Lecture Two
The Sounds of Language—Consonants

Scope: The English alphabet gives a very approximate sense of what the actual sounds of the language are. There are 26 letters but 44 sounds. For this reason, linguists transcribe the sounds of languages with a special alphabet that represents sounds as they actually are, the International Phonetic Alphabet, inspired in part by the linguist on whom *Pygmalion*'s Henry Higgins was based. This lecture presents the consonant sounds.

Outline

I. The problem: The alphabet gives the appearance that there are 26 sounds in the English language. In fact, there are quite a few more: English has 44 different sounds. The alphabet only reflects the reality of English sounds approximately.

- **A.** *Cough*, *tough*, *bough*, *through*, and *though* demonstrate five different pronunciations of the sequence of letters *-ough*. In English, the relationship between spelling and pronunciation is inconsistent, sometimes even arbitrary.
- **B.** Because linguists are engaged in analyzing language as it actually is, they transcribe languages according to the actual sounds in them, rather than approximations. English—and all languages—as transcribed in this way can be counterintuitively different from the way we are used to writing.
- **C.** Yet to think as a linguist is to conceive of language as transcribed as it actually is, and this way of "seeing" language will be basic to all of the schools of thought and concepts we will engage in this course.
- **D.** In this lecture, I will introduce you to what consonants look like to a linguist.

II. Linguists have come up with a way of looking at the sounds of language that is more honest to the sounds we make and hear by emphasizing the place and manner of articulation.

- **A.** We will start with *p*. A *p* requires stopping an airflow by putting your lips together. Now notice that *b* is actually a lot like *p*; *b* is,

basically, an *s* with a hum in it, and you may feel as if making a *b* requires more of a push.

B. Now think about *t*. A *t*, like *p*, requires stopping an airflow, only here you stop it by putting your tongue on the ridge behind your upper teeth, the alveolar ridge. Then, just as *b* is a kind of *p*, *d* is a kind of *t*. Thus we have another pair of related consonants, even though the alphabet makes them look unrelated.

C. You can also stop an airflow further back in the mouth. That creates a *k* sound, and the humming/pushy equivalent is a *g*. This time, you stop the airflow with the soft palate, called the *velum*.

D. We now have the beginnings of a chart of how consonants are actually produced in a human mouth. Imagine a cross section of a human head, facing left; *p* and *b* are produced by the lips at the front of the mouth, *t* and *d* in the middle, and *k* and *g* in the back.

E. Now try *m* and notice that you use your lips to produce it just as you do with *p* and *b*. Then you produce *n* on the alveolar ridge just like *t* and *d*.

F. However, when we get to the soft palate, we meet the first sound that English has no written letter for. That sound is the one written as *ng* in *singer*. It may seem that *ng* actually is a sequence of an *n* and a *g*, but it is actually a single sound. We know this because there are words where we actually do produce a sequence of *n* and *g*, such as *finger*. The one in *singer* is a sound in its own right and is called *engma*.

G. At this point we can see that consonants pattern in a way that the sequence of the alphabet obscures. The sounds that halt an airflow are called *stops*. The buzzing ones are called *nasals*. The ones produced with the lips are called *bilabials*, the ones on the alveolar ridge are called *alveolars*, and the ones on the soft palate are called *velars*.

H. Among the stops, the ones that have a hum or a push in them are called *voiced*, while the others are called *voiceless*.

I. Thus there are three features that consonants can have. Where in the mouth they are produced—lips, alveolar ridge, palate—is called the *place of articulation*. How they are produced—such as a stop or a nasal hum—is called the *manner of articulation*. Then there is the *voicing*. Thus linguists refer to *b* as a *voiced bilabial stop*.

III. Linguists use a different kind of alphabet called the International Phonetic Alphabet (IPA). Note: In the IPA, sounds and transcriptions of words and sentences are indicated in brackets, like so: [b], [ð], [æ].

A. The first person to develop an alphabet of this kind was linguist Henry Sweet (1845–1912), who wrote in 1876 that "We must learn to regard language solely as consisting of groups of sounds, independently of the written symbols."

