


Michael Kenstowicz

PHONOLOGY IN GENERATIVE GRAMMAR



Phonology in Generative Grammar

Michael Kenstowicz


BLACKWELL
Cambridge MA & Oxford UK

Copyright © Michael Kenstowicz 1994

The right of Michael Kenstowicz to be identified as author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

First published 1994
Reprinted 1994, 1995

Blackwell Publishers, the publishing imprint of
Basil Blackwell Inc.
238 Main Street
Cambridge, Massachusetts 02142, USA

Basil Blackwell Ltd
108 Cowley Road
Oxford OX4 1JF
UK

All rights reserved. Except for the quotation of short passages for the purposes of criticism and review, no part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher.

Except in the United States of America, this book is sold subject to the condition that it shall not, by way of trade or otherwise, be lent, resold, hired out, or otherwise circulated without the publisher's prior consent in any form of binding or cover other than that in which it is published and without a similar condition including this condition being imposed on the subsequent purchaser.

Library of Congress Cataloging-in-Publication Data

Kenstowicz, Michael J.

Phonology in generative grammar / Michael Kenstowicz.

p. cm.

Includes bibliographical references and index.

ISBN 1-55786-425-X.—ISBN 1-55786-426-8

1. Grammar, Comparative and general—Phonology. 2. Generative grammar. I. Title.

P217.6.K46 1993

92-37749

414—dc20

CIP

British Library Cataloguing in Publication Data

A CIP catalogue record for this book is available from the British Library.

Typeset in 10 on 12 pt Times Roman
by Maryland Composition Co., Inc.
Printed in the United States of America

This book is printed on acid-free paper

Preface

My aim in writing this book has been to provide a thorough introduction to generative phonology. By reading the text and working through the exercises, the student will develop a basic understanding of the fundamental concepts of phonology and the ability to apply them to the analysis of novel data. It is my hope that the reader will acquire the background and tools not only to read the current literature with critical understanding but also to become an active participant in phonological research.

I see the book as being useful to three sorts of readers. First, it can serve as the core text for a one- or two-semester course in phonology at the advanced undergraduate or beginning graduate level. Second, it can be used to provide an overview or to fill in the gaps for instructors who wish to develop a different approach to the subject or to treat a particular topic in depth. Finally, it is suitable for self-study, giving a sense of the kinds of questions phonologists ask and how they go about answering them.

The first three chapters introduce the basic descriptive concepts and analytic techniques of “classical” generative phonology, covering some of the same ground as Kenstowicz and Kisseberth 1979. In conjunction with chapter 4, they are suitable for an undergraduate phonology course. Remaining chapters survey the major lines of generative research that have opened up in the ensuing period. Each motivates the questions being explored, presents the basic results, and then surveys the contemporary scene. These chapters can be read more or less independently. All chapters in the book are accompanied by a list of suggested readings and a set of exercises ranging in difficulty and type. There are more exercises than could reasonably be worked through in a single course. Instructors are urged to pick and choose the ones that work best.

Acknowledgments

This book grew out of the lecture notes for a course on generative phonology I gave in 1987 at the Scuola Normale Superiore in Pisa. I wish to thank my hosts Pier-Marco Bertinetto and Pino Longobardi for their support and interest. Thanks also to Cinzia Avesani, Alessandra Giorgi, Jim Higginbotham, Michele Loporcario, Giovanna Marotta, and Mario Vaira. Much of the data and many of the ideas for the book result from work over the years with students at the University of Illinois and later at MIT. I wish to thank in particular Mohammad Abasheikh, Kamal Abdul-Karim, Issām Abu-Salim, Mohammad Alghazo, Irja Alho, Yousef Bader, Christina Bethin, Rakesh Bhatt, Zoann Branstine, Maria Carreira, Farida Cassimjee, Raung-fu Chung, Abigail Cohn, Laura Downing, Donna Farina, Anne Garber, Ghassan Haddad, Abdel Halim Hamid, Omar Irshied, Omar Ka, Mairo Kidda, Scott Krause, Shlomo Lederman, Michal Livnat, Hassan Marchad, David Odden, Meterwa Ourso, Trudi Patterson, Elizabeth Pearce, Pilar Prieto, Teoh-Boon Seong, Hyang-Sook Sohn, Wafaa Wahba, Uthaiwan Wong-opasi; and Zhiming Bao, Eulàlia Bonet, Lisa Cheng, Chris Collins, Hamida Demirdache, San Duanmu, Alicja Gorecka, Tom Green, Mark Hewitt, Bill Idsardi, Peter Ihionu, Janis Melvold, Rolf Noyer, Brian Sietsema, Kelly Sloan, Tony Bures, Colin Phillips, and Vijayanthi Sarma. Colleagues at Illinois, Venice, and MIT who provided input and support include C. C. Cheng, Guglielmo Cinque, Rodolfo Delmonte, Morris Halle, Jim Harris, Hans Hock, Jay Keyser, Charles Kisseberth, Philip Khoury, Pino Longobardi, Alec Marantz, Wayne O'Neil, Donca Steriade, and Ken Stevens.

