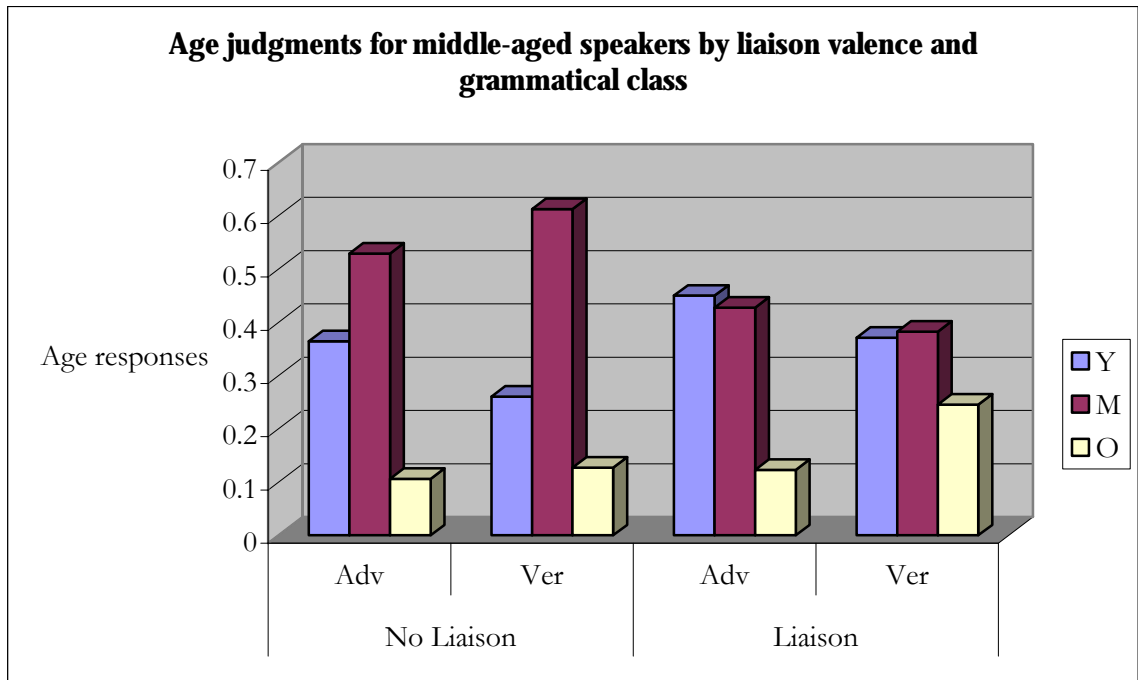


Figure 15 - Age judgments for middle-aged speakers as a function of liaison valence and grammatical class

The statistical significance tests of these data were conducted in the same manner as those for the data that included age variation, above. That is, ‘young’ and ‘old’ responses were separately compared with ‘middle’ responses. Each was subjected to both a forward and backward stepwise model building procedure. Thus, there were four tests of significance. In three of the four, all but the backward ‘young’ to ‘middle’ comparison, the interaction of valence by grammatical class was deemed significant, less than 0.05. In the other case, it was not included in the model. Grammatical class and valence were also included in a subset of the models.

To conclude, whether the speaker’s age is controlled for or included in a statistical analysis, the majority of logistic regression models (six of eight) include a term representing

the statistically significant interaction between liaison valence and liaison word grammatical class.

4. Discussion

The results discussed above suggest strongly that individual listeners are at least unconsciously aware of the correlations between liaison on the one hand and a number of other phonological, syntactic, and social factors on the other. The cross-modal matching task provided a forum for liaison identity, liaison word length, and liaison word frequency to prove their status in the individual's linguistic system. The age perception task added to these results a demonstration of the relevance of interactions between factors on liaison for language processing.

The existence of subtle, probabilistic correlations between sound and other domains of knowledge runs counter to the assumptions we normally make about language. Linguistic theories often view the human language capacity as modular, deterministic, and discrete. In Chapter 6, I seek explanations for why the linguistic system would in fact be probabilistic, cross-modal, and interactive, from a neural perspective.

The psychological reality of probabilistic effects on liaison and other phonological processes from phonology, syntax, semantics, morphology and the social realm, and of interactions between these effects, imposes strict requirements on what our models of linguistic knowledge can look like. In the next chapter (Chapter 5), I discuss these constraints and outline a computational model that is capable of meeting the challenge.

Chapter 5. Probabilistic computational models of non-categoriality

Outline

1. Introduction
2. Belief Networks
3. Liaison in a Belief Network model
4. Phonaesthemes in a Belief Network model
5. Properties of the models

As far as the laws of mathematics refer to reality, they are not certain; and as far as they are certain, they do not refer to reality.

Albert Einstein

1. Introduction

As we have seen in the preceding chapters, language users put to use a broad range of linguistic and extra-linguistic knowledge when they process phonological input. Some of this knowledge is non-categorial in nature, such as the correlations between liaison valence and word length or grammatical class. Among these correlations are statistical ones between sound and meaning, as demonstrated by non-productive English phonaesthemes in Chapter 2. These sorts of non-categorial information from different sources can be brought to bear in a combined fashion, as demonstrated by interactions between social variables and linguistic ones in Chapter 4.

These aspects of language processing have previously fallen outside the usual realm of linguistic inquiry, because they run counter to the conventional assumption that language is a deterministic, modular, and rule-governed system. From the classical perspective, the

behaviors described above seem non-optimal. Why would we language users needlessly burden ourselves with detailed knowledge of statistics where deterministic rules would suffice? Why mix different types of knowledge when a modular system would be cleaner and simpler? Why complicate processing by including interactions between factors?

While these properties are quite difficult to explain or model from the perspective of a deterministic, modular, rule-based modeling enterprise, when language is considered in its biological context, they are anything but surprising. There are clear indications from studies of learning and processing in the brain that probability, cross-modularity, and schematicity are anything but aberrant. Quite the opposite, in fact - given brain facts to be discussed in the next chapter (Chapter 6), it would be surprising if we didn't employ knowledge of probabilistic, interacting correlations across domains of linguistic and extralinguistic knowledge.

A growing body of recent work has progressively begun to document the relevance of probability in phonology (e.g. Anttila 1997, Boersma and Hayes 2001, Frisch 2001, Guy 1991, Hayes 2000, Nagy and Reynolds 1997, and Pierrehumbert In Press). Intriguing and challenging new models have emerged from this research, including gradient versions of OT (Boersma and Hayes 2001), for example. None of these linguistic models, however, provides a sufficiently general architecture to allow the range of influencing factors and interactions between those factors to be captured. Moreover, aside from a brief account of one version of OT (Prince and Smolensky 1997), there has been no attempt that I know of to ground these models in neural structure, either in terms of their architecture or their behavior. They don't include clear computational correlates of the neural properties we need to incorporate in order to explain the cognitive and linguistic behavior. Even if they did produce functionally identical behavior at the cognitive/linguistic level, they would nevertheless not

be able to capture the biological explanation for that behavior.

On the other hand, a large quantity of work has been done from the perspective of local connectists models of phonological processing (e.g. Dell and Juliano 1996). These models, being more psycholinguistically oriented, tend to deal with phonological and lexical access, rather than on morphophonology per se. Contrary to probabilistic linguistic models, however, their descriptions are usually well annotated with neural justifications. The model presented in this chapter is very similar in a number of respects to these localist processing models. Most centrally, both can capture probabilistic effects like the ones demonstrated in relation to phonaestheme and liaison processing, based on frequency of co-occurrence. Both are also computational representations capable of being neurally grounded.

In this chapter, I propose a computational model, based on a computational architecture known as Belief Networks. This type of model has two desirable properties. First, it can represent the subtle statistical knowledge humans demonstrate and can learn that knowledge from linguistic data. Second, as will be shown in Chapter 6, it can be directly mapped to the neural mechanisms responsible for those cognitive and linguistic behaviors it captures.

2. Belief Networks

A model capable of capturing the perception behavior that subjects demonstrated in response to phonaesthemes, liaison, and the numerous other linguistic structures reported on in the preceding chapters must meet certain functional criteria. It has to incorporate probability of some sort, allow for (the learning of) structured relations between pieces of knowledge of different sorts, and capture interactions, potentially between knowledge from

different domains. It must also be mappable to the neural level of explanation, such that it captures the biological explanations for these cognitive and linguistic phenomena.

Belief Networks (BNs - Pearl 1988, Jensen 1996) are a formalism that allows for the representation of potentially interacting probabilistic relations between any arbitrary set of pieces of propositional knowledge. They also come with a set of inferencing algorithms. They are thus able to capture probabilistic, cross-modal, interactive knowledge and the use of that knowledge. As we will see below, this formal mechanism is also groundable in neural structures. That is, there are neural structures that can accomplish the same computations, to which pieces of the BN architecture can be mapped.

BNs are a concise and powerful computational representation of uncertain propositional knowledge. Specifically, BNs consist of (1) a set of *nodes* representing propositions or variables, each of which has (2) a set of possible *values*, (3) *links* between causally-related propositions (where causation can be interpreted either ontologically or epistemically), and (4) *conditional probabilities*, specifying the probability of each value of every node given a value assignment to its parents. Using probability theory, inferences can be made about the probability of the value of any node given any (set of) observed values for any other nodes.

In a simple example, five propositions, each with multiple possible values, are represented by nodes (circles) in Figure 1. There is a node $\text{Rain}(t,f)$, which represents the belief that it has rained, and has two values, t and f , which in this case stand for true and false. The $\text{Rain}(t,f)$ node stands in a causal relation to $\text{Lawn_Wet}(t,f)$ (indicating that the lawn is believed to be wet), as indicated by the link connecting the two. Causality is indicated by the unidirectionality of the link; $\text{Rain}(t)$ causes $\text{Wet_Lawn}(t)$, and not the reverse.

Each node is also associated with a conditional distribution table (CPT), the small tables next to each node in Figure 1. CPTs are the specifications of the probabilities of each of a node's values along the y-axis, given possible values of its parents on the x-axis. Orphan nodes, those with no parents like Rain(t,f) and Sprinkler(t,f), have simple prior distributions that express the initial probability of each of their values. In the example in Figure 1, there network represented the generalization that there is a 0.3 likelihood that Rain will take the value true (that, that it will has rained at any point), and a 0.7 likelihood that it will be false. Similarly, next to the Sprinkler(t,f) node, we see that there is a 0.5 chance that Sprinkler will be true (meaning that the sprinkler has been on). The sum of probabilities for all the values of a proposition is always 1.

The relationship between two causally-linked propositions is encoded in the CPT of the downstream node. Briefly, the CPTs for nodes with parents specifies the probability of each of the vales of the child node, given the values of each of its parents. Lawn_Wet, for example, has two parents, and since each of them has two possible values, the probability of each of its two values is specified for the four possible causal states, thus giving eight possible configurations. In Figure 1, what we know about Lawn_Wet is that if Rain is true, and Sprinkler to be false, then the probability of Lawn_Wet(t) is 0.95, while if Rain is false and Sprinkler is true, then Lawn_Wet(f) has a 0.1 probability.

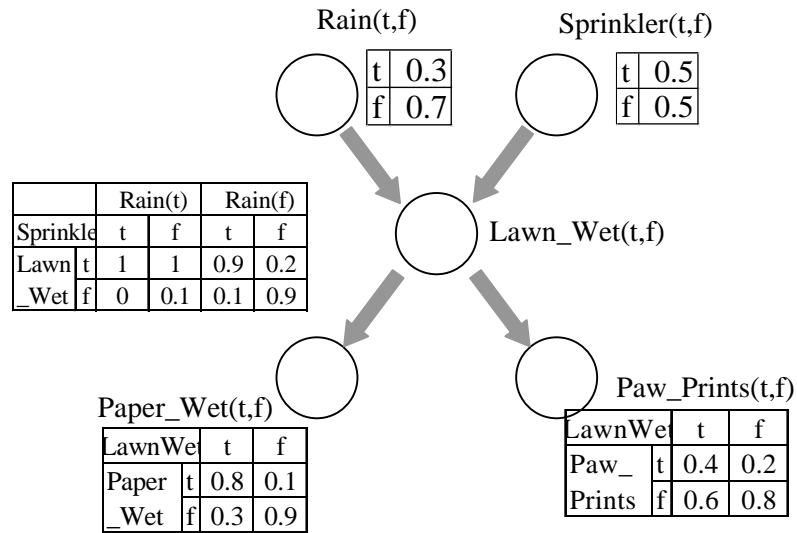


Figure 1 - A simple Belief Network, with five nodes, each with an associated conditional probability table (CPT)

The real interest of BNs lies not in their representational power but more importantly in their inferential power. Given a network and full set of CPTs as in Figure 1, inferences can be made given observations about the values of propositions. The simplest sort of inference is *causal inference*, which is the prediction of effects given that values for causes are observed. In the case of Figure 1, we might observe Sprinkler to be false and then let the inference algorithm determine the probability of Paper_Wet (meaning that the newspaper is wet) also being false. A second sort of inference, diagnostic inference, involves the propagation of evidence from an observed effect (a child) to an unobserved cause (a parent). For example, given that we observe Paw_Prints to be true (indicating that there are an animal's paw tracks on the living room carpet), we might ask what the probability is of Lawn_Wet being true. Finally, hybrid types of inference are possible: for example, we can ask what the probability of Lawn_Wet being true would be if Rain and Paw_Prints are both true.

How are BN's built? There are two aspects of BNs that can be learned automatically or constructed by hand: the structure (the nodes and links) and the CPTs. Historically, it has been most common for BN structure to be hand-crafted by experts in the domain being modeled. For example, I constructed the network in 14 on the basis of prior knowledge about the causal relationships between the variables it includes. Recently, though, a number of algorithms have been proposed to allow the automatic induction of BN structure on the basis of a large database. There are some problems with this. For example, if two variables are correlated, how can you automatically determine which is a cause of which? Even more imposing is the the problem of the number of possible solutions: for 5 variables, there are on the order of 100 million networks that could relate them, and for a network of 10 nodes, there are 10^{43} or ten million billion billion billion billion solutions. This makes it infeasible in practice to consider all possible solutions en route to a best solution. While various search methods have been proposed, including entropy methods, score metrics, simulated annealing, and genetic algorithms (Jordan 1998), none have rendered the induction of large-scale networks really practicable.

Learning the conditional probabilities of a BN is a much simpler process. Given a network, the distributions of co-occurrent proposition states are straightforwardly extracted from raw data, or from measures of the correlations between the various variables.

As we will see below, computational models of linguistic knowledge in a BN architecture are fully capable of representing the statistical relationships between any given variables, including interactions between independent variables on a dependent one. Additionally, the inferencing procedures that can be applied to a BN allow us to make predictions about human behavior on the basis of this structured statistical knowledge. This will become apparent through models of French liaison and English phonaesthemes.

3. Liaison in a Belief Network model

We'll start with the more complicated model - the one treating French liaison. You'll remember that our analysis of a corpus of spoken French demonstrated that a total of 13 autonomous and interacting factors influence whether or not a liaison consonant is produced. Moreover, during perception, language hearers are aware of these correlations and make linguistic and social judgments on the basis of them. In this section, we will see how these structured probabilistic relationships between variables can be learned by a BN model from corpus data, and how closely the inference procedures in such a BN model match the experimental perceptual results that we've seen.

We can capture the relationships between the various variables that affect French liaison in the language of the community with a BN model. In such a model, each variable will be expressed as a node, with a set of possible values assigned to it. For example, the dependent variable, the liaison consonant's valence, will be expressed as a node with two values, true and false. Liaison word grammatical class will similarly be expressed as a node, this one with 10 possible values, representing the ten most frequent grammatical class distinctions described in Chapter 3 above. (Abbreviation and interrogatives are omitted for clarity.)