B. In 1886, today's IPA was invented by a French professor of English, Paul Édouard Passy.

IV. Now we can fill in the grid with other consonants.

A. Some of the remaining consonants are produced not by stopping the airflow but by obstructing it, making a hissing sound.

1. Take [s]; notice that you produce it where you produce [t], [d], and [n]. It is an alveolar sound. Then [z] is its voiced counterpart. Hissy sounds are called *fricatives*. [s], then, is a voiceless alveolar fricative.

2. Another hissing sound is [f]. It is produced with the lips, but also with the teeth. This happens somewhat further back in the mouth than bilabially but further forward than the alveolar ridge. [f], then, is a voiceless *labiodental* fricative. Its voiced counterpart is [v].

3. You can also make a fricative by putting your tongue between your teeth. These are the *interdental* fricatives. In English, an interdental fricative is rendered with *th*. However, *th* is pronounced differently in *thin* than in *weather*. That difference is one of voicing. Thus we have two interdental fricative sounds, a voiceless one, [θ], and a voiced one, [ð].

4. In the same way, you can make a fricative between the alveolar ridge and the palate—an *alveopalatial* fricative—which leads to what is usually written as *sh*. However, there are two *sh* sounds—the one in *show* and the one in *pleasure*. Linguists transcribe the first, voiceless, as [ʃ] and the other, voiced, as [ʒ].

B. The sound indicated in English by *ch* is actually a stop followed by a fricative. Linguists indicate it as [tʃ], and its voiced counterpart, the *j* in *judge*, is [dʒ]. These two are called *affricates*.

C. The sounds indicated by *l* and *r* are called *liquids*. Both are produced on the alveolar ridge. In terms of manner of articulation, [l] is called *lateral* and [r] is called *rhotic*.

D. Finally, in English there are two sounds called *glides*. They are the first sounds in *wash* and *yard*. [w] is a bilabial sound, so it is a bilabial glide. What is written as *y* in English is rendered by linguists as [j]. It is an alveolar sound and thus an alveolar glide.

E. [h] is a fricative but is produced further back than the soft palate, in the glottis. It is, therefore, a *glottal* fricative.

F. English has another sound that the alphabet lacks. The first sound in *uh-oh* is not a vowel but a catch in the throat. Thus it is called a *glottal stop*; it is indicated with [ʔ].

G. Notice that there is no symbol for the English letter *x*, which is just a sequence of [k] and [s]. Meanwhile, there are two symbols for consonants the alphabet lacks: the engma, [ŋ], and the above-mentioned glottal stop, [ʔ].

V. There are consonants beyond those in English.

A. Each language fills in only part of this "grid." Other languages have consonants that fill in places in this grid that are empty in English. For example, the *ch* sound in *Bach* is a velar fricative, indicated with [x].

B. Japanese people pronounce *Fuji* in a fashion that sounds, to us, as if they simply "have an accent." However, their sound is not an [f] at all but a bilabial fricative, indicated with [ɸ].

C. The throaty *r* we learn in French classes is, in terms of place of articulation, *uvular*—it is a voiced uvular fricative, indicated with capital *r* written upside down: [ʁ].

Essential Reading:

Pullum and Ladusaw, *Phonetic Symbol Guide*.

Supplemental Reading:

Ladefoged and Maddieson. *Sounds of the World's Languages*.

Questions to Consider:

1. Another consonant that English does not have is a variant of the *r* sound, the trilled version familiar from Spanish words like *perro* ("dog"). In fact, although there is an informal tradition of representing the English *r* sound as [r], in the IPA the English *r* sound is technically represented by [ɹ]. An [r] represents the trilled *r* sound of *perro*. Given that trill sounds are part of the IPA, think of the sound we can make by flapping our lips to express that it's cold ("Brrrr!") or the more raucous sound known as the Bronx cheer. After this lecture, where would you place each of these sounds in our grid? What is the place of articulation, manner of articulation, and voicing? There are, in fact, languages in New Guinea where the Bronx cheer is not an expressive interjection but as normal as [t] and [m]!
2. English spelling has *ng* in the words *hangar*, *linger*, *longer*, *headbanger*, *hunger*, and *ringer*. How would you render what is spelled as *ng* in each of these words? The lesson here is that English spelling is deceptive; The IPA transcribes the actual sounds, which are what linguists study.