For checking data and (dis)confirming various facts I thank Paola Benincà, José Hualde, Greg Iverson, Aditi Lahiri, Utpal Lahiri, Joan Mascaró, Carole Paradis, Jean-François Prunet, Jerzy Rubach, Miklós Törkenczy, Hubert Truckenbrodt, Laura Vanelli, and Draga Zec. Abigail Cohn, Nick Clements, John Harris, and Iggy Roca offered valuable comments and criticism on various chapters. Andrea Calabrese, Bruce Hayes, David Odden, Keren Rice, and Moira Yip read the entire manuscript and provided very penetrating comments that have markedly improved the presentation. I am especially grateful to Morris Halle for help with chapter 4. Also, I wish to thank Anne Mark and Ruth Myott for turning the manuscript into a book and to Philip Carpenter for his persistent interest in the project. Words cannot express the debt I owe to my wife Tamara for her enduring patience and indulgence over the course of this too long project.

Finally, I dedicate the book to the three individuals from whom I have learned the most about phonology: Theodore Lightner, who introduced me to the subject; Charles Kisseberth, fellow graduate student, faculty colleague, and research collaborator; and Morris Halle, colleague and mentor, the founder and leading proponent of phonology in generative grammar.

Contents

Preface	ix
Acknowledgments	xi
Introduction	1
1 The Sounds of Speech	12
1.1 The Saussurean Sign	12
1.2 Speech Sounds: An Internal Dimension	13
1.3 Natural Classes	18
1.4 Distinctive Features	19
1.5 Round and Central Vowels	23
1.6 Features for Consonants	27
1.7 Stricture	34
1.8 Major Class Features	36
1.9 Laryngeal Features	38
1.10 Secondary Articulations and Complex Segments	41
1.11 Prosodic Features	44
Suggested Readings	48
Exercises	49
2 Phonological Rules and Representations	57
2.1 Distinctive versus Redundant Features	57
2.2 Two Levels of Representation	59
2.3 Default Rules and Unmarked Feature Values	61
2.4 Excursus on Unmarked Values	64
2.5 The Phoneme	65
2.6 One or Two Levels of Representation?	69
2.7 Dual-Level Model: Additional Evidence	74
Suggested Readings	81
Exercises	81
3 Phonological Alternations and Derivations	89
3.1 Ordered Rules	90
3.2 Feeding and Bleeding	94
3.3 Intermediate Representations	100
3.4 Multiple Alternation	103