Now, if we wished to automatically induce the structure of a BN best capturing the relationships between these variables, the algorithms would have to choose between an unearthly 10^{95} possible models. By comparison, there are estimated to be only about 10^{16} stars in the universe and about 10^{57} atoms in a medium sized stars, like our sun. So there are as many possible models here as there would be atoms in as many stars as our sun has atoms. Obviously, then, structure induction is functionally intractable for a full network.

I will proceed in two ways. In the first, I will take a small subset of the variables that account for liaison, and use a structure induction tool to attempt to discern the causal structure of the limited subset. I will then use an automatic learning algorithm to extract the conditional probabilities for each node of the small network and test the behavior of the network, to see how well it captures the generalizations we have drawn about the language of the community. In a second path, I will address the larger question of the full list of effects on liaison, by hand-crafting a BN that my prior knowledge indicates to be the most appropriate description of the relationships between the full set of variables.

For the learning task, I chose to investigate the induction of the structure capturing just three variables: liaison valence, liaison word grammatical class, and speaker age. These variables allow us to test the encoding of both autonomous and interacting effects.

A BN model of liaison is only as useful as the data it can capture. We can evaluate a BN model trained on a given corpus by asking it to perform inference, of any of the varieties described above (causal, diagnostic, or hybrid). We can then compare the predicted probabilities for a given node or set of nodes with the actual distributions of the values of those nodes in the corpus. But, like other statistical models, BNs capture correlations between variables straightforwardly as probabilities of their co-occurrence. This means that comparing the predictions of a BN model with the distribution in the very corpus that that model was learned from is not particularly informative - the match will always be very close.

The usual solution to this problem is to randomly split the relevant data into a training set and a test set. The BN is trained on the training set, the size of which can vary. The BN's predicted probabilities are then compared with the actual distributions in the test set. In this way, the statistical model constructed through training can be compared with novel data, and its power to predict those data can be ascertained.

I randomly selected 80% of the liaison corpus to act as the training corpus, with the remaining 20% constituting the test set. This allows us to ask the following questions:

- How do the BN's predictions about the autonomous (non-interacting) effects of speaker age and liaison word grammatical class on valence correlate with the test data?
- How do the BN's predictions about the interacting effects of speaker age and liaison word grammatical class on valence correlate with the test data?

The BN simulator used was the Tetrad III program, available from CMU (www.hss.cmu.edu/html/departments/philosophy/TETRAD/tetrad.html). Three variables were taken from the corpus for the structure induction task: valence, liaison word grammatical class, and speaker age. Aside from being shown the data themselves, the structure induction tool in Tetrad was only told that valence could not be a parent of either of the other two variables. On the basis of this constraint and the data, a BN was induced that had the structure in Figure 2. Given this network, a CPT for each node was then automatically learned from the training data.

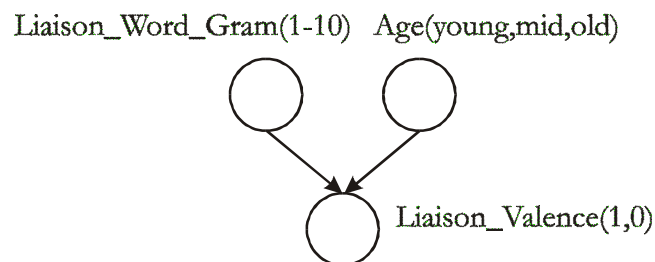


Figure 2 - BN structure induced from training set including three variables

Once a BN has been constructed and probabilities assigned to each of its nodes, we can compare the model’s predictions for the variables’ behavior with the patterns in the test set. Remember that the BN was trained on a separate training set. To test the BN’s predictive power, we can clamp nodes of the network to particular values. That is, we observe certain nodes to have particular values. For examples, we can tell the network that we have observed the liaison word’s grammatical class to be Adjective (in our notation, L_Gram is observed to have the value Adj or is “clamped” to that value). We can then ask the network to perform inference, and enquire about the predicted values of other nodes. So, given that the liaison word is a adverb, what does the network predict the probability of a liaison consonant being produced to be? Clamping the value of the Liaison_Word_Gram node at each of its values yielded the predicted probabilities for liaison valence shown in the top row of Figure 3. By comparison, the actual valence distributions for these same categories in the test set is shown immediately below the BN numbers, also in Figure 3.

	Lgram									
	Adj	Adv	Conj	Det	Noun	PropN	Prep	ProN	Verb	
Test	0.45	0.49	0.5	1	0.06	0	0.82	0.92	0.47	
BN	0.44	0.44	0.55	0.98	0.07	0.33	0.94	0.93	0.41	

Figure 3 - Predicted and observed liaison valence as a product of lgram

At first glance, the correlation between the BN’s predicted probabilities in the first row and the distributions observed in the test set (the second row) seems very tight, but we need some statistical measure of the closeness of this tightness. The degree of correlation between two sets of numbers is measured by the correlation coefficient. This measure varies

between -1 (inverse correlation) and 1 (direct correlation). The correlation coefficient for the BN and test set in Figure 3 is shown in column (a) of Figure 4. As shown there, when we graph predicted probabilities along the x-axis and observed distributions along the y-axis, we find that the data roughly describe a straight line, whose intercept is slightly below 0 and whose slope is approximately 1. In fact, if there were a complete correlation between the two sets of values, then the slope would be exactly one and the intercept exactly 0. It can be seen from Figure 4a that the predicted and observed values are closely correlated with slope and intercept nearly at 1 and 0 respectively.

	Clamped node(s)		
	(a) lgram	(b) age	(c) lgram & age
Correlation	0.94	0.99	0.77
Slope	1.03	1.78	0.87
Intercept	-0.06	-0.4	0.01
Average Error	0.04	0.02	0.07
Chance Error	0.24	0.04	0.27

Figure 4 – Goodness of fit between predicted and observed liaison valence as a product of (a) lgram, (b) age, and (c) lgram and age

These three measures don't tell us anything about the significance of the relationship, they just tell us whether there is one, and what its shape is. This is why a measure of significance, a t-test, is usually included with the correlation coefficient. In this case, though, a t-test isn't appropriate, because part of the t-test's metric for significance is the number of tokens it's provided. In the case we are considering, only nine numbers are being compared,

so the t-test can falsely conclude that the figures are highly insignificant, even though 500 tokens were actually evaluated. For this reason, the t-test result is not included in the following figures. As an alternative, we can complement the correlation measure with an average error measure, which is the average of the absolute value of the difference between the BN probability prediction and the test distribution by weighted condition. For example, the error for the adjective condition, the first column in Figure 3 above, is 0.01 (0.45 minus 0.44) and for the adverb condition it's 0.05. All these errors are weighted by the number of instances of that condition in the test set, and then are averaged. The average error is shown in the second to last row of Figure 4 above. For comparison, the average error of chance is shown in the final row of the same figure. This number is the result of assuming that the probability for valence is the same across all grammatical contexts. In the training set, the unconditional liaison valence probability was 0.49. The chance error, then, reflects the absolute value of the difference between 0.49 and the observed test value, by weighted condition.

Moving along to the effect of age on liaison valence, we can now clamp only the BN's speaker age node at a particular value, and ask it to predict the probability of a liaison consonant being produced. This test also yields strong predictions of the test data. Column (b) of Figure 4 shows the BN's probabilities and the distributions observed in the test set by age. The correlation here is nearly 1, and although the slope and intercept deviate from what we would expect for a perfect correlation, there are only three data points to consider. This renders the line-drawing task difficult. Notice, though, that the average error for valence is a miniscule 0.02, half of the chance error, 0.04.

Now that we've seen the predicted and observed effects of age and liaison word grammatical class alone, we can examine their combined effects, in column (c) of Figure 4.

While the differences between the BN and test values for each condition are greater than in the previous comparisons, these differences fall predominantly within the least frequent conditions. We can determine this by the small size of the average error, 0.07

We've established the strong correlation between the model's predictions for probabilities of valence and the actual observed distributions from the test set in both autonomous and interacting effects. We can now move on to examine the implications of the BN model for perception. Of particular interest are:

- How do the BN's predictions about age inference on the basis of valence and liaison word grammatical class correlate with the test data?
- How do they correlate with the results of the perception experiment?

The main difference between these evaluations and those described above is that in this case, the BN is not doing causal (forward) reasoning, but rather diagnostic (backwards) reasoning: reasoning about causes on the basis of observed effects. The BN's predicted probabilities therefore match the test data less well than do those described above. The BN's predicted results barely correlate with those of the test set (a coefficient of 0.7 is usually the minimum accepted for a positive correlation). The average error is similarly just barely better than that of the chance error rate (as seen in column (a) of Figure 5).

	Clamped node(s)		
	(a) BN vs. test set	(b) BN vs. experiment	(c) Normalized BN vs. experiment
Correlation	0.65	0.43	0.91
Slope	0.84	0.35	1
Intercept	0.05	0.22	0
Average Error	0.06	0.09	0.05
Chance Error	0.07	0.08	0.03

Figure 5 – Goodness of fit for predictions of age, between (a) the BN’s and the test set, (b) the BN and the experimental results, and (c) the normalized output of the BN and the experimental results.

Let’s move on now to a comparison between the BN’s predictions and the results of the perception experiment discussed in Chapter 4. As you will recall, subjects were asked to guess the age of speakers they heard, where two factors were varied: the liaison word grammatical class and the liaison valence. In essence, then, these speakers were performing the same sort of hybrid inference that the BN makes when those two variables are clamped. When we look at the relation between the predicted values for speaker age and the measured responses by subjects in the perception experiment the model seems to be even a worse predictor than it was of the test set, as seen in column (b) of Figure 5. There was no significant correlation between the two data sets, and that the average error was greater than that of chance.

But looking more closely at the cause of the differences between the BN’s predicted values and the subjects actual age evaluations, we see that there is a different overall distribution across the three speaker age categories (Figure 6).

	BN	Test	Experiment
Old	0.33	0.41	0.15
Middle	0.49	0.41	0.49
Young	0.19	0.17	0.36

Figure 6 - Average probability by age group for the BN, test set, and experimental responses

While the BN model follows the distribution in the liaison corpus at large, identifying half of subjects as middle aged, a third as young and only a fifth as old, the subjects in the perception experiment essentially reversed this trend in guessing speaker ages. They were about twice as likely to identify speakers as young than was the BN and about half as likely to label them old. One plausible explanation for this behavior is that since the subjects fell predominantly into the young class themselves (only 5 of the 63 being 25 or older), they were perhaps more likely therefore to identify speakers as falling into that same age range. Or perhaps the corpus was not drawn from an entirely representative sample. There could have been some self-selection in who contributed to the corpus: a selection bias towards older speakers. Because of this, the subject’s expectations could more closely reflect the real age distribution in Switzerland than does the BN, which only has access to the liaison corpus.¹

¹ There is some evidence against this second hypothesis. According to the Swiss Federal Statistical Office (<http://www.statistik.admin.ch/eindex.htm>), old speakers make up a much larger part of the Swiss population

Whatever the reason for it, this particular experimental task seems to have elicited response tendencies that diverged from the distribution of speaker ages in the corpus, and the population in general. We need to account for this experimental bias when comparing the predictions of the BN and the experimental results. Otherwise, we will run the risk of falsely rejecting the hypothesis that there is a correlation between the two groups of figures. One way to solve this dilemma is to normalize the output of the BN so that the probabilities of each of the age groups matches those of the experiment. In other words, multiply the output of the BN by a coefficient that serves to make the BN's average output per age group the same as that of the distribution of the subjects' responses. The values for the coefficient and the resulting normalized average responses for the BN are shown in Figure 7.

	Experiment average	BN average	Normalization coefficient	Normalized BN average
Old	0.15	0.33	1.92	0.15
Middle	0.49	0.49	0.99	0.49
You	0.36	0.19	0.47	0.36

Figure 7 – (a) Average probabilities for the experiment, (b) the BN's predictions, (c) Normalization coefficients (average percentage of total age responses in the experiment divided by the average percentage of age predictions by the BN), and (d) normalized values for the BN's predictions

than they do of the population of the corpus. This makes the difference between old speaker in the population and the age judgments by subjects in the perception task even more striking.

We can now compare these normalized BN probabilities with the observed experimental age judgments. This normalization process did away with much of the difference between the BN diagnoses and those of subjects. In fact, as column (c) of Figure 5 shows, the correlation between the BN's predictions and the human responses is extremely strong. It's stronger, for instance, than the one between the BN predictions and the distribution in the test set (column (a)). On the other hand, the chance error, in this case the average weighted absolute difference between the average responses in the age judgment task and the actual responses by condition, is still smaller than the BN's average error. In other words, it seems as though subjects' average age guesses are more closely to their actual values than the BN's predicted values.

The automatic construction of a BN model on the basis of a large corpus is useful for constructing a model of human language use and linguistic knowledge. For the particular task we have been looking at, while the BN makes very good predictions of test data on the basis of training data, the predictions for a human perception task have had to be normalized because subjects seem to have had a task-specific bias about speaker age. Nevertheless, once adapted to this different skewing of responses, the BN is able to capture human age judgments to a relatively good degree.

The structure of a network including all those factors found in Chapter 3 to be relevant to liaison cannot be induced, due to the size of the potential search space. Nevertheless, a BN that captures this large number of features can be constructed by hand. In such a network, all independent variables that directly (either autonomously or interactingly) affect the liaison consonant's production are represented as having direct causal links to the node representing liaison valence. Further structure can be included in such a network model. For example, as depicted in Figure 8, many of the factors on liaison

are correlated because they share common causes. For example, the identity of the liaison word and following word influence the orthography, grammatical class, length, and other aspects of the two words. The speaker's identity influences the variables age and education.

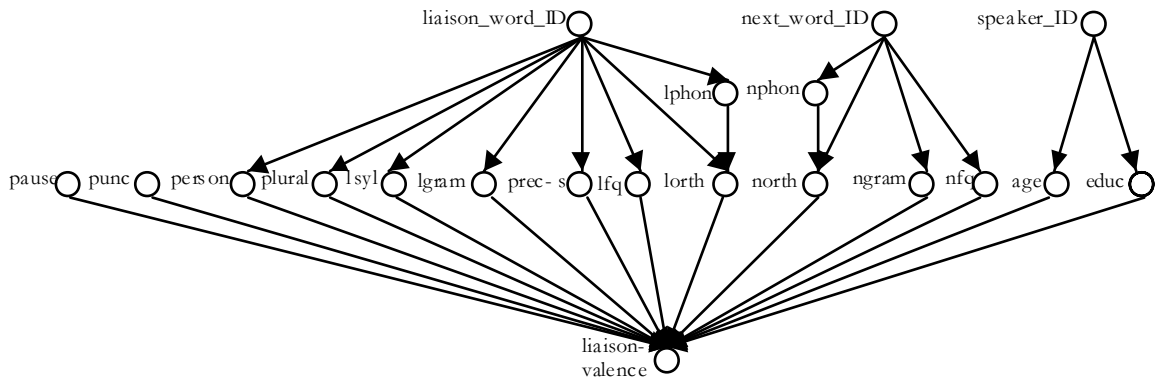


Figure 8 - Structured BN model of liaison

4. Phonaesthemes in a Belief Network model

In this section, I will describe a BN model for phonaesthemes. Three phonaestheme-related behaviors were discussed in Chapter 2. First, when presented with a definition for which there is an appropriate phonaestheme, people are more likely than chance to use that phonaestheme in a neologism. Second, given a novel word with a phonaestheme, people are more likely than chance to access a meaning appropriate to that phonaestheme. Third and finally, after hearing a phonaestheme-bearing word, a person more quickly identifies another word also bearing that phonaestheme than they do a word sharing a form and meaning that don't constitute a significant class of the lexicon. A simple BN model of phonaesthemes described below is able to easily account for the first two of these phenomena, while the temporal dimension must be incorporated into it to account for the third.