3.5 Linguistic Reconstruction	115
3.6 Algonquian	117
Suggested Readings	125
Exercises	125
4 The Phonetic Foundations of Phonology	136
4.1 Introduction	136
4.2 The Articulators	139
4.3 The Feature Tree	145
4.4 Acoustic Phonetics	168
4.5 Speech Perception	183
Suggested Readings	191
Exercises	192
5 Lexical Phonology	195
5.1 Introduction	195
5.2 Derived Environment Rules	200
5.3 Strict Cyclicity	203
5.4 Morphological Preliminaries	210
5.5 Lexical Phonology	213
5.6 Structure Preservation	221
5.7 Multistratal Rules	224
5.8 Outstanding Problems	227
Suggested Readings	243
Exercises	244
6 The Syllable and Syllabification	250
6.1 Motivations for the Syllable	250
6.2 Syllable Structure and Syllabification: Basic Concepts	252
6.3 The Syllables of English: A Sketch	256
6.4 Medial versus Marginal Clusters	261
6.5 Core Syllables in Ancient Greek and Sanskrit	264
6.6 Continuous Syllabification and Syllable Templates	269
6.7 Continuous versus Staged Syllabification	274
6.8 Resyllabification	280
6.9 Prosodic Licensing and Exhaustive Syllabification	285
6.10 Syllable Quantity and Weight	291
Suggested Readings	298
Exercises	299
7 Autosegmental Phonology	310
7.1 Margi	312
7.2 Shona and the Obligatory Contour Principle	324
7.3 Venda	338

7.4 Vowel Harmony	346
7.5 Makua	359
7.6 Latent Tones in Ga'anda	363
7.7 Parameters of Association	369
7.8 Tone in Asian Languages	376
Suggested Readings	385
Exercises	385

8 The Phonological Skeleton 395

8.1 Arabic Root-and-Pattern Morphology	396
8.2 Some Modern Arabic Dialects	398
8.3 Semitic Biradicals and Multiple Linking	407
8.4 The Integrity of Geminates	410
8.5 Tier Conflation	416
8.6 Assimilation as Spreading	420
8.7 Two Views of the Skeleton	424
8.8 A Comparison	431
8.9 Subsyllabic Constituency	436
Suggested Readings	442
Exercises	442

9 Feature Geometry, Underspecification, and Constraints 451

9.1 Feature Geometry: A Review	452
9.2 Gutturals	456
9.3 Vocalic Place	462
9.4 Vowel Height	475
9.5 Stricture Features	480
9.6 Nasal	490
9.7 Laryngeal Features	493
9.8 Contour Segments	498
9.9 Underspecification	506
9.10 Underspecification and Blocking in Nonderived Contexts	513
9.11 The Coronal Syndrome	516
9.12 Phonetic Underspecification	521
9.13 Rules versus Constraints	524
9.14 Repair Rules and Optimization	528
9.15 The Obligatory Contour Principle and Soft Universals	532
Suggested Readings	537
Exercises	537

10 Stress 548

10.1 Background and Basic Properties	548
10.2 The Metrical Grid	553
10.3 Some Basic Parameters	555
10.4 Heavy Syllables	562

10.5	Unbounded Constituents	565
10.6	Metrical Constituency	566
10.7	Line Conflation	577
10.8	Conflation: Some Extensions	581
10.9	Iambic versus Trochaic Rhythm	587
10.10	Outstanding Issues	593
10.11	Bracket Matching and Edge Parameters	605
	Suggested Readings	613
	Exercises	614
11	Prosodic Morphology	622
<hr/>		
11.1	Reduplication	623
11.2	Prosodic Circumscription	633
11.3	Minimal Word Phenomenon	640
11.4	Extensions and Refinements	646
	Suggested Readings	658
	Exercises	659
	References	663
	Abbreviations	685
	Language Index	686
	Subject Index	693

Introduction

Generative grammar is an approach to linguistics developed at MIT (Massachusetts Institute of Technology) by Noam Chomsky and Morris Halle in the 1950s. It has since become the received theory of phonology and syntax in the United States. Many generative linguists are also active in Canada, Western Europe, Japan, Korea, Australia, and elsewhere. The major goal of generative linguistics has been to solve what Chomsky (1986) has termed *Plato's problem*: any speaker knows many surprising things about the structure of his or her language, things whose internalization is difficult to understand if based solely on evidence from the linguistic environment. Chomsky illustrates this point with the paradigm in (1).