Phonaesthemes in a BN context involve fewer factors than liaison. Three factors are involved: the meaning to be expressed, the identity of the word, and the phonological identity of the onset. Each of these factors will be represented by a single node in a BN. Thus, I will represent meanings that can be expressed by words as values of a single Meaning variable. This variable can have values like ‘WET’ or ‘LIGHT’. Using this simplified representation makes meanings mutually exclusive. In actuality, this is not accurate, since multiple, compatible meanings can be co-expressed by a word (like *glisten*, which evokes both ‘WET’ and ‘LIGHT’). Word identities are similarly values of a single Word node. This is not a simplification at all - a given word really is selected in a particular context to the exclusion of all others. Finally, onsets are represented on a single node.

Rather than taxing the network as well as the reader’s attention by running a simulation that includes all the words in the lexicon that start with a particular set of onsets, I selected a subset that will adequately make the theoretical point. By selecting four words, *glisten*, *glitter*, *glee*, and *smile*, we can capture the following generalizations. There is a large class of words that start with *gl-* and have a meaning related to light (i.e. in this case “glisten” and *glitter*). There are other words sharing the onset, like *glee*, that have some other semantics, but each semantic class of these constitutes a small minority.²

Tetrad was instructed to build a network on the basis of these four words and their semantic and phonological values. It was also told that Meaning could not be a child of either Word or Onset, and that Word could not be a child of Onset. Tetrad proposed two potential models, shown in Figure 9. In both, Meaning links to Word and Word to Onset.

² For present purposes, the possibility that the Conceptual Metaphor HAPPINESS IS LIGHT is responsible for glee taking a gl- onset isn’t relevant, since in this simplified model, meanings are unique. However, the possibility that metaphors could play a role in structuring phonaestheme distributions is an intriguing one. (See Lakoff 1993 and Grady et al. 1999 for descriptions of HAPPINESS IS LIGHT.)

The difference is that in the one on the right, Meaning is also a contributing cause to Onset. That is, the meaning to be expressed directly affects the onset to be selected. While this is a reasonable hypothesis, the two models have effectively the same inferencing properties for the limited data set we are working with, so I will proceed by looking exclusively at the simpler model on the left.

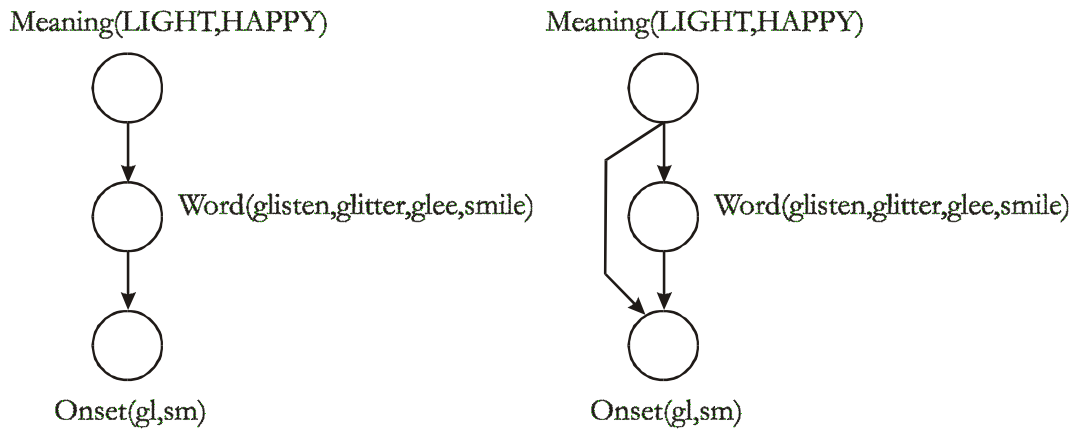


Figure 9 - BNs for a simplified phonaestheme domain

Asking Tetrad to estimate the conditional probabilities for each node results in the CPTs shown in Figure 10 below. We can see from these numbers that the meanings LIGHT and HAPPY are equally likely, as there were two words with each meaning in the data set. By the same token, given each of these meanings, each word is equally likely. For example, *glisten* and *glitter* are equally probably given that the meaning is LIGHT, while *glee* and *smile* are not at all likely. Finally, considering the chart in figure 10c, the probability of each onset given the word it occurs in is assigned a probability of 1.

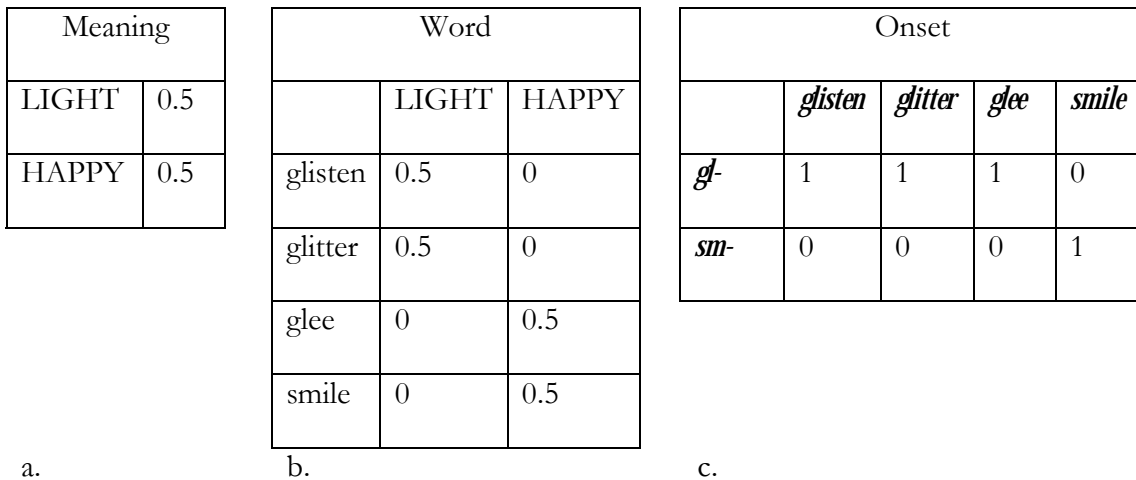


Figure 10 - Conditional probability tables for (a) Meaning node, (b) Word node and (c) Onset node

Now we can probe the network to see how these probabilities change when certain facts are known. Let’s start by asking it to guess a word’s semantics on the basis of its onset. When Onset is clamped at *gl-*, the words beginning in *gl-* should be equally likely, while *smile* shouldn’t be at all likely. This is precisely the result of inference shown in Figure 11a. Moving further up the network, with the same clamping of Onset at *gl-*, we find that the probabilities for the Meaning node become those in Figure 11b. Here, we see that LIGHT is twice as likely a meaning as HAPPY, due directly to the distribution of words starting with *gl-* that mean LIGHT relative to those that mean HAPPY. This evocation of a phonaestheme-related meaning when presented with only the phonological cue associated with that phonaestheme matches subjects reactions to neologisms, described in Chapter 2.

Word	
<i>glisten</i>	0.33
<i>glitter</i>	0.33
<i>glee</i>	0.33
<i>smile</i>	0.33

Meaning	
LIGHT	0.67
HAPPY	0.33

a. b.

Figure 11 - Probabilities of (a) Word and (b) Meaning values given Onset = gl-

We can also assess the network’s prediction of a word’s form on the basis of its semantics. If we clamp the Meaning node at LIGHT, then the Word node has the values in Figure 12a. Here, *glisten* and *glitter* are equally likely. These words share the onset *gl-*, so the value of Onset trivially selected is *gl-*, as seen in Figure 12b. This behavior mimics subjects’ responses to novel phonaestheme-related definitions.

Word	
<i>glisten</i>	0.5
<i>glitter</i>	0.5
<i>glee</i>	0
<i>smile</i>	0

Onset	
<i>gl-</i>	1
<i>sm-</i>	0

a. b.

Figure 12 - Probabilities of (a) Word and (b) Onset values given Meaning = LIGHT

We’ve now seen that the network will tend to predict a phonaestheme-related semantics on the basis of the onset identifying that phonaestheme, and will predict a

phonaestheme-related onset on the basis of appropriate semantics. It does so solely on the basis of the distribution in the lexicon of words sharing that form and meaning. We can now turn to what the network predicts about tasks like the priming experiment in Chapter 2. In that task, subjects were presented with two word in quick succession, where those words could share a phonaestheme. If they did share a phonaestheme, the second word was identified more quickly than target words following a pseudo-phonaesthemically related prime.

As it stands, a phonaestheme model like the one outlined up to now is purely static. This is problematic for modeling the priming behavior described in Chapter 2. While priming involves the continued activation of some neural structure over time, the BN models described so far do one-shot inference in a single, static time slice. Whenever a new node is clamped, inference is re-initialized.

A general solution to the problem of modeling dynamic systems in BNs involves incorporating time information into the structure of the BN. Dynamic Belief Networks (DBNs) view time as being broken into time slices, where each slice includes a representation of the variables that persist through time. So the dynamic equivalent of the phonaestheme model in Figure 9 would look something like the one presented in Figure 13. A DBN involves as many state descriptions as it allows time slices. Temporally-contingent influences are represented as connections between the state descriptions in different time slices. For example, in Figure 13, the meaning value at time 1 (T_1) influences the meaning value at time 2 (T_2).

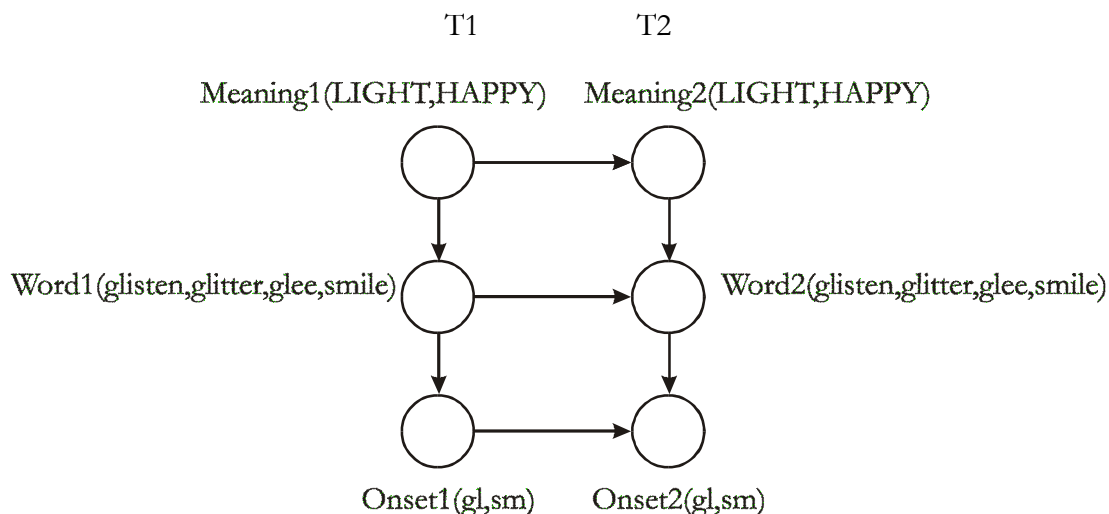


Figure 13 - A simple DBN for phonaesthemes.

The model in Figure 13 is an example of one of the simplest sorts of DBN - one in which the value of each node at each time is only influenced (aside from co-temporal variables) by its own value at the immediately preceding time. This sort of model can account for the priming effects phonaestheme-users demonstrated as described below.

Let's assume that each node bears a relationship to its own incarnation in the next time slice such that it increases the probability that the same value will be true in the future as is true in the present.

In a situation parallel to the one created by the experiment reported on in Chapter 2, subjects observe Onset to have a particular value in T1, in the phonaestheme case, *gl*. This makes words sharing that onset more likely in T1. That is, Word1's values *glisten*, *glitter*, and *glad* become more likely. Observing *gl* in Onset1 also makes *gl* a more likely value of Onset2, the onset in T2. Additionally, the words in T1 that have become more likely due to having the onset *gl* in T1 make their equivalents more likely in Word2. Given that three of the four words in Word1 have become more likely, their meanings will also become more

likely. More of the active words (the ones sharing *gl*) bear the meaning LIGHT than HAPPY, so LIGHT will become more likely in Meaning1 than will HAPPY. This will make the meaning HAPPY more likely in Meaning2, and as inference continues to propagate through the network, the increased likelihood of LIGHT in Meaning2 and *gl* in Onset2 will make each value of Word2 that shares these Meaning and Onset values more likely than its counterparts that may share one or none of these features.

We can test this process of spreading inference quantitatively using the network in Figure 13. Notice that the CPTs for T2 will now be slightly more complex, since each node now has two parents, rather than one. Persistence of activation can be represented if we assume that each value of a node in T2 is 0.25 more likely if the same value is observed for the same node in T1. For example, then, the CPT of Meaning2 will look as in Figure 14. The other nodes can follow the same pattern.

Meaning 1		
	LIGHT	HAPPY
LIGHT	0.75	0.25
HAPPY	0.25	0.75

Figure 14 - CPT for Meaning2 in a DBN for phonaesthemes

In such a model, given that Onset1 is observed to have value *gl*, the probabilities of the words in Word2 will be skewed to reflect phonaesthetic distribution, as seen in Figure 15a. Although *glee* shares the onset observed in Onset1 with *glisten* and *glitter*, it remains less likely. Similarly, when both Onset1 and Onset2 are clamped at *gl*, the probabilities of *glisten* and *glitter* are slightly higher than that of *glee*, as shown in Figure 15b. In both of these

simulations, it is the distribution in the lexicon of words sharing form and meaning that leads to increased likelihood of words sharing this form and meaning when a phonaesthetic prime is presented.

Word		Word	
“glisten”	0.3	“glisten”	0.33
“glitter”	0.3	“glitter”	0.33
“glee”	0.27	“glee”	0.3
“smile”	0.13	“smile”	0.03

Figure 15 - Probabilities of Word2 when (a) Onset1 only and (b) Onset1 and Onset2 are clamped at *gl-*

Both the static and dynamic BN models of phonaesthemes presented above make the prediction that the simple distribution in the lexicon of shared form and meaning will give rise to processing asymmetries. Those asymmetries are the same ones observed in the priming experiment and neologism experiments described in Chapter 2.

From a broader perspective, these models demonstrate how the full productivity of a linguistic unit can be unnecessary for that unit to have a psychological status. A model built up simply from statistical correlations between observable variables can model human behavior, whether or not those correlations are fully productive.

5. Properties of the models

Relations to other cognitive models

Models based on BNs, like connectionist models (e.g. Feldman 1988, Rumelhart et al. 1986) are bidirectionally grounded. They are expected to mimic cognitive/linguistic behaviors while simultaneously being responsible to the neural explanations for those behaviors, and incorporating them into a computational model.

BN models bear striking and not unintended similarities to other usage-based models of language (e.g. Bybee 1999, Langacker 1991, Kemmer and Israel 1994, Tomasello In Press). They are constructed on the basis of linguistic experiences by a person in the world. They are based on abstractions over grounded experiential knowledge, knowledge which is not discarded, but rather represented simultaneously with the abstract versions. In usage-based models like the current one, representations are schematic, probabilistic, and redundant.

Usage-basedness is a necessary consequence of taking a broad perspective on factors that can influence linguistic structure. Obviously, most statistical correlations between variables in a linguistic utterance cannot be inborn in the human species, and certainly their quantitative details cannot be either. Rather, they can only arise from language experience and from abstractions over that experience. Nevertheless, even if abstractions are drawn, they must remain closely tied to the perceptuo-motor details they are derived from, or else they could not be used. A usage-based perspective is inclusive in that it can also capture otherwise apparently categorical, and thus potentially innate or top-down generalization, like the phonologically-based allomorphies described in section 4.

The models presented above are also similar to other embodied models of language and cognition (e.g. Lakoff 1987, Johnson 1987). Their structure is strongly constrained along three dimensions. The models described above are grounded in the neural structures responsible for language; they are neurally embodied. They also ground out in (are abstractions over) perceptual and articulatory linguistic knowledge; they are corporeally embodied. Finally, they are grounded in the actual communicative use language is put to, by being generalizations over actually observed utterances, including their social correlates; they are embodied through use.

Unlike most usage-based and embodied models, however, BNs provide a quantitatively detailed computational architecture. This architecture can represent large-scale problems, and importantly, learning in such models.

Power of the model and potential explanation

A useful metric for linguistic models is the power that they bring to bear on a particular task. In general, models experience a tradeoff between their representational power and their explanatory power. A model capable of capturing more complexity is usually seen as less able to explain data it captures than is a less powerful model. This aspect of models is relevant to the present work since BNs are much more powerful than most other phonological modeling architecture that has been proposed, and certainly more powerful than all mainstream models.

There are two reasons why the power argument should not affect the decision to use BNs as tool for building linguistic models. First, qualitatively less powerful models than BNs are unable to capture the behaviors we've looked at above. The next most powerful

computational architecture to BNs is known as QPNs - Qualitative Probabilistic Networks (Wellman 1990). These encode relations between variables not in terms of specific conditional probabilities, but rather as simple qualitative effects of parents on children. Rather than a CPT for each node, a QPN has a table indicating whether the effects of a particular parent variables values are positive or negative - whether they increase or decrease the likelihood of the child nodes' values. Figure 55 compares a simple BN CPT with the qualitative influence table for a node of a QPN.

V1.1				
A	1	2	3	4
One	1	0	0.7	0.1
Two	0	1	0.3	0.9

V1.1				
A	1	2	3	4
One	+	-	+	-
Two	-	+	-	+

Figure 55 - The CPT for a BN and one for a QPN

It should be clear that QPNs fail to capture the quantitative details of causal relations. There is no way to distinguish with just two nodes in a QPN like the one in Figure 55 between the effects of node A having value 1 and value 3. And yet the numerical details of these relations are essential for computing the relative effects, for example, of a liaison consonant being an /r/ or an /z/. Both disfavor liaison, as shown in Chapter 3, but to radically different degrees. This generalization would be lost in a QPN model. We see then that machinery at least as powerful as BNs is needed to capture the data presented in this thesis.

A second reason why the computational power argument does not immediately discount BNs is that BNs only constitute a computational architecture for building linguistic

models. It remains to be seen exactly in what way this architecture would be implemented in a general theory of phonological and morphological knowledge. Presumably, a full BN model would have to be further constrained by general human cognitive properties, like attention, the time course of linguistic processing, and short and long term memory (the “Head-filling-up problem” - Johnson 1997).

A BN linguistic model would also have to be further constrained by human inductive and deductive power. There is some indication that people actually reason in linguistic and non-linguistic tasks in ways quite similar to BN inference (Tenenbaum 1999, 2000, Tenenbaum and Griffiths 2001). But it is unreasonable to assume that human reasoning would match an arbitrarily complex BN.

Even assuming, though that a BN model of some aspect of language that closely conformed to human cognitive power could be constructed, we would still might run into the argument of representational power. After all, even a humanly constrained BN would still be able to learn most any superficial correlation in the linguistic or extralinguistic evidence it was confronted with. Can a BN ever predict what correlations will be extracted and which ones not? Can it make predictions about linguistic typology?

In answer to the first question, I think that a BN model rightly predicts that most every correlation that is significant beyond an as of yet undetermined threshold will have a measurable effect on linguistic processing. After all, look at the bizarre and overwhelming array of subtle statistical correlations described in the preceding chapters, the knowledge of which, given appropriately acute measures, are apparent in language processing. From trends in the phonology of first names to generalizations about age effects on liaison, human language users pick up on statistical generalizations in their environment. I haven’t yet seen any limitations on language users’ statistical correlation extraction capacities, only limitations

on which ones have been detected and studied!

Now, there must be some quantitative limit on what correlations individuals are able to extract from their environment. This could be in terms of the minimum strength a trend must have to be detected or the number of decimal points to which a hearer can predict the distribution of a variable's values. But these restrictions remain to be assessed empirically. Only through the development of models that can outperform humans can we ascertain the limits of the human cognitive potential.

I also believe the answer to the second question above - whether a BN model can make predictions about linguistic typology - is 'yes', despite the overwhelming power of a BN. The reason for this belief is that I don't think a synchronic model of linguistic knowledge - an individual representational model - will ever succeed as the sole explainer of linguistic typology. Rather, it must be combined with acquisitional, historical, and contextual models to provide an explanation; it can model proximate causation and must be a prominent but not sufficient part of a model of ultimate causation.

Importantly, a usage-based BN model allows us to understand an important aspect of language change. When statistical tendencies are extracted from ambient language, these trends will tend to affect language processing. We've seen above that language perception reflects statistical distributions through speed of processing and response tendencies in forced-choice tasks.

Remember that a BN model is not only a model of perception, in its diagnostic mode. BNs also can serve to model language production, through causal reasoning. Mere tendencies observed in ambient language will inevitably come to taint language production. For example, given that a speaker wishes to express a meaning related to vision, the BN model of phonaestemes above predicts that that person will be more likely to select a word

with an initial gl- than with some other less well represented onset. In fact, the miniature model shown in section 3 above assigns a probability of 1 to the production of a gl- when the semantics is 'LIGHT'. Neologisms and speech errors should therefore both follow the patterns that already exist in the language.

In other words, the BN model would predict that the statistically rich should get statistically richer. This sort of diffusion over the lexicon has been shown to have historically taken place for all sorts of statistical trends, like phonaesthemes (Marchand 1959), Austronesian roots (Blust 1988), and strong verbal morphology (Bybee and Moder 1983).

Moreover, BNs present a convenient framework for representing what could be the basis for intentional manipulation of linguistic variables for social purposes. It has been widely documented that the use of linguistic variables can depend upon non-linguistic social attitudes of speakers. Subjects in Labov's (1963) study of Martha's Vinyard, for example, were more likely to produce the local (non-standard) centralized variants of diphthongs /ay/ and /aw/, the more resistant they were to the development of tourism in their historically isolated community.

The intentional manipulation of linguistic variables for social effects is tantamount in a BN model to assigning particular values to nodes representing social factors, and allowing inference to skew the subsequent linguistic effects appropriately. Presumably, the only reason a speaker would reasonably assume this could be an effective means for achieving a social goal is the knowledge that other hearers make inferences about social causes from the character of the linguistic signal. In other words, someone with a causal model can artificially manipulate the hidden variables such that the effects are interpreted by hearers in a way that that speaker intends.

I have very subtly transitioned here into a discussion of individual language

production, about which I have very little else to say. Although there are indications, as cited in Chapter 3 above, that individuals production follows the patterns of the social groups those individuals belong to, from the evidence I have presented in this thesis, there is little evidence for or against this belief. If it were the case that individual language production reflected the production of a social group, then a BN would also be an extremely effective tool for capturing this variation, in the same way as it captures the production of the community. The degree of fit here, though, remains to be evaluated.

Chapter 6. Neural bases for unruly phonology

Outline

1. Introduction
2. Levels of representation
3. Neural explanation
4. A neural theory of Belief Networks

The throughput principle: That which goes in at the ear, and out from the mouth, must somehow go through the head.

Mattingly and Studdert-Kennedy

1. Introduction

In the previous chapter, I presented a computational mechanism that is able to model probabilistic effects on phonology and interactions between these factors, as well as the temporal effects observed in priming. We now ask why it is that the five properties demonstrated in this thesis (and summarized in (1) below) should be part of the cognitive processing of language. The answer will be found in the neural bases for language processing, learning, and representation. Then we can ask whether the computational mechanisms proposed in Chapter 5 can help bridge the gap between cognitive and linguistic behaviors and their neural explanations.

(1) Properties to be explained

- Language users acquire probabilistic linguistic knowledge.
- Language users encode probabilistic linguistic knowledge.
- Some of this probabilistic knowledge involves correlations among different domains of knowledge, such as between phonology and social cognition.
- Full productivity is not essential for form-meaning pairings to play a role in language processing, as shown by phonaesthemes.
- Language users make use of probabilistic interactions between factors, as in the case of French liaison.

While the properties in (1) are quite difficult to explain from the perspective of a deterministic, modular, rule-based modeling enterprise, they are anything but surprising when language is considered in its biological context. There is clear evidence from studies of learning and processing in the brain that probability, cross-modularity, and schematicity are anything but aberrant. Quite the opposite, in fact - given brain facts to be discussed in the next section, it would be surprising if we didn't use knowledge of probabilistic, interacting correlations across domains of linguistic and extralinguistic knowledge.

2. Levels of representation

Although neural considerations are crucial to an understanding of linguistic knowledge and behavior, most modelers of human language and cognition find working at the level of neural structure difficult. For this reason, it seems useful to work at an intermediate Computational level, as shown schematically in Figure 1 below. As we will see below, not

just any neurally plausible computational model is a valid abstraction over the neural level. While any bridging model necessarily abstracts somewhat from the details of neural structure, an explanatory one must in particular display bidirectional *isomorphy*. That is, computational mechanisms posited to account for a given cognitive behavior must themselves be directly grounded in the neural mechanisms that are themselves directly responsible for the cognitive behavior.

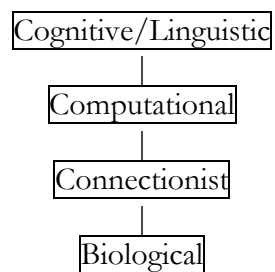


Figure 1 - The placement of a computational level of analysis

Because of the complexity of neural systems, an additional level will intervene in the following discussion between the computational and biological levels, as seen in Figure 1. This connectionist level is a low-level abstraction over neural processing that picks out the computational properties of neural systems. The connectionist level helps to bridge the gap between the purely computational and the purely biological.

The idea of representing a mental phenomenon at different levels is by no means new. Probably the best known breakdown of functional systems is Marr's (1982) three-way distinction between the computational (or functional) level, the algorithmic/representational level, and the implementational level. On Marr's view, any machine performing some computational task must be seen as having these three distinct levels of representation, which respectively describe the purpose or highest-level function of the system, the

algorithms used to implement these functions, and finally the actual hardware, software, or wetware it is implemented on.

The notion of level I am proposing here, following Bailey et al. (1997), differs in its purpose. Unlike Marr's purely descriptive schema, the current proposal is intended to bear some explanatory power. In the present conception, levels of representation mediate between two qualitatively distinct, empirically observable types of evidence about some human comportment: cognitive/linguistic behavior and neural functioning. The observations we make at these levels will be described as pertaining to the cognitive/linguistic and biological levels respectively.

Directly mapping these observations to each other is difficult for two reasons. First, the data we have at the two levels is of fundamentally different types. Neural synapse functioning is not of the same stuff as perception data. This means that some mapping has to be constructed. By necessity, this mapping transforms observations and structure from the two levels into some common language. It seems that for problems of the scale of language, the most efficient common ground is information processing. Any model of the functioning of a neural system requires metaphorical understanding of electrochemical patterns as information propagation.

A second reason why an intermediate level of representation is necessary to the enterprise of aligning mind and body is that the scale and complexity of the neural structures in question is impractical for an analyst to manipulate. People excel at manipulating physical and mental objects, in performing elementary arithmetic and spatial reasoning, and in doing basic causal reasoning. Analysts, therefore need models that satisfy these restrictions. Conceptualizing even a small piece of the brain's 100 billion neurons, firing up to 1000 times a second along 100 trillion synapses far exceeds our conceptual or perceptual potential as

human scientists.

In order for a computational model to act as a bridge between the neural and cognitive/linguistic levels, it must to the greatest extent possible capture those properties of the neural system that are considered explanatory of cognitive/linguistic behavior. For our purposes, it must have a way of representing the aspects of language processing enumerated in (1) above. Additionally, the mechanisms in an explanatory bridging model that give rise to these cognitive and linguistic behaviors must themselves be mapped to the neural mechanisms that give rise to the behaviors. In other words, it's not enough to show that a model can capture interactions between factors. The computational level mechanism responsible for interactions must also be mappable to the neural mechanism responsible for interactions. This specificity of the mappings from above and below to the computational level can be referred to as their relational *isomorphism*.

A bridging computational theory is more tightly empirically constrained than a functional model of a single behavioral domain. When the modelling enterprise is just constrained by cognitive/linguistic phenomena, any one of a large number of theories of varying neural plausibility would be equivalently possible. For example, phonological theories based on language-internal and typological distributions of phonological units do not necessarily have any grounding in the biological substrate responsible for individual language learning and use. A large number of qualitatively and quantitatively different functional models are thus behaviorally equivalent.

3. Neural explanations for non-categorical phonology

There are very clear neural explanations for the sorts of linguistic behavior identified in the preceding chapters. I will demonstrate the utility of considering biological explanations for these behaviors through neural answers to the following questions, which are adapted from the properties listed in (1):

(2) Questions addressed in this section

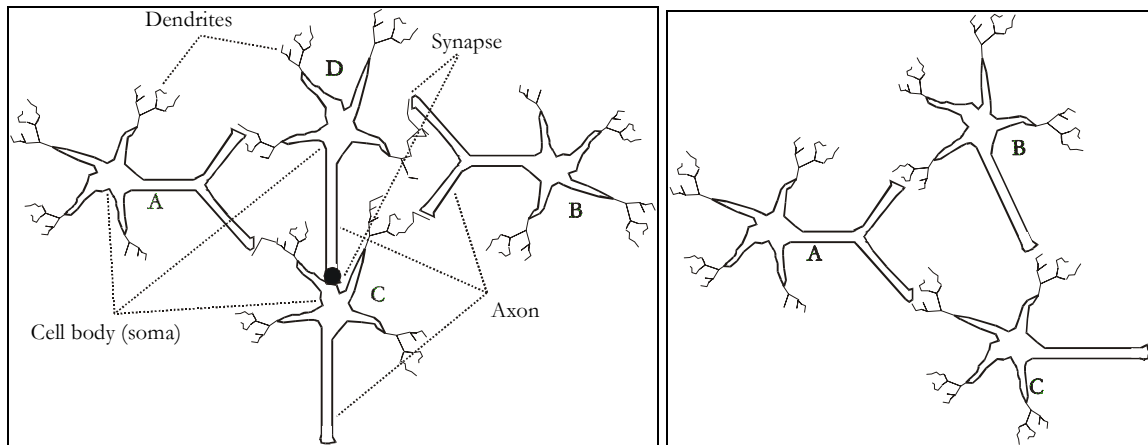
- a. How is probabilistic linguistic knowledge encoded?
- b. How do language users acquire probabilistic linguistic knowledge?
- c. How can some of this probabilistic knowledge cross over domains?
- d. Why isn't full productivity essential for form-meaning pairings to play a part in language processing?
- e. How do perceivers make use of probabilistic interactions between factors?

In addressing each of these concerns in turn in the rest of this section, I will also be constructing a list of mappings between cognitive and linguistic behaviors on the one hand and neural mechanisms on the other.

a. How is probabilistic linguistic knowledge encoded?