- (1) a. I wonder who [the men expected to see them]
b. [the men expected to see them]
c. John ate an apple
d. John ate
e. John is too stubborn to talk to Bill
f. John is too stubborn to talk to

In (1a) the pronoun *them* may take *the men* as antecedent; but this anaphoric relation is impossible in (1b), even though exactly the same sequence of words is involved. In (1b) *the men* and *them* must be understood to refer to different individuals. These judgments are quite sharp and immediate and will be readily assented to by any speaker of the language. How are they possible? No one has ever taught them to us; in fact, they were only recently discovered. Nevertheless, every mature speaker comes to share these intuitions. A traditional explanation of language learning is that it proceeds by analogy from the most frequent or salient patterns and structures. It is easy to show that analogy is incapable of explaining many such facts. For example, comparing (1c) and (1d), one might infer that if the object of the verb is dropped, the verb is interpreted as taking an arbitrary, unspecified object (John ate something). But this explanation immediately runs afoul of examples such as (1e) and (1f), where the object of the verb *talk to* is also missing. Here object deletion does not result in an unspecified noun phrase interpretation; the missing object in (1f) must refer to *John*.

Chomsky's point is that these judgments are immediate, intuitive, and natural. A child learning English would never generalize from (1c,d) to (1e,f). But "mistakes" of this kind are just what we should expect to find if language acquisition

proceeds by analogy, stimulus generalization, and the other learning mechanisms proposed in classical empiricism.

Over 30 years' study of Plato's problem suggests a different picture. Rather than emphasizing how different from one another languages are, linguists are much more impressed with their basic similarity. For many deep-seated properties such as those illustrated in (1), languages seem to be identical or exhibit only a small range of variation. The knowledge of language revealed in speakers' judgments of coreference in (1) and in many other syntactic structures is quite rich and specific in character. Since this knowledge is so different and remote from overt speech behavior, the most promising and natural explanation for its existence and origin is that humans share a common core of Universal Grammar (UG) as part of their genetic endowment. It is in terms of this common core that language develops in the individual in response to the language of the environment, evolving into Korean, French, English, and so on, as depicted in (2). Languages are so similar or even identical in their underlying structure because they develop from this common core. What systematic differences do exist involve choosing among a constrained set of options (parameter fixing).



Phonology is the component of our linguistic knowledge that is concerned with the physical realization of language. Possession of this knowledge permits us to realize words and the sentences they compose as speech (or as gestures in the language of the deaf) and to recover them from the acoustic signal (or visual sign display). Phonology differs from syntax in that knowledge of the pronunciation of lexical items must be stored in memory; since the relation between sound and meaning is arbitrary, there is no way to predict that 'man's best friend' is *dog* in English, *Hund* in German, *chien* in French, and so on. This basic fact allows for more unpredictable differences among languages that must be learned in the development of an individual grammar from UG. Nevertheless, for a significant range of phenomena, the same general picture depicted in (2) appears accurate. There are many recurrent aspects of phonological structure of a highly specific and rich character whose acquisition cannot be explained on the basis of analogy or stimulus generalization in any useful sense of these terms. These properties are also most naturally explained as reflections of UG.

In this Introduction we discuss a few examples to motivate this general point of view. We then summarize the kinds of questions generative phonologists have been asking and preview the book's survey of the answers that have been offered. In his celebrated (1933) paper "The Psychological Reality of Phonemes," the American linguist Edward Sapir (1884–1939) pointed to the existence of several puzzles posed by the sound structure of language that are as striking as the syntactic ones illustrated in (1). In the course of his fieldwork on American Indian languages, Sapir was struck by the fact that his language consultants would hear

sounds that were not objectively present in their speech. Let us consider one of Sapir's examples in some detail. While studying the Canadian Athabaskan language Sarcee, he was puzzled by his informant John Whitney's insistence that there was a difference between *diní* 'this one' and *diní* 'it makes a sound' even though the two were phonetically homophonous to Sapir's trained ears.