Neural processing is inherently probabilistic, both at the scale of the neurons, the brain's individual information manipulating units, and at the scale of neural ensembles. Neurons are connected to one another through *synapses*, gates across which where they pass information

to other neurons (Figure 2). Synapses are located at the end of the (potentially branching) axon of the *pre-synaptic cell* (the cell sending the information) and usually on the dendrites of the *post-synaptic cell* (the cell receiving the information). That information takes the form of electrical or chemical signals. To a large extent, the only information a neuron can pass to its downstream neighbors across synapses is its degree of activation.



a.

b.

Figure 2 – (a) Four neurons, with synapses from A to C, A to D, B to C, and B to D. Each neuron has a body, dendrites, and a (potentially branching) axon. D is an inhibitory interneuron, accomplishing a diminutive interaction. (b) Three neurons; A has excitatory synapses on both B and C, creating an augmentative interaction between A and B on C.

Neurons, unlike the machinery of digital computers, are not best viewed as binary on-off switches. Their output is in the form of a sequence of *spikes*, periods of greatly changed electrochemical polarity. While a neuron's spikes are quite uniform, they time between spikes is not. Spikes are usually separated by at least 1 msec, but the time between spikes is generally not very regular. As a consequence, the input to a downstream neuron

reading information about the activation status of its predecessor must integrate the spikes it takes in over time. In other words, information passed between neurons is continuous over time, although it is passed in discrete packets.

A given neuron's reaction to a particular environmental stimulus is not deterministic either. For example, Figure 3 shows the response of a neuron in a monkey's somato-sensory cortex (responsible for haptic perception) to an edge placed in its hand at different angles. The vertical lines on the right represent spikes emitted by the cell over time. This particular cell fires most when the edge is oriented horizontally, but less with other orientations.

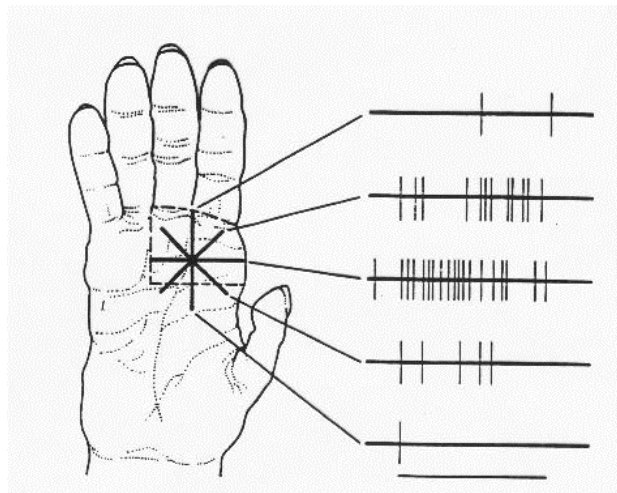


Figure 3 - The receptive field of a neuron in the monkey's somato-sensory cortex that responds when an edge is placed in the hand. This cell responds well when the edge is oriented horizontally, but less well to other orientations (From Hyvarinen and Poranen 1978).

The same gradient response is characteristic of groups of neurons as they respond to a stimulus. The information passed between neurons and groups of neurons is graded.

Essentially every aspect of a neural system, from activation patterns to information passing, is non-categorical. This explains why it is that some linguistic knowledge should be non-categorical.

b. How do language users acquire probabilistic linguistic knowledge?

Synapses, the connections between neurons, determine the organization of the brain at the super-neuronal level. They are also the locus of the large part of learning. Most learning in the adult brain involves the adjusting of synapses between neurons. One central mechanism responsible for the long-term recalibration of synapses is known as Hebbian learning, named for the neuroscientist Donald Hebb. The idea of Hebbian learning is extremely simple:

When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased. (Hebb 1949, p.62; his italics)

In other words, if there exists a potentially useful but currently weak synapse between two cells, for example, A and C in Figure 2a, then if A and C tend to be co-active, the synapse between them will be strengthened such that electrochemical signals are more readily passed from A to C. This scenario can play out, for example, when A and B are commonly co-active and there is a strong connection from B to C, which causes C to fire when B does. Imagine that A represents the perception of a bell ringing, B the perception of food, and C the mechanism responsible for activating salivation. In this case, B will be (perhaps innately) linked to C, such that when food is perceived, the animal salivates. When

a bell is heard along with the presentation of food (that is, when A and B are co-active), then A and C fire together. Hebbian learning ensures that the A->C synapse is strengthened, such that A can now activate C; perception of a bell leads to salivation, even when no food is perceived.

While the neurobiology of Hebb's day couldn't provide a precise chemical explanation for this sort of learning, Long Term Potentiation (LTP) has recently been demonstrated to serve precisely the purpose of strengthening synaptic weights when a pre-synaptic and a post-synaptic cell are co-active (Kandel 1991). Because it involves the incremental strengthening of connections between cells that are co-active (giving rise to the neurobiologist's mantra "cells that fire together wire together"), LTP is believed to be responsible for associative learning. On this widely held view, the *recruitment* of latent potentially viable connections for associative purposes provides us with the ability to notice and remember correlations between perceptions.

Given even this brief introduction to neural structure, we can already see the extent to which probabilistic processing of linguistic knowledge is inevitable. Language users are constantly bombarded with linguistic input. This input varies along many dimensions and co-variance among variables characterizing this input is rampant, as we saw in the corpus study in Chapter 3. When we concurrently perceive two environmental factors, like the grammatical class of a word and the production of a liaison consonant, or like the character of an onset and a particular semantics, then Hebbian learning will ensure that over time a connection between the neural structures responsible for the detection of those factors will be strengthened. Hebbian learning explains how probabilistic correlations are learned.

c. How can some of this probabilistic knowledge cross over domains?

To answer the question of cross-domain knowledge, we must move to a higher level of brain structure. Since the beginning of the nineteenth century, researchers have been aware that many brain functions are mostly localized in specific processing regions (e.g. Gall and Spurzheim 1810). Carl Wernicke's (1908) and Pierre Broca's (1865) early work on patients with brain lesions in specific brain areas identified two regions of the brain that are responsible for certain aspects of linguistic processing. Broca's area is partially responsible for the processing of grammatical information and the production of sentences. Wernicke's area is classically seen to be responsible for speech comprehension and the ability to choose and use words in a meaningful way. Since the brain computes locally, domain-internal associations should dominate cross-domain associations.

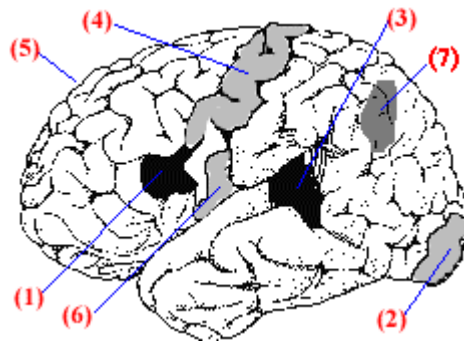


Figure 4 - The left hemisphere of a human brain, showing (1) Broca's area, (2) Visual cortex, (3) Wernicke's area, (4) Motor cortex, (5) Frontal cortex, (6) Auditory cortex, and (7) Angular Gyrus (from www.ohiou.edu/~linguist/L550ex/neurolingvid.htm)

But recent work has shown that language functions are not as localized as previously thought (Mayeux and Kandel 1991). First, classic studies that identified Broca's and

Wernicke's areas as language-specific processing areas in fact dealt with subjects whose lesions damaged more than just the restricted area that was identified. In fact, it seems that sub-cortical areas, such as the left thalamus and caudate nucleus are also critically responsible for language processing.

Second, although brain lesions in certain areas tend to correlate with loss of particular language functions, areas dedicated to language processing also perform other cognitive functions. For example, Maess et al. (2001) have recently demonstrated using magnetoencephalography that Broca's area reacts not only to syntactic abnormalities, but also to musical abnormalities, specifically harmonically inappropriate chords.

Third and finally, the simple observations that brain region is associated with a given cognitive behavior does not imply that it does so independent of other types of information. In fact, both Broca's and Wernicke's areas receive long distance projections from other brain regions, specialized for other functions. Broca's area collects information, among other things, from Wernicke's area. Wernicke's area itself has afferent (incoming) connections from at least auditory cortex, visual cortex, and somatosensory cortex (responsible for tactile, pain, cold/hot, and position sensation). Both are also likely to employ information from other processing centers. The data suggest that even if Broca's and Wernicke's areas don't receive afferent connections from all the information processing areas responsible for the factors demonstrated in the preceding chapters, this information is consolidated somewhere.

We can see that in the brain areas most responsible for language processing, information from different domains is available through incoming long-distance connections. Moreover, the non-encapsulation of linguistic functions in Wernicke's and Broca's area increases the likelihood of functional overlap between linguistic and

extralinguistic information processing. If neurons encoding different modalities fire together, then Hebbian learning suggests that they will tend to wire together.

d. Why isn't full productivity essential for form-meaning pairings to play a part in language processing?

According to the phonaestheme experiment reported on above in Chapter 3, productivity is not an essential prerequisite for a sub-lexical unit like a phonaestheme to achieve some psychological status, as measured by a priming task. As we will see below, this is due to Hebbian learning over repeated experiences that excite common neural structure. Before leaping into priming, though, we need to discuss how sounds and meaning are represented neurally.

To a first approximation, there are two major schools of thought as to how concepts are represented in the brain. The localist view (e.g. O'Seaghdha and Marin 1997) sees individual neural structures as responsible for particular concepts. For example, in order to represent all four members of the Beatles, we can assume four groups of neurons, each of which is dedicated to one of the concepts. In the simplified model in Figure 5, a given cell or group of cells (cells A, B, C, and D) is active (has a value of 1) in association with the concept it represents. This localist view has been criticized because in the limit it is untenable. Certainly we don't have just one neuron dedicated to representing the idea of 'Ringo Starr' - otherwise losing that neuron would completely remove the notion of the misunderstood percussionist from our universe. Additionally, there are many aspects of Ringo Starr that contribute to our knowledge of him. His appearance, the sound of his voice, the sound of a particularly clunky riff in "All You Need is Love"; these are all parts of what

make up the concept of Ringo Starr. It would be difficult to see them as being all localized in a given neuron or even a group of local neurons.

		Neurons			
		A	B	C	D
Beatles	John	1	0	0	0
	Paul	0	1	0	0
	George	0	0	1	0
	Ringo	0	0	0	1

Figure 5 - A localist encoding of the four Beatles

An alternative perspective sees representations as distributed over shared neural circuits (e.g. Plaut and Gonnerman 2000). For example, the four Beatles could be represented with different activation patterns over just three neurons as in Figure 6. Here, neuron A is active to represent John, while both A and C are active for George.

	A	B	C
John	1	0	0
Paul	1	1	0
George	1	0	1
Ringo	0	0	1

Figure 6 - A distributed encoding of the four Beatles in three neurons

Distributed representations have several advantages over localist ones. They make more efficient use of neuronal resources than localist representations, as the Beatles example demonstrates. They also degrade more gracefully. At the same time, though, when more than one concept is active, a distributed representation will experience cross-talk, which will confound the representation. For example, in the examples above, it's impossible to distinguish between situations the situation where both John and Ringo are active and the one where only George is active; both yield activation pattern [1 0 1].

Whether words, their meanings and their phonological elements are represented locally, distributedly, or by some alternative method such as coarse-coding (Tienson 1987), words that share semantics will share activation over some neural structure representing that semantics, and words sharing a phonological component will share activation over the neural structure responsible for that component. For example, in Figure 7, the words *glisten* and *glimmer* share an onset *gl-* as well as some semantics, in this case, LIGHT. This is rendered in the localist model on the right as connections between a neuron representing each word and a common phonological and a common semantic neuron. In the distributed model on the left, the pattern of activation over the phonological layer for each word will share some features, as will the pattern of activation over the semantic layer.

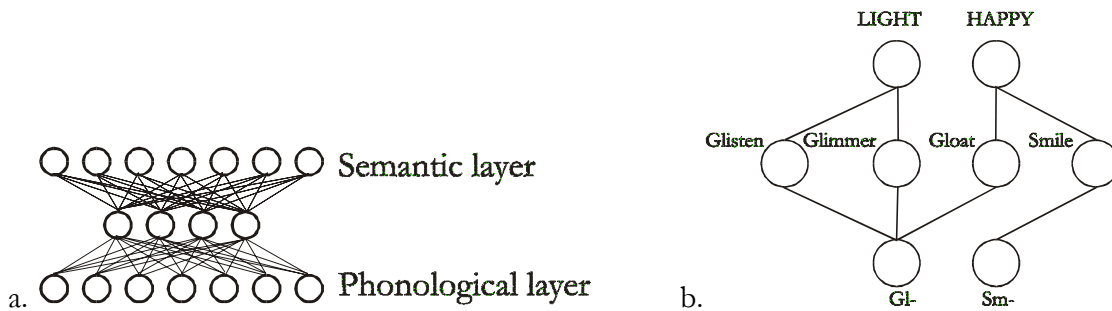


Figure 7 – (a) A distributed and (b) a localist representation of phonology and semantics

Given, then, that *glisten* and *glimmer* will both activate shared semantic and phonological neural structures, it is straightforward that new links between these units will be recruited through Hebbian learning.

The phenomenon of priming is evidence for the existence shared structure. In both localist (e.g. Dell and Juliano 1996) and distributed models (e.g. Rumelhart et al. 1986), priming of the type found in the experiments in the previous chapters is the product of ambient activation remaining on shared units.

Supposing then that the neural basis for the priming effects found for phonaesthemes is in neural connections, and knowing that such connections arise through an exceedingly common neural learning process, our outlook on phonaesthemes changes. If human brains are automatically restructured on the basis of co-occurring perceptual experiences, and if this restructuring takes the form of strengthening of connections, then learning phonaesthemes no longer appears to be an unlikely and difficult process. Rather, phonaesthemes, as statistically over-represented form-meaning pairings, must inevitably gain a foothold in the human cognitive system, whether or not they combine with other morphemes to construct the words they occur in. In fact, we should expect that any tendency in the lexicon above some minimum threshold should give rise to Hebbian learning. This process would also give rise to the knowledge of phonology-syntax correlations cited in Chapter 2 above.

Now we can see why productivity isn't an essential prerequisite for associations like phonaesthemes to play a part in language processing. Words that share form and meaning will share neurons or firing patterns over neurons. With frequent exposure to those words, that is, simultaneous activation of neural structures responsible for the shared the form and meaning, links between those neural structures will be strengthened through Hebbian

learning. Subsequently, when the neural structures responsible for either the form or meaning are activated, activation will also be passed to the other pole of the association. Thus, hearing *gl* should increase the likelihood that the semantics ‘LIGHT’ or ‘VISION’ will be evoked, and vice versa. The productivity of the sub-lexical association is not essential for an association to be learned in a neural system, so it should not be surprising that non- or partially-productive form-meaning pairings should be part of human linguistic knowledge.

e. How do perceivers make use of probabilistic interactions between factors?

In Chapters 3 and 4, we saw two types of interaction above that we need to account for. In the first, the combination of two particular values of independent variables gives rise to a greater likelihood of a value of the dependent variable than would be otherwise expected. We’ll call this an *augmentative* interaction. For example, while increasing age in general increases the likelihood that a liaison consonant will be produced, when a speaker is young and the liaison word is an adverb, the probability of liaison is greater than otherwise expected. In the second, opposite, type of interaction, the combination of two values yields a lower likelihood of a given value of the dependent variable: an *diminutive* interaction. These two sorts of interaction both have their roots in the integrative function of brain cells.