When I asked him what the difference was, he found it difficult to say, and the more often he pronounced the words over to himself the more confused he became as to their phonetic difference. Yet all the time he seemed perfectly sure that there was a difference. At various moments I thought I could catch a slight phonetic difference, for instance, (1) that the *-ní* of "this one" was on a slightly lower tone than the *-ní* of "it makes a sound"; (2) that there was a slight stress on the *dí-* of "this one" . . . and a similarly slight stress on the *-ní* of "it makes a sound"; (3) that the *-ní* of "this one" ended in a pure vowel with little or no breath release, while the *-ní* of "it makes a sound" had a more audible breath release, [and] was properly *-ní'*. These suggestions were considered and halfheartedly accepted at various times by John, but it was easy to see that he was not intuitively convinced. The one tangible suggestion that he himself made was obviously incorrect, namely, that the *-ní* of "it makes a sound" ended in a "t". John claimed that he "felt a t" in the syllable, yet when he tested it over and over to himself, he had to admit that he could neither hear a "t" nor feel his tongue articulating one. We had to give up the problem, and I silently concluded that there simply was no phonetic difference between the words . . . (p. 26; all page references are to Makkai 1972)

Sapir's discussion concerns what he calls a *phonetic illusion*. Two objectively identical stimuli *diní* are perceived as different when they are associated with different meanings. Furthermore, the difference is of a precise form; Whitney felt that the *-ní* of 'it makes a sound' ended in a [t] while *diní* 'this one' did not. What could be the basis of this strange illusion? Clearly, unlike other well-known illusions in which identical stimuli are interpreted differently depending on context (e.g., the Müller-Lyer illusion in which a line looks longer between outward-pointing than between inward-pointing arrowheads: \longleftrightarrow vs. $\rightharpoonup\leftarrow$), the *diní* case depends on having very specialized knowledge – being a speaker of Sarcee.

Sapir goes on to say that as his understanding of the structure of Sarcee increased, the source of the "mysterious t" that his informant intuited became clear. The final syllables of *diní* 'this one' and *diní* 'it makes a sound' behave differently when suffixes such as the inferential *-la* and the relative *-i* are added, as shown in (3). *-la* is unchanged after 'this one' but devoices its initial lateral consonant to [t] after 'it makes a sound'. Before the relative suffix, a [t] appears in 'it makes a sound' while *i + i* contract to a long [a] in 'this one'.

(3)	<u>inferential</u>	<u>relative</u>
	-la	-i
<i>diní</i> 'this one'	<i>diníla</i>	<i>diná:^a</i>
<i>diní</i> 'it makes a sound'	<i>diníṭa</i>	<i>dinít'í</i>

Sapir explains,

There are phonologically distinct types of final vowels in Sarcee: smooth or simple vowels; and vowels with a consonantal latency, i.e., vowels originally followed by a consonant which disappears in the absolute form of the word but which reappears when the word has a suffix beginning with a vowel or which makes its former presence felt in other sandhi phenomena. (p. 26)

In other words, the postulated latent consonant intuited by John Whitney for *diní* 'it makes a sound' actually emerges before the relative *-í* in *dínít'-í*. Its presence is also detectable in the devoicing of the inferential *-la* suffix: *dínít'-la*.

In order to explain his informant's intuition, Sapir postulates that the final vowels of words like *diní* 'it makes a sound' have a "latent" consonant. What this means is that there is another, psychologically more accurate representation of the word that records the presence of this intuited sound: [dinit]. Following Sapir, let us refer to this representation as the *phonological representation*. Sapir concludes,

It is clear that, while John was phonetically amateurish, he was phonologically subtle and accurate. His response amounted to an index of the feeling that *diní* "this one" = *díní*, that *díní* "it makes a sound" = *díní'*, and that this *-ní'* = *-nít'*. (p. 27)

Lest we conclude that such phonological illusions are only to be found in "exotic" languages, Sapir discusses an analogous case from English. He states that

[his informant] John's certainty of difference in the face of objective identity is quite parallel to the feeling that the average Englishman would have that such words as *sawed* and *soared* are not phonetically identical. It is true that both *sawed* and *soared* can be phonetically represented as *sɔ·d*, but the *-ing* forms of the two verbs (*sawing*, *soaring*), phonetically *sɔ·-iŋ* and *sɔ·r-iŋ*, and such sentence sandhi forms as "Saw on, my boy!" and "Soar into the sky!" combine to produce the feeling that the *sɔ·d* of *sawed* = *sɔ·d* but that the *sɔ·d* of *soared* = *sɔ·r-d*. In the one case zero = zero, in the other case zero = r. (p. 27)

Thus, in this particular *r*-less dialect of English, *saw* and *soar* are homophonous as [sɔ:]([:ɔ:] indicates a long back vowel). But the native speaker nevertheless distinguishes them as [sɔ] vs. [sɔr], the [r] in *soar* being supplied in the same way as the [t] that John Whitney heard in Sarcee *díní* 'this one'.

Sapir concludes his discussion of the intuitions of native speakers with respect to the pronunciation of their language with the remark that

Among educated but linguistically untrained people who discuss such matters differences of orthography are always held responsible for these differences of feeling. This is undoubtedly a fallacy, at least for the great mass of people, and puts the cart before the horse. Were English not a written language, the

configuratively determined phonologic difference between such doublets as *sawed* and *soared* would still be “heard,” as a collective illusion, as a true phonetic difference. (p. 27)

In other words, phonetic illusions are found in languages that lack an orthography. Indeed, the phenomenon arose for Sapir in the context of the linguistic description practiced at that time in which large bodies of text (myths, riddles, stories, etc.) were gathered to serve as a data base for linguistic analysis. An efficient way to collect the data was to train the native informants to transcribe them; designing an orthography was thus a prerequisite to linguistic analysis. The phonetic illusions were manifested in the informants’ tendency to spell sounds that were objectively absent in their pronunciation. Since the phenomenon happens in unwritten languages, the orthography cannot be responsible. Rather, the reverse is true: it is precisely because of the existence of such illusions that native speakers feel compelled to represent such latent sounds.

Granted the existence of the phenomenon, the important question then becomes how to explain it. Why are Sarcee speakers compelled to represent *dini* ‘this one’ as [dɪnɪt] and *r*-less English speakers [sɔː] ‘soar’ as [sɔɹ]? How does a child growing up in the Sarcee/English language environments where the relevant sounds are always absent in the pronunciation of these words discover them in the course of acquiring the language? A reasonable first guess is that the phantom [t] in Sarcee and the [ɹ] in English derive from the fact that these sounds actually appear in the related suffixed words *dinit’-i* and [sɔɹɪŋ] *soaring*. But what precisely do we mean by a related word? How is it possible for the pronunciation of one word to influence another? Why doesn’t the influence run in the opposite direction: that is, why doesn’t the absence of a final consonant in the isolation form induce deletion in the suffixed forms? Providing a serious scientific answer to these kinds of questions forms an essential part of the research program of generative phonology.

Let us consider some additional examples of phonetic illusions. English speakers tend to perceive the intersyllabic consonantal material in *camper* and *anchor* as analogous to *clamber* and *anger*. This is an illusion, however. In most dialects (Malécot 1960) the nasal consonant is phonetically absent before such sounds as [p,t,k,s], so that *camper* and *anchor* have the same gross phonetic shape (C)VCVC (V a nasal vowel) as (C)VCVC *wrapper* and *acre*. While VCVC *anchor* belongs with VCVC *acre* phonetically, English speakers have the strong intuition that psychologically it belongs with VCCVC *anger*. Somehow the nasality of the vowel in *anchor* signals the presence of a following latent nasal consonant. But why do we hear the nasal consonant after the vowel instead of before? Why is the sequence VCVC interpreted as VNCVC and not as NVCVC? Notice that in this case, unlike the examples from Sarcee and *r*-less English, there is no suffixed form of the word in which the latent nasal emerges. For these dialects, under normal circumstances *anchor* is always pronounced without a nasal consonant. Nevertheless, the child acquiring these English dialects comes to interpret [æ̃kər] as having the VNCVC shape of *anger* [æ̃ŋgər] and not the VCVC shape of *acre* [ekər]. On what basis is this perception acquired?