Neurons collect inputs from other neurons that have synapses on them. Those inputs can enter the neuron through synapses that are located at different points on the cell, and they can be received by the neuron at different times. A post-synaptic cell must therefore *spatially* and *temporally integrate* information. It must additionally integrate *excitatory* connections, like the ones discussed up to the present, with *inhibitory* ones.

Inhibitory synapses are links between neurons that have the following effect. When a pre-synaptic cell like cell D in Figure 2a fires, post-synaptic cell C is less likely to fire. Figure 2a shows an inhibitory synapse between these two cells; inhibition is usually marked with a black head on the pre-synaptic axon. Inhibitory synapses are most often located on the body, or soma, of the post-synaptic cell.

In general, a given neuron will have only excitatory or inhibitory outgoing synapses, because the biochemistry of the two types of cell is radically different. Therefore, for a given neuron to have excitatory effects on some neurons and inhibitory effects on others, so-called inhibitory interneurons are recruited. These cells serve the purpose of essentially reversing the polarity of a cell's effect. As shown in Figure 2a, excitatory neuron A can inhibit C if it has an excitatory connection on inhibitory interneuron D. When A becomes active, its inhibitory synapse on D has the effect of inhibiting C.

We can now see how that diminutive interactions could be effected by a single interneuron. Two neurons, like A and B in Figure 2a, have a diminutive interaction on cell C if they both have excitatory synapses not only on C, but also on inhibitory interneuron D. This inhibitory neuron D will decrease the likelihood that C will fire whenever both A and B are active. A necessary feature of interneuron D is that it require input from both A and B to fire. If this is the case, then the diminutive interaction, that is, the inhibition of C, will not be felt when only A or B is active.

In a similar vein, an augmentative interaction can be realized through an excitatory synapse between the two presynaptic cells. In Figure 2b, we see that if cell A has an excitatory synapse on cell C and also on cell B, and if this synapse is insufficient to bring B to fire by itself, then the effect of this synapse will be felt on C only if B is also active. By the

same token, if only A is active, there will be no influence of B on C, thus making the A → B synapse useless.

The overall point here is that the two types of interaction between factors can be accomplished at the neural level through the recruitment of either inhibitory interneurons or simple excitatory synapses between the two interacting factors. They will be recruited with repeated correlations in linguistic experience.

Summary

We have seen neural-level mechanisms that can account for each of the types of linguistic phenomenon that were documented in Chapters 2 through 4, as summarized in Figure 8. Each of these biological mechanisms is well-documented, and widely distributed throughout the human central nervous system.

<i>Neural Phenomenon</i>	<i>Cognitive/Linguistic Phenomenon</i>
Hebbian learning	associative learning of probabilistic correlations
graded computation	probabilistic knowledge
converging afferent connections	cross-modal information combination
overlapping neural structure	priming of partially productivity units
inhibition and excitation	diminutive and augmentative interactions between factors

Figure 8 - Summary of neural bases for cognitive/linguistic phenomena

I have not demonstrated instrumentally that these particular mechanisms are in fact responsible for the linguistic behaviors in question. All of these same mechanisms have been observed to be responsible for similar behaviors in different domains, however (Kandel et al. 2000):

- Hebbian learning is known to be responsible for long term-associative learning.
- Spreading graded activation is responsible for priming.
- Graded degrees of activation do correlate with degrees of approximation of a given experience.
- The same neural structures are activated when related concepts or percepts are evoked.
- Inhibitory interneurons are responsible for the repression of post-synaptic activation.

In the next section, I will show how BNs can be grounded in neural structure. This will complete the bridge from behavior, through a computational modeling level, to a plausible neural substrate.

4. A neural rendition of Belief Networks

Recall that BNs were introduced in Chapter 5 to make detailed explanations for and predictions of cognitive behaviors. BNs display system-level properties that make them particularly neurally plausible. First, just like a neuron, all of a BN node's information is stored locally: everything a node needs to compute the effects of events elsewhere in the network is available in the node representation itself. Inference in BNs is performed through the propagation of beliefs from one node to another, in a fashion similar to the propagation of activation in neural systems. Finally, the result of inference in a BN is a probabilistic result, which is analogous to the graded output of a neuron or batch of neurons responding to an input over time.

You may remember that at the outset of this thesis, I expressed the goal of identifying a computational architecture that could capture cognitive and linguistic phenomena using mechanisms that mapped to the neural mechanisms believed to be responsible for them. That is, there should be direct mappings between the BN mechanisms responsible for each of associative learning of probabilistic correlations, probabilistic knowledge, cross-modal information combination, partial productivity, and interactions between factors and their neural correlates. These are Hebbian learning, graded computation, converging afferent connections, overlapping neural structure, and inhibition and excitation, respectively.

There has been little effort on the part of BN researchers to seek out neural explanations for BNs. This is surprising since many researchers share the belief that no obvious obstacles stand in the way (e.g. Pearl and Russel 2001). But scalability as a limiting factor in mapping out full-blown neural models of BNs. For example, diagnostic inference is

particularly taxing for large-scale neural models of BNs. Luckily, the phenomena modeled in this thesis require BNs of very limited complexity. This class of BN has been shown, as we will see in the remainder of this chapter, to have a straightforward connectionist implementation.

One line of recent work articulates the neural plausibility of BNs through a structured connectionist model (Wendelken and Shastri 2000). Structured connectionist models (Feldman 1988) are computational models whose architecture and functioning closely replicate that of neural systems. They are composed of nodes, which usually represent neurons or groups of neurons, and connections between these nodes, which represent synapses or groups of synapses. Activation (the connectionist equivalent of electrochemical signals) is passed along connections between active nodes. It is thus through the grounding of the BN formalism in a structured connectionist model, itself directly mapped to neural structure, that the BN model gains its neural plausibility.

The connectionist realization of BNs presented by Wendelken and Shastri (2000) is part of a larger structured connectionist model of structured knowledge known as SHRUTI (Shastri and Ajjanagadde 1993). Although it is somewhat tangential to the main point of this dissertation, a summary description of a SHRUTI rendering of BNs is essential to the demonstration that they satisfy the isomorphism constraint (the requirement that a bridging theory's internal structure map neural explanations directly to their behavioral consequences.) Readers who are willing to take my word for it that BNs can be realized in a connectionist model can skip to the chapter's last paragraph for a summary.

The SHRUTI system represents relational knowledge through internally and externally connected *focal clusters* of nodes, where each cluster represents a proposition (as seen in Figure 9). Clusters have a number of nodes that are distinguished by their function.

A cluster's collector node (+) represents the amount of evidence that has been collected in favor of the predicate expressed by the cluster. The enabler node (?) indicates the strength of search for information about that predicate, that is, how actively verification of the predicate is being sought out.

Nodes in SHRUTI have variable degrees of activation. The activation values for nodes are analogous to the activation state of a neuron. When they are used to represent causal relational knowledge, node values are interpreted as probabilities. In other words, the value of the collector node (+) is the probability of the proposition that that node's focal cluster represents. A connection between the enabler node (?) and the collector node (+) indicates the probability of the proposition given no focal-cluster-external knowledge, in other words, the prior probability of the proposition.

Causal relations between two or more nodes are represented through connections between the enabler nodes of the two predicates and between their collector nodes. Taking two nodes, A and B, the connection from +A to +B represents the degree to which predicate B is true, given that A and only A is known to be true. And for reasons beyond the present scope, the connection from ?B to ?A represents the probability of B given A and perhaps other knowledge. These relations are depicted in Figure 9.

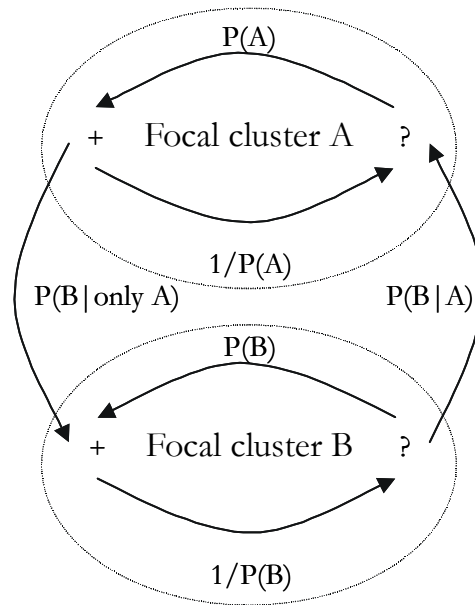


Figure 9 - SHRUTI representation of links between focal clusters A and B

This connectionist translation of BNs also includes a neurally plausible learning algorithm, based on Hebbian learning. Hebbian learning, as you will remember, simply involves the strengthening of connections between neurons that are co-active. In the case of causal learning in a SHRUTI-based BN model, connections need to be strengthened unidirectionally. In other words, if A is a cause of B, then the the +A to +B connection and ?B to ?A connection must be strengthened, but not the +B to +A or the ?A to ?B connections. Shastri and Wendelken’s solution is to bootstrap off of the temporal relationship of prototypical causal relations. In general, a cause temporally precedes an effect. This is certainly true for causes and effects in the networks we have described so far in this chapter; identifying a meaning to be expressed precedes the selection of a word to express that meaning. Because causes precede effects, a Hebbian learning algorithm that takes time of firing into account can strengthen connections asymmetrically; it can strengthen connections from cause to effect, and not the reverse. For example, if +A is

active before +B, then such a Hebbian learning mechanism can strengthen the link from +A to +B but not from +B to +A.

Finally, each node in SHRUTI computes its own activation as a function of the inputs that it collects. The functions that it uses to perform this calculation, Evidence Combination Functions (EFCs) can differ from node to node. EFCs can, for example, take the maximum value of their inputs, or the average of their inputs. Although such has not been implemented, there is no technical reason why EFCs could not also perform interactive calculations. An EFC could thus approximate a CPT that defined some interaction between factors.

<i>Neural Phenomenon</i>	<i>Connectionist BN mechanism</i>	<i>BN mechanism</i>	<i>Cognitive/Linguistic Phenomenon</i>
Hebbian learning	Hebbian-based learning	Structure induction and CPT building	learning probabilistic correlations
graded computation	graded activation and information passing	graded activation and information passing	probabilistic knowledge
converging afferent connections	links between nodes in different focal clusters	links between nodes	cross-modal information combination
overlapping neural structure	shared nodes	shared nodes	priming of partially productive units
inhibition and excitation	EFCs	CPTs	interactions between factors

Figure 10 – Mappings between mechanisms at four levels of representation

As we have seen, the five neural mechanisms cited in section 2 are each realized through different machinery in this connectionist implementation of BNs. And each connectionist mechanism maps to a unique BN mechanism. The specific mappings between the biological and cognitive/linguistic levels as mediated by the BN models described above are presented in Figure 10.

Chapter 7. Conclusion

Outline

1. Debates
2. Future directions
3. Summary

The trouble with our times is that the future is not what it used to be.

Paul Valery

1. Debates

The study of language is notoriously divisive, and any description of the field of Linguistics inevitably begins with a characterization of camps that differ along some dimension. There's the nativists versus the empiricists, the symbolists against the connectionists, the modularists and the integrationists. We can describe these three pairs of positions as defining the endpoints of three continua: nativism, symbolism, and modularism.

The perspective taken in the present work is one of methodological and empirical flexibility. The computational models developed to capture phonaestheme and liaison behavior were formulated in an architecture that is sufficiently flexible to capture a range of points along these three continua. While BNs are able to learn structure and probabilities from data they are exposed to, they can also be furnished with prior structure. This prior structure could presumably be of an innate nature. That stated, it doesn't seem likely that the subtle, statistical, interacting effects on liaison that we saw in this thesis could be innate.

Rather, just like the BNs that were subjected to thousands of tokens of liaison, so humans most likely extract statistical correlations from ambient language.

The nodes of a BN are inherently symbolic - they represent propositions with a limited set of possible values. At the same time, though, they share structural and procedural characteristics with connectionist models. Local inference, spreading information, and parallel computation all characterize both connectionist and BN systems. The solution presented here, then is a sort of compromise between symbolic rule systems and unstructured probability machines.

Finally, the BN model of liaison proposed above treats a phenomenon that is inherently integrated - it require knowledge from different domains to be brought to bear on a problem simultaneously. Nevertheless, modular phenomena, too, can be modeled without difficulty in a BN system. From a broader perspective, though, I believe that a model's degree of cross-modularity correlates directly with its potential breadth. Restricting modeling enterprises strictly to modular phenomena simply limits the range of phenomena that can be considered.

2. Future directions

The present work creates more new questions than it provides answers. I have already identified some of these that relate to phonaesthemes in Chapter 2. Among these are the precise relationship between phonaesthetic priming and morphological priming, and the effects of degree of semantic overlap and frequency on phonaesthetic priming.

Another empirical question is raised by the observation of interactions between factors. Since these have rarely been studied, very little is known about the range of their

distribution. If, as suggested in Chapter 5, there are limits on the human ability to pick up on interactions between factors, then these limits need to be established and incorporated as constraints on models like the ones presented in this thesis.

A final line of empirical inquiry involves the place of the models described in this thesis in a larger theory of phonology. Two issues are relevant to this discussion: the relation between static and dynamic knowledge and the relation between categorical and non-categorical knowledge.

The BNs presented in this thesis, with one exception, capture only purely static knowledge. That is, they are representations of a particular set of causal relations between variables at an instant. And yet, much of what we know about the details of the phonological forms of our language is inherently dynamic in that the variables that describe the vocal tract or perceptual space change over time. This dynamicity can be modeled to a limited extent through the use of DBNs, as shown in Chapter 5. But the detailed relative and absolute temporal information that characterizes phonetics cannot be captured in models that only break time up into regular slices.

This sets up the question, then, of exactly what BNs might be an abstraction over, if they are incapable of representing the lower level phonetic detail that they could be learned from. A solution to this problem in the domain of motor control has been proposed by Baily (1997) and Narayanan (1997). This work breaks the task of motor control into two components - the low-level, dynamic, active motor knowledge on the one hand, and parameterized abstractions on the other. In Baily's and Narayanan's models, BNs are used to represent the parameterized abstractions, while another formalism, more appropriate for dynamic systems, is conscripted to model motor control.

A similar solution could prove fruitful in the domain of phonology. If phonetic details are in fact abstracted over in parameterizations, then this might serve as a partial explanation for how phonologization can effectively take phonetically natural processes and generalize them such that they are no longer natural. Once an abstraction has been made, it can take on its own, abstract existence. Exactly how phonologization would proceed in a hybrid model remains to be determined.

The second issue pertinent to incorporating this model into a larger picture of phonology is the relation between categorical and non-categorical knowledge. As was suggested in Chapter 5, categorical knowledge can be very easily handled in BNs. The conditional probability tables of BNs need not include exclusively probabilistic effects - values of 1 and 0 are also possible, as seen for example in the phonaestheme model. BNs additionally offer the possibility of hybrid categorical and non-categorical effects.

3. Summary

The main goal of this thesis was to contribute to the growing catalogue of non-categorical effects on phonology. The burgeoning literature on this subject has prior to this thesis not considered the psychological reality of two important patterns. Phonaesthemes and social correlates of phonology were shown here through perception experiments to be psychologically real.

Also neglected in this literature, and in more widespread sociolinguistic work, is the study of interactions between factors. Speaker of French were shown in this thesis to be unconsciously aware of interactions between two factors that influence French liaison: the age of the speaker and the grammatical class of the liaison word. Interactions between

factors are interesting from a theoretical perspective. Only a restricted class of model can capture interactions between different knowledge domains, and I have shown that Belief Networks provide a useful architecture for creating such models.