Lahiri and Marslen-Wilson (1991) report experimental evidence that corrobo-

rates the contention that English speakers interpret nasality in the vowel as reflecting a following nasal consonant. In a word completion task in which subjects hear stimuli of the form CV and are asked to supply a continuation as rapidly as possible, English-speaking subjects tend to complete C \bar{V} with C \bar{V} N (N = [m,n,ŋ]); their performance differs significantly from that of Bengali subjects, who tend to complete the C \bar{V} stimuli with an oral consonant: for instance, [p,t,k]. This consistent difference between the Bengali- and English-speaking subjects reflects the different phonological status of vowel nasality in the two languages: \bar{V} implies a following nasal in English but not in Bengali, where nasal vowels comprise an important subset of the vocalic inventory. Interestingly, the difference is also reflected in the orthography of the two languages. Bengali has a diacritic mark for nasal vowels while English does not and, if our analysis is correct, could not have an analogous device such as the nasal tilde. Another important difference between the two languages is that in English a nasal vowel is always followed by a nasal consonant unless a consonant from the voiceless [p,t,k,s] series follows. In other words, while English has lexical items such as *camp* with the shape C \bar{V} X (X = [p,t,k,s]), it lacks items with the shapes C \bar{V} or C \bar{V} X (X = [b,l,r,v], etc.). In Bengali no such restriction holds. We will see later why factors of this kind could be relevant to the differing interpretations of the same phonetic stimuli in the two phonological systems.

Phonetic illusions do not just involve the addition of ‘phantom’ sounds. We often misperceive sounds that are objectively present in the speech signal. Consider the well-known pairs *write* vs. *ride* and *writer* vs. *rider* in (4), taken from many American English dialects.

(4)	write	ride	writer	rider
phonological representation	raɪt	raɪd	raɪt-er	raɪd-er
phonetic representation	raɪt	raɪ̯d	raɪ̯D-ər	raɪ̯D-ər

Most speakers of this dialect hear a clear difference between *writer* and *rider* and localize it in the medial dental consonant: *writer* has a medial [t] while *rider* has a medial [d]. The pronunciations of *writer* and *rider* are indeed different. But phonetic study reveals that the medial dental consonants are in fact pronounced identically – as a rhotic, *r*-like consonant called a *flap* (found in Spanish *pero* ‘but’) and symbolized here as [D]. The difference instead resides in the vowel – it is longer in *rider* than in *writer*. Unless one has had phonetic training, however, this vowel length difference is ignored because it is below the threshold of consciousness.

The *writer-rider* pair thus presents us with another phonetic illusion. Our ears must detect the vowel length difference: we do hear these words as different, and this is where the objective difference lies. But somehow our consciousness interprets the contrast as located in the following consonant. Furthermore, it does so in a very precise way – the medial flap of *writer* is perceived as [t] while that of *rider* is perceived as [d]. Once again, how is this possible?

It appears that we must recognize at least two different conceptualizations or *representations* for phonological information: a *phonetic* one indicating how the lexical item is to be realized in speech and an additional *phonological* one that

helps to explicate the “illusions” we have been discussing. These contrasting representations play a significant role in the development of an orthography. Phonetic differences that are overlooked (such as the vowel length in [ra:ɪd]) will typically fail to receive an orthographic registration. Phonetically accurate (but psychologically misleading) orthographies have frequently had to be abandoned as impracticable – precisely because they attempt to record distinctions that are below the level of consciousness of ordinary speakers. Clearly it would be absurd to require English spelling to indicate the vowel length difference in *writer-rider*. But just as clearly, rational orthographies will often require spelling differently what is phonetically the same sound. For example, orthographic reformers would correctly point to the irrational spelling distinction between the “wr” in *writer* and the “r” in *rider* for what is clearly the same sound. But no one would ever propose spelling the medial dental consonants the same in *writer* and *rider* – precisely because they are perceived as different by English speakers.

There is an interesting regularity that supports this interpretation. As just noted, the vowel of *rider* is longer than that of *writer*. This vowel length difference runs throughout English and is strictly correlated with the type of consonant that follows the vowel. Examination of the data in (5) shows that short vowels are associated with such consonants as [p,t,k,s] while long ones are associated with [b,d,g,z].