Many of the phenomena documented in this work have a neural basis, and biological explanations for five cognitive and linguistic phenomena were given. Moreover, since Belief Networks are neurally plausible and hold to the isomorphism constraint, they act as a bridging model that allows neural explanations for cognitive and linguistic behavior to be formalized. For the cases studied in this thesis, the details of how the brain works matter to language.

References

- Abelin, Asa. 1999. *Studies in Sounds Symbolism*. Ph.D. Thesis, Department of Linguistics, Goteborg University.
- Abney, Steven. 1996. *Statistical methods and linguistics*. In J. Klavans and P. Resnik (Eds.) *The Balancing Act*. Cambridge: MIT Press.
- Albright, Adam and Bruce Hayes. 2000. *Distributional encroachment and its consequences for morphological learning*. In A. Albright and T. Cho (Eds.) *UCLA Working Papers in Linguistics 4 (Papers in Phonology 4)*:179-190.
- Anderson, Stephen. 1971. *On the role of deep structure in semantic interpretation*. *Foundations of Language* 6: 197-219.
- Anderson, Steven. 1992. *A-Morphous Morphology*. *Cambridge Studies in Linguistics* 62. Cambridge: Cambridge University Press.
- Anttila, Arto. 1997. *Deriving Variation from Grammar*. In F. Hinskens et al. (Eds.) *Variation, Change, and Phonological Theory*. Amsterdam: John Benjamins: 35-68.
- Arnold, Jennifer, Thomas Wasow, Anthony Losongco, and Ryan Ginstrom. 2000. *Heavyiness vs. newness: the effects of structural complexity and discourse status on constituent ordering*. *Language* 76(1): 28-55.
- Aronoff, Mark. 1976. *Word Formation in Generative Grammar*. *Linguistic Inquiry Monograph One*. Cambridge, MA: MIT Press.
- Ashby, William. 1981. *French Liaison as a Sociolinguistic Phonemenon*. In Cressey-William-W. (Ed.); Napoli-Donna-Jo (Ed.). *Linguistic Symposium on Romance Languages*: 9. Washington, DC : Georgetown UP: 46-57.

- Bailey, David, Jerry Feldman, Srinivasa Narayanan, George Lakoff. 1997. Modeling Embodied Lexical Development , Proceedings of the Nineteenth Annual Meeting of the Cognitive Science Society COGSCI-97, Aug 9-11, Stanford: Stanford University Press.
- Bates, Elizabeth, Antonella Devescovi, Arturo Hernandez, and Luigi Pizzamiglio. 1996. Gender Priming in Italian. *Perception and Psychophysics* Oct, 58:7: 992-1004.
- Bergen, Benjamin K. 2000a. Probabilistic associations between sound and meaning: Belief Networks for modeling phonaesthemes. Paper presented at the 5th Conference on Conceptual Structure, Discourse, and Language, Santa Barbara.
- Bergen, Benjamin K. 2000b. Probability in phonological generalizations: modeling French optional final consonants. In Alan Yu et al. (Eds.) Proceedings of the 26th Annual Meeting of the Berkeley Linguistics Society. Berkeley: BLS
- Berko, Jean. 1958. The child's acquisition of English morphology, *Word* 14: 150-177.
- Bianchi, Marta. 1981. La liaison en Frances. *Revista de Lingüística Teórica y Aplicada* 19: 67-74.
- Biber, Douglas, Susan Conrad, and Randi Reppen. 1998. *Corpus linguistics: Investigating language structure and use*. Cambridge: Cambridge University Press.
- Blust, Robert. 1988. *Austronesian root theory: an essay on the limits of morphology*. Philadelphia: J. Benjamins Pub. Co.
- Bochner, Harry. 1993. *Simplicity in generative morphology*. New York: Mouton de Gruyter
- Bod, Rens. 1998. *Beyond grammar*. Stanford: CSLI.
- Boersma, Paul and Bruce Hayes. 2001. Empirical tests of the Gradual Learning Algorithm. *Linguistic Inquiry* 32:45-86.

- Bolinger, Dwight. 1949. The Sign Is Not Arbitrary. *Boletín del Instituto Caro y Cuervo*, 5: 52-62.
- Bolinger, Dwight L., 1980, *Language; The Loaded Weapon: The Use and Abuse of Language Today*. New York: Longman.
- Booij, Geert and Daan de Jong. 1987. The domain of liaison: theories and data. *Linguistics* 25: 1005-1025.
- Bradlow, Ann, Lynne Nygaard, and David Pisoni. 1999. Effects of talker, rate, and amplitude variation on recognition memory for spoken words. *Perception & Psychophysics* 61(2):206-219.
- Bromberger, Sylvain and Morris Halle. 1989. Why Phonology is Different. *Linguistic Inquiry* 20: 51-70.
- Burzio, Luigi. 1994. *Principles of English stress*. Cambridge: Cambridge University Press.
- Bybee, Joan. 1999. Usage based phonology. M. Darnell, E. Moravcsik, F. Newmeyer, M. Noonan and K. Wheatley (Eds.) *Functionalism and formalism in linguistics, volume I: General papers*. Amsterdam: John Benjamins. 211-242.
- Bybee, Joan. 2000. The Phonology of the Lexicon: Evidence from Lexical Diffusion. In Barlow, Michael and Suzanne Kemmer (Eds.). *Usage-Based Models of Language*. Stanford: CSLI: 65-85.
- Bybee, Joan and Carol Moder. 1983. Morphological Classes as Natural Categories. *Language* 59:251-270.
- Bybee, Joan and Dan Slobin. 1982. Rules and Schemas in the development and use of the English past tense. *Language* 58: 265-289.
- Bybee, Joan and Joanne Scheibman. 1999. The effect of usage on degrees of constituency the reduction of *don't* in English. *Linguistics* 37(4): 575-596.

- Campbell, Ruth and Derek Besner. 1981. THIS and THAP--constraints on the pronunciation of new, written words. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 1981 Nov, v33 (n4):375-396.
- Carroll, Lewis. 1897. *Alice's adventures in wonderland, and Through the looking-glass*. New York, Macmillan.
- Cassidy, Kimberly and Michael Kelly. 1991. Phonological information for grammatical category assignments. *Journal of Memory and Language* 30: 348-369.
- Cassidy, Kimberly, Michael Kelly, and Lee-at Sharoni. 1999. Inferring gender from name phonology. *Journal of Experimental Psychology: General* 128(3): 362-381.
- Cedergreen, H. 1973. *The interplay of social and linguistic factors in Panama*. Ph.D. Dissertation, Cornell University.
- Chambers, Jack. 1995. *Sociolinguistic theory: Linguistic variation and its social significance*. Oxford: Blackwell.
- Chollet, C., J.-L. Chochard, A. Constantinescu, C. Jaboulet and Ph. Langlais. 1996. *Swiss french polyphone and polyvar: Telephone speech database to model inter- and intraspeaker variability*. Technical Report RR-96-01, IDIAP, Martigny, April 1996.
- Chomsky, Noam. 1964. "Current issues in linguistic theory," in Jerry A. Fodor and Jerrold J. Katz, Eds., *The Structure of Language*, Prentice-Hall, Englewood Cliffs, NJ
- Chomsky, Noam. 1981. *Lectures on Government and Binding: The Pisa Lectures*. Dordrecht: Foris.
- Chomsky, Noam. 1995. *The Minimalist Program*. Cambridge: MIT Press.
- Chomsky, Noam and Morris Halle. 1968. *Sound patterns of English*. Harper and Row: New York.

- Davis, S. and Kelly, Michael 1997. Knowledge of the English noun-verb stress difference by native and nonnative speakers. *Journal of Memory and Language* 36:445-460.
- De Jong, Daan. 1989. A multivariate analysis of French Liaison. In M. Schouten and P. van Reenen (Eds.) *New Methods in Dialectology*. Dordrecht: Foris: 19-34.
- De Jong, Daan, E. Poll, and W. Woudman. 1981. *La liaison: l'influence sociale et stylistique sur l'emploi de la liaison dans le français parlé à Tours*. Unpublished Masters Thesis, University of Groningen.
- Delattre, P. 1966. *Studies in French and comparative phonetics*. The Hague: Mouton: 39-62.
- Dell, Gary, L. Burger, and W. Svec. 1997. Language Production and Serial Order: A Functional Analysis and a Model. *Psychological Review* 104(1):123-147.
- Dell, Gary, and C. Juliano. (1996). Phonological encoding. In T. Dijkstra and K. DeSmedt (Eds.), *Computational psycholinguistics: Symbolic and connectionist models of language processing*. London: Harvester-Wheatsheaf.
- Di Sciullo, Anna-Maria, and Edwin Williams. 1987. *On the Definition of Word*. Cambridge: MIT Press.
- Duez, Danielle. 1985. Perception of silent pauses in continuous speech. *Language and speech*, 28:4: 377-389.
- Encreve, Pierre. 1988. *La liaison avec et sans enchainement: phonologie tridimensionnelle et usages du francais*. Paris: Editions du Seuil.
- Fasold, Ralph. 1978. Language variation and linguistic competence. In David Sankoff (Ed.) *Linguistic variation: models and methods*. Academic Press: New York:85-95.
- Fasold, Ralph. 1996. The quiet demise of variable rules. In R. Singh (Ed.) *Towards a critical sociolinguistics*. Amsterdam: John Benjamins: 79-98.

- Feldman, Jerome. 1988. Connectionist Representation of Concepts. In D. Waltz and J. Feldman (Eds.) *Connectionist Models and Their Applications*. Ablex Publishing Company.
- Feldman, Laurie, 1995. *Morphological aspects of language processing*. Hilldale, NJ: Erlbaum.
- Feldman, Laurie and Emily Soltano. 1999. Morphological Priming: The Role of Prime Duration, Semantic Transparency, and Affix Position. *Brain and Language* 68, 33-39.
- Fillmore, Charles. 1968. The case for case. In Emmon Bach and R. Harms (Eds.) *Universals in linguistic theory*. New York: Holt, Rinehart, and Winston: 1-88.
- Firth, J. R., 1930, *Speech*. London: Oxford University Press.
- Fodor, Jerry. 1983. *Modularity of mind*. Cambridge, Mass.: MIT Press.
- Francis, W. Nelson and Henry Kucera. 1982. *Frequency analysis of English usage : lexicon and grammar*. Boston: Houghton Mifflin.
- Frisch, Stephen and Bushra Zawaydeh. 2001. The psychological reality of OCP-Place in Arabic. *Language* 77.
- Gall, Franz and Johann Spurzheim. 1810. *Anatomie et physiologie du système nerveux en général, et du cerveau en particulier, avec des observations sur la possibilité de reconnoître plusieurs dispositions intellectuelles et morales de l'homme et des animaux, par la configuration de leurs têtes*. Paris: Schoell.
- Garrote, Ernesto. 1994. Características fonéticas y funcionales de la "liaison" en frances. *Revista de Lingüística Teórica y Aplicada*, 32:101-113.
- Goldinger, Stephen D.; Kleider, Heather M.; Shelley, Eric. 1999. The marriage of perception and memory: Creating two-way illusions with words and voices. *Memory & Cognition* 27(2):328-338.

- Good, Jeff. 1998. H-aspiré Under Close Scrutiny: How Regular Can an Irregular Phenomenon Be? Unpublished MA thesis. University of Chicago.
- Grady, Joseph., Todd Oakley, and Seana Coulson. 1999. Conceptual Blending and Metaphor. In G. Steen and R. Gibbs (Eds.) *Metaphor in cognitive linguistics*. Philadelphia: John Benjamins.
- Grammont, Maurice. 1938. *La prononciation française*. Paris: Librairie Delgrave.
- Grimes, Barbara (Ed.). 2000. *Ethnologue*, 14th Edition. Dallas: Summer Institute of Linguistics.
- Guy, Gregory. 1977. A new look at -t, -d deletion. In R. Shuy and R. Fasold (Eds.) *Studies in language variation*. Washington, D.C.: Georgetown University Press:1-11.
- Guy, Gregory. 1980. Variation in the group and the individual: The case of final stop deletion. In Labov, William (Ed.) *Locating language in time and space*. New York: Academic Press: 1-36.
- Guy, Gregory. 1991. Explanation in variable phonology: An exponential model of morphological constraints. *Language Variation and Change* 3: 1-22.
- Halle, Morris; and Alec Marantz. 1993. Distributed Morphology and the Pieces of Inflection. Pg. 111-176 in K. Hale and S. Bromberger (Eds.) *The View from Building 20*. Cambridge, MA: MIT Press.
- Hawkins, John. 1994. *A Performance Theory of Order and Constituency*. Cambridge University Press: Cambridge.
- Hayes, Bruce. 1980. *A metrical theory of stress rules*. M.I.T. Ph.D. Dissertation.
- Hayes, Bruce. 1998. On the Richness of Paradigms, and the Insufficiency of Underlying Representations in Accounting for them. Paper presented at Stanford U.

- Hayes, Bruce. 2000. Gradient Well-formedness in Optimality Theory. In Joost Dekkers, Frank van der Leeuw and Jeroen van de Weijer, (Eds.), *Optimality Theory: Phonology, Syntax, and Acquisition*, Oxford University Press: 88-120.
- Hebb, Donald. 1949. *The Organization of Behaviour*. New York: John Wiley & Sons.
- Hutchins, Sharon Suzanne. 1998. *The Psychological Reality, Variability, and Compositionality of English Phonesthemes*. Ph.D, Dissertation. Emory University.
- Hyvarinen, J. and A. Poranen. 1978. Movement-sensitive and direction and orientation-selective cutaneous receptive fields in the hand area of the post-central gyrus in monkeys. *Journal of Physiology* 283: 523-527.
- Inkelas, Sharon. 1995. The consequences of optimization for underspecification. In E. Buckley and S. Iatridou (Eds.), *Proceedings of the Twenty-Fifth Northeastern Linguistics Society*. Amherst: GLSA.
- Inkelas, Sharon and Draga Zec. 1995. syntax-phonology interface. In John Goldsmith (Ed.) *The Handbook of Phonological Theory*. Oxford: Blackwell: 535-549.
- Jackendoff, Ray. 1997. *The architecture of the language faculty*. Cambridge, Mass.: MIT Press
- Jaeger, Jeri and John Ohala. 1984. On the structure of phonetic categories. In Claudia Brugman et al. (Eds.) *Proceedings of the Tenth Annual Meeting of the Berkeley Linguistics Society*. Berkeley: BLS: 15-26.
- Jensen, Finn V. 1996. *An introduction to Bayesian networks*. New York: Springer.
- Jespersen, Otto. 1922. Symbolic value of the vowel I. *Philologica* 1.1-19.
- Jespersen, Otto. 1942. *A modern English grammar on historical principles, IV: Morphology*. Copenhagen: Munksgaard.

- Johnson, Keith. 1997. Speech Perception Without Speaker Normalization: An Exemplar Model. In Keith Johnson and John Mullennix (Eds.) Talker Variability in Speech Processing. San Diego: Academic Press: 145-165.
- Jordan, Michael I. (Ed.). 1998. Learning in Graphical Models. Kluwer Academic Press.
- Jurafsky, Daniel. 1996. A probabilistic model of lexical and syntactic access and disambiguation. *Cognitive Science* 20:137-194.
- Kaisse, Ellen. 1985. Connected Speech. San Diego: Academic Press.
- Kandel, Eric. 1991. Cellular mechanisms of learning and the biological basis of individuality. In Eric Kandel, James Schwartz and Thomas Jessel (Eds.) Principles of Neural Science. Norwalk: Appleton and Lange: 1009-1031.
- Kandel, Eric, James Schwartz, and Thomas Jessel (Eds.). 2000. Principles of neural science, 4th Ed. New York: McGraw-Hill.
- Kawamoto, A., W. Farrar, and M. Overbeek. 1990. Effect of syntactic context on naming bisyllabic words. Poster Presented at the 31st annual meeting of the Psychonomic society, New Orleans.
- Kay, Paul. 1978. Variable rules, community grammar, and linguistic change. In David Sankoff (Ed.) Linguistic variation: models and methods. Academic Press: New York:71-82.
- Kelly, Michael. 1988. Phonological biases in grammatical category shifts. *Journal of Memory and Language* 27:343-358.
- Kelly, Michael. 1992. Using sound to solve syntactic problems: The role of phonology in grammatical category assignments. *Psychological Review* 99(2):349-364.