(5)	tăp	ta:b
	bět	be:d
	büçk	bu:g
	döse	do:ze

The former (voiceless) consonants are articulated with a stiff glottis position while the latter (voiced) consonants have the glottis in a slack (vibrating) position. If *writer* [raɪDər] is represented as [raɪtər] and *rider* [ra:ɪDər] is represented as [raɪdər], then the same regularity observed in (5) explains the vowel length difference in *writer-rider* as well: [t] belongs to the voiceless set and hence calls for a preceding short vowel while [d] belongs to the voiced set and hence calls for a long vowel.

These considerations suggest that English speakers conceptualize the phonological information comprising the lexical units of their language in two different ways. A phonetic representation such as that in (4) indicates how the word is actualized in speech – the instructions sent to the vocal apparatus to articulate the sounds and acoustic properties that are isolated in order to decode the speech signal. The phonological representation is more abstract. It is called into play when speakers have occasion to represent the word in spelling; it may be revealed in language games (e.g., “Say *writer* or *anchor* backward”) and judgments of poetic rhyme. The phonological representation also allows us to explain regularities in the phonetic signal (e.g., the vowel length in *rider-writer*). Phonologists believe that it is essentially the form in which the lexical item is stored in memory. The two representations are systematically related by phonological rules that delete, insert, or change sounds in precise contexts. In the *rider-writer* case, the two rules are first a rule that introduces a vowel length distinction as a function

of the nature of the following consonant ([p,t,k,s, . . .] vs. [b,d,g,z, . . .]), followed by a process that transforms [t] and [d] to the flap [D] in an intervocalic context (V____V).

It is clear that these rules and representations are under the active control of the speakers and not simply a legacy of the historical development of General American English. This point is seen in the ability of many persons to shift from one dialect or accent to another in response to a change in geographic locale or social setting. For example, in his study of English accents, Wells (1982:31–2) says of one of his best Cockney informants,

For telling a joke involving an upper class figure, he can put on quite a flawless U-RP accent. But his everyday speech is rather broad Cockney. Clearly he has within his competence not only the local working class accent but also the local upper class accent. . . . It seems likely that most speakers, perhaps all except younger children, have the ability to “raise” or “lower” the apparent social class characteristics of their speech in this way.

To develop this point further, consider a person who moves from a nonflapping dialect area to a flapping one. Although the process is acquired more or less effortlessly and unconsciously, flapping represents a computational feat of considerable complexity. One must be able to isolate the consonant [t] and then substitute for (better, modify) this speech segment another of a very precise character – the flap [D] – a sound that speakers of the flapping dialect find virtually impossible to pronounce in other contexts such as the beginning of a word. One is not taught to do this – and could not be if it is correct that speakers of the flapping dialect are not even conscious that they make this sound. Furthermore, the [t] is changed only in a particular context whose exact character involves not only the features of the surrounding sounds but also the stress contour: for instance, *átom* has a [D] but in *atómic* the flap is impossible. Finally, the change is extended quite naturally to [d] (as in *rider*) but not to [p] or [k]. For example, in learning the flapping accent, no one would generalize the rule so that *wiper* is changed to *wi[D]er*. Similarly, no child learning this dialect as his or her first language will ever make this false analogy. But why not?

The general answer that has emerged from over 30 years’ study of the problem is that speech sounds are represented in memory as *distinctive features* – linguistic categorizations with precise phonetic correlates. Sound changes are defined over these feature structures. The [t] and [d] (as well as flap [D]) of *writer* and *rider* share the feature of being articulated with the tongue tip and thus form a natural class; the [p] of *wiper* is articulated with the lips. The change of [t] and [d] to the flap [D] proceeds along the same phonological dimension; [p] → [D] cuts across this dimension and thus is an unnatural change. In other words, what is a possible sound change is constrained by the feature structure inherent in the sound’s linguistic representation. Because speech sounds are encoded in memory as feature bundles, the range of modification a sound may undergo in phonetic realization is considerably narrowed – so much so that [p] → [D] will not even be considered as a possible phonological rule.