- Kelly, Michael. 1996. The role of phonology in grammatical category assignments. In James Morgan and Katherine Demuth (Eds.) *Signal to syntax: bootstrapping from speech to grammar in early acquisition*. Lawrence Erlbaum, Mahwah NJ: 249-262.
- Kelly, Michael and Kathryn Bock. 1988. Stress in time. *Journal of Experimental Psychology: Human Perception and Performance* 14: 389-403.
- Kelly, Michael, Ken Springer, and Frank Keil. 1990. The relation between syllable number and visual complexity in the acquisition of word meanings. *Memory and Cognition* 18(5): 528-536.
- Kemmer, Suzanne and Michael Israel. 1994. Variation and the Usage-Based Model. *CLS* 30.
- Kempey, S. and J. Morton. 1982. The Effects of Priming with Regularly and Irregularly Related Words in Auditory Word Recognition. *British Journal of Psychology* 73, 441-445.
- Kessler, Brett. 1994. Sandhi and syllables in Classical Sanskrit. In E. Duncan, D. Farkas, and P. Spaelti, (Eds.), *The proceedings of the Twelfth West Coast Conference on Formal Linguistics*. Stanford: CSLI: 35-50.
- Koller, Daphne and Ave Pfeffer. Object-oriented Bayesian Networks. In *Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence*: 302-313.
- Kucera, Henry and W. Nelson Francis. 1967. *Computational analysis of present-day American English*. Providence, Brown University Press.
- Labov, William, P. Cohen, C. Robins and J. Lewis. 1968. *A study of the non-standard English of Negro and Puerto Rican speakers in New York City*. Final report, Cooperative Research Project 3288. 2 vols. Philadelphia: U.S. Regional Survey.
- Labov, William. 1963. The social motivation of a sound change. *Word* 19:273-309.

- Labov, William. 1972. *Sociolinguistic Patterns*. Philadelphia: University of Pennsylvania Press.
- Lakoff, George. 1993. How Metaphor Structures Dreams: The Theory of Conceptual Metaphor Applied to Dream Analysis. *Dreaming* 3(2): 77-98.
- Langacker, Ronald. 1991. *Concept, Image, and Symbol: The Cognitive Basis of Grammar*. New York: Mouton De Gruyter.
- Lieber, Rochelle. 1980. *On the Organization of the Lexicon*. Ph.D. dissertation, MIT; published 1981 by the Indiana University Linguistics Club.
- MacWhinney, B., & Leinbach, J. 1991. Implementations are not Conceptualizations: Revising the Verb Learning Model. *Cognition*.
- Maess, Burkhard, Stefan Koelsch, Thomas Gunter, and Angela Friederici. 2001. Musical syntax is processed in Broca's area: An MEG study. *Nature Neuroscience* 4(5):540-545.
- Magnus, Margaret. 2000. *What's in a Word? Evidence for Phonosemantics*. Ph.D. Dissertation, University of Trondheim, Norway.
- Marchand, Hans. 1959. Phonetic symbolism in English word formations. *Indogermanische Forschungen* 64: 146-168.
- Marr, David. 1982. *Vision*. New York: W.H. Freeman & Co.
- Martinet, Jeanne. 1988. Un Traitement fonctionnel de la liaison en francais. *Folia Linguistica Acta Societatis Linguisticae Europaeae*, 22(3-4): 293-299.
- Mayeux, Richard and Eric Kandel. 1991. Disorders of language: the aphasias. In Eric Kandel, James Schwartz and Thomas Jessel (Eds.) *Principles of Neural Science*. Norwalk: Appleton and Lange: 839-851.

- McCune, Keith. 1983. *The Internal Structure of Indonesian Roots*, Ph.D. dissertation, University of Michigan.
- Meunier, Fanny and Juan Segui. 1999. Morphological Priming Effect: The Role of Surface Frequency. *Brain and Language* 68: 54-60.
- Morel, Eliane. 1994. The treatment of liaison in the child: experimental studies. *Travaux neuchatelois de linguistique*, 21:85-95.
- Morgan, James, Rushen Shi, and Paul Allopenna. 1996. In James Morgan and Katherine Demuth (Eds.) *Signal to syntax: bootstrapping from speech to grammar in early acquisition*. Lawrence Erlbaum, Mahwah NJ: 263-281.
- Morin, Yves-Charles and Jonathan Kaye. 1982. The syntactic bases for French liaison. *Journal of Linguistics* 18: 291-330.
- Nagy, Naomi and Bill Reynolds. 1997. Optimality Theory and word-final deletion in Faetar. *Language Variation and Change* 9: 37-55.
- Narayanan, Srinivas and Daniel Jurafsky. 1998. Bayesian Models of Human Sentence Processing. *Proceedings of CogSci98*.
- Newman, Stanley. 1944. *Yokuts Language of California*. New York: Viking Fund.
- Nygaard, Lynne C.; Pisoni, David B. 1998. Talker-specific learning in speech perception. *Perception & Psychophysics* 60(3):355-376.
- Ohala, John. 1984. An ethological perspective on common cross-language utilization of F0 of voice. *Phonetica* 41:1-16.
- O'Seaghdha, Padraig and Joseph Marin. 1997. Mediated Semantic-Phonological Priming: Calling Distant Relatives. *Journal of Memory and Language* 36: 226-252.
- Partee, Barbara. 1965. *Subject and object in Modern English*. New York: Garland.

- Pearl, Judea and Stuart Russel. 2001. Bayesian networks. In M. Arbib (Ed.), Handbook of Brain Theory and Neural Networks, MIT Press.
- Pearl, Judea. 1988. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. San Mateo, CA: Morgan Kaufman Publishers.
- Pierrehumbert, J. (in press) Stochastic phonology. GLOT
- Pierrehumbert, Janet. 2000. What people know about sounds of language. Studies in the Linguistic Sciences 29(2).
- Pinker, Steven. 1989. Learnability and cognition: the acquisition of argument structure. Cambridge, Mass.: MIT Press.
- Pinker, Steven. 1994. The language instinct. New York: W. Morrow and Co.
- Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a Parallel Distributed Processing Model of language acquisition. Cognition 29: 73-193.
- Plaut, David and Laura Gonnerman. 2000. Are non-semantic morphological effects incompatible with a distributed connectionist approach to lexical processing? Language and Cognitive Processes 15: 445-485.
- Prince, Alan and Paul Smolensky. 1993. Optimality Theory: Constraint Interaction in Generative Grammar. Rutgers University Center for Cognitive Science Technical Report 2.
- Prince, Alan and Paul Smolensky. 1997. Optimality: From neural networks to universal grammar. Science 275: 1604-1610.
- Radford, Andrew. 1990. Syntactic theory and the acquisition of English syntax. Oxford: Basil Blackwell.

- Rand, David and David Sankoff. M.s. GoldVarb Version 2: A Variable Rule Application for the Macintosh. Available at:
<http://www.crm.umontreal.ca/~sankoff/GoldVarbManual.Dir/>
- Rhodes, Richard. 1981. On the Semantics of Ojibwa Verbs of Breaking, in W. Cowan, (Ed.) Papers of the Twelfth Algonquian Conference, Ottawa: Carleton University Press.
- Rhodes, Richard. 1994. Aural Images, in Leanne Hinton, Johanna Nicols and John J. Ohala, (Eds.), Sound Symbolism, Cambridge University Press.
- Rice, Keren D. 1990. Predicting rule domains in the phrasal phonology. In Sharon Inkelas and Draga Zec (Eds.) The phonology-syntax connection. Chicago: Univ. of Chicago Press: 289-312
- Rummelhart, David, J. McClelland and the PDP Research Group. 1986. Parallel Distributed Processing. Cambridge: MIT Press.
- Sankoff, David. 1987. Variable Rules. In Ammon, Dittmar & Mattheier (Eds.) Sociolinguistics: An international handbook of the science of language and society, Vol. I: 984-997.
- Sankoff, Gillian. 1974. A quantitative paradigm for studying communicative competence. In R. Bauman and J. Sherzer (Eds.) The ethnography of speaking. London: Cambridge University Press: 18-49.
- Santerre, Laurent and Jean Milo. 1978. Diphthongization in Montreal French. In David Sankoff (Ed.) Linguistic variation: models and methods. Academic Press: New York:173-184.
- Saussure, Ferdinand de. 1916. Cours de linguistique generale. Paris: Payot.

- Schmid, Helmut. 1994. Probabilistic part-of-speech tagging using decision trees. In International Conference on New Methods in Language Processing Manchester, UK: pp. 44-49.
- Schwartz, Randal and Tom Christiansen. 1997. Learning PERL, 2nd Edition. Sebastopol, CA: O'Reilly & Associates.
- Selkirk, Elisabeth. 1974. French liaison and the X' notation. *Linguistic Inquiry* 5: 573-590.
- Selkirk, Elisabeth. 1978. On prosodic structure and its relation to syntactic structure. Unpublished Ms. University of Massachusetts, Amherst.
- Selkirk, Elisabeth. 1984. *Phonology and Syntax: The Relation between Sound and Structure*, Cambridge: MIT Press.
- Sereno, Joan. 1994. Phonosyntactics. In L. Hinton, J. Nichols, and J. Ohala (Eds.), *Sound Symbolism*. Cambridge: Cambridge University Press, 263-275.
- Shastri, Lokendra and V. Ajjanagadde. 1993. *From simple associations to systematic reasoning* Behavioral and Brain Sciences Vol. 16, No. 3, 417-494.
- Shastri, Lokendra and Carter Wendelken. 1999 Soft Computing in SHRUTI. In Proceedings of the Third International Symposium on Soft Computing.
- Sherman, Donald. 1975. Noun-verb stress alternation: An example of lexical diffusion of sound change in English. *Linguistics*, 159: 43-71.
- Siegel, Sidney and N. John Castellan, Jr. 1988. *Nonparametric statistics for the social sciences*. Second Edition. McGraw-Hill: New York.
- Smith, Jennifer. 1997. Noun faithfulness: On the privileged behavior of nouns in phonology. Available on the Rutgers Optimality Archive (#242).
- Stefanowitsch, Anatol. In Prep. Sound Symbolism in a Usage-Driven Model.

- Swingle, Kari. 1993. The role of prosody in right node raising. In Geoffrey Pullum and Eric Potsdam (Eds.) *Syntax at Santa Cruz*, Volume 2: 83-110.
- Tenenbaum, Joshua. 1999. Bayesian modeling of human concept learning. *Advances in Neural Information Processing Systems* 11. MIT Press.
- Tenenbaum, Joshua. 2000. Rules and similarity in concept learning. *Advances in Neural Information Processing Systems* 12. MIT Press: 59-65.
- Tenenbaum, Joshua and Thomas Griffiths. 2001. Structure learning in human causal induction. In *Advances in Neural Information Processing Systems* 13.
- Thompson-Schill, Sharon L., Kenneth Kurtz, and John Gabrieli. 1998. Effects of Semantic and Associative Relatedness on Automatic Priming. *Journal of Memory and Language* 38: 440-458.
- Tienson, John. 1987. *An Introduction to Connectionism*. The Southern Journal of Philosophy. Vol. 26 Suppl.
- Tomasello, Michael. In press. A usage-based approach to child language acquisition. *Proceedings of the Berkeley Linguistics Society*.
- Tranel, Bernard. 1981. *Concreteness in generative phonology: evidence from French*. Berkeley: University of California Press.
- Trudgill, Peter. 1972. Sex, Covert Prestige and Linguistic Change in the Urban British English of Norwich. *Language in Society* 1(2): 179-95.
- Trudgill, Peter. 1983. *On dialect: social and geographical perspectives*. Oxford: Blackwell.
- Tucker, G. Richard, Wallace Lambert, Andre Rigault, and Norman Segalowitz. 1968. A Psychological Investigation of French Speakers' Skill with Grammatical Gender. *Journal of Verbal Learning & Verbal Behavior* 7(2): 312-316.

- Wahlster, Wolfgang (Ed.). 2000. *Verbmobil: foundations of speech-to-speech translation*.
New York: Springer
- Walker, Douglas. 1980. Liaison and rule ordering in Canadian French phonology.
Linguisticae Investigationes IV:1: 217-222.
- Wallis, John. 1699. *Grammar of the English Language (Fifth Edition)*. Oxford: L. Lichfield.
- Wasow, Thomas. 1997. Remarks on grammatical weight. *Language variation and change*
9:81-105.
- Watbled, Jean-Philippe. 1991. Les processus de sandhi externe en français de Marseille.
French language studies 1: 71-91.
- Webster's Seventh Dictionary. 1965. Springfield, MA: Merriam-Webster Inc., 1965.
- Wellman, Michael. 1990. Fundamental concepts of qualitative probabilistic networks.
Artificial Intelligence 44:257-303.
- Wendelken, Carter. and Lokendra Shastri. 2000. Probabilistic Inference and Learning in a
Connectionist Causal Network. In *Proceedings of the Second International
Symposium on Neural Computation*. Berlin.
- Whitney, William Dwight. 1891. *A Sanskrit grammar including both the classical language,
and the older dialects, of Veda and Brahmana*. Leipzig, Breitkopf and Hartel:
Boston.
- Wolfram, Walt and D. Christian. 1976. *Appalachian Speech*. Arlington, VA: Center for
Applied Linguistics.
- Wolfram, Walt. 1973. Sociolinguistic aspects of assimilation: Puerto Rican English in East
Harlem. Arlington, VA: Center for Applied Linguistics.
- Zimmer, Karl. 1969. Psychological Correlates of Some Turkish Morpheme Structure
Conditions. *Language* 45: 309-321.

- Zubin, David and Klaus-Michael Kopcke. 1986. Gender and Folk Taxonomy: The Indexical Relation between Grammatical and Lexical Categorization. In Craig-Colette (Ed.), Noun Classes and Categorization. Amsterdam: Benjamins: 139-180.
- Zwicky, Arnold M. 1985. How to describe inflection. BLS 11: 372-386.
- Zwitserslood, Pienie. 1996. Form Priming. Language and Cognitive Processes 11(6), 589-596.

Appendix: Phonaestheme examples

These words were automatically extracted from the Brown corpus, and were assigned phonaesthetic meanings if one of their senses in Webster's 7th referred to 'LIGHT' or 'VISION' for the case of *gl-* and 'NOSE' or 'MOUTH' for *sn-*.

gl-	glazes	gloom	sneering
glance	glazing	gloomily	sneers
glanced	gleam	gloomy	sneezed
glances	gleamed	gloss	sneezing
glancing	gleaming	glossy	snickered
glare	glimmer	glow	sniff
glared	glimmering	glowed	sniffed
glaring	glimpse	glowered	sniffing
glaringly	glimpsed	glowering	sniggered
glass	glimpses	glowing	snippy
glass-bottom	glint	glows	snivelings
glasses	glinted		snoring
glass-fiber	glinting	sn-	snorkle
glassless	glisten	snack	snort
glass-like	glistened	snacks	snorted
glassy	glistening	snarled	snout
glaucoma	glitter	snarling	snuffboxes
glaze	glittered	sneer	snuffed
glazed	glittering	sneered	snuffer