Lecture Three
The Other Sounds—Vowels

Scope: The actual vowel sounds in English vastly outnumber the five vowels in the alphabet—*a*, *e*, *i*, *o*, and *u*—and include one that is not distinguished in the alphabet at all. The way the vowels are produced in the mouth also explains just why the way we spell vowels in English is so different from the way they are pronounced. At the end of this lecture, we will be able to transcribe entire words and sentences in the IPA.

Outline

I. In the English alphabet, there are five vowels: *a*, *e*, *i*, *o*, and *u*.

- **A.** In school we learn that there are actually "long" and "short" versions of vowels, e.g., "short *a*" in *cat* and "long *a*" in *father*.
- **B.** The reality of English vowels is, in fact, even richer, as we can see from how vowels are rendered in the IPA.

II. There are five "basic vowels" (especially if you've had Spanish or Italian).

- **A.** Like consonants, the vowels are produced in the mouth in ways that have nothing to do with the order *a*, *e*, *i*, *o*, *u*.
- **B.** Make the sound "ee" and notice that you make it high in the mouth and near the front. Now notice that the sound "oo" is just as high but requires shifting to the back.
  - **1.** Those two sounds are [i] and [u] in the IPA (which will be familiar to those who have studied many other languages). The [u] sound is produced behind the [i] one, and so we will place [i] in the front and [u] in the back.
  - **2.** The sounds "ay" and "oh," on the other hand, are produced lower in the mouth. Their symbols are [e] and [o]. The [e] sound is produced in the front, like [i], while the [o] sound is produced in the back, like [u]. We place them below [i] and [u], respectively. This represents where these sounds are produced in the mouth.

3. Finally, the sound "ah" is produced even lower than [e] and [o], and in the back like [u] and [o]. When we place it on our grid, we have a fundamental representation of vowels in human languages. Most languages have at least these five vowels. Some have just these; most have many more.

III. Three kinds of additional vowels bring us to the sound of English itself.

A. Most of the "basic" vowels have what are called *lax* alternates, as opposed to the *tense* ones we have seen. For example, the vowel sound in *hit* is not [i]—in which case the word would be *heat*—but a different one in the IPA: [ɪ]. It is pronounced somewhat less frontward than [i].

B. In the same way, the vowel in *bed* is not the same one as the one in *made*. It is "eh" rather than "ay," and in IPA it is [ɛ], again less frontward than [e].

C. The lax alternative to [o] is "aw": [ɔ]. It is produced somewhat less toward the back than [o]. In the same way, the lax alternative to [u] is the "uuh" sound in *foot* and *soot* (not *boot* or *coot*). It is [ʊ] in the IPA.

D. A characteristically English vowel is the *a* in *cat*. This is not the "ah" sound in *father*, and it has its own symbol: [æ]. It is produced as low in the mouth as [a] but in the front rather than the back.

IV. Vowels are determined by three parameters: *height*, *frontness*, and *roundedness*.

A. [i], for example, is *high* and *front*. [o] is *mid* and *back*. [æ] is *low* and *front*.

B. All of the back vowels except [a] involve rounding the mouth. Thus we specify whether a vowel is *rounded* or not.

1. [o] is a mid back rounded vowel; [i] is a high front unrounded vowel.
2. Roundedness is important because front vowels can be rounded as well, while back ones can be unrounded even when they are not [a]. The vowel in French *lune* ("moon") or German *Schlüssel* ("key") is a high front rounded vowel, [y] in the IPA.

3. Japanese people pronounce what is written in Roman transcription as *u* as a sound that is transcribed as [ɯ] in the IPA, a kind of "ih" pronounced further back, "ueh," which is a high back unrounded vowel.

V. Vowels in between the front and back vowels are called the central vowels. The last two English vowels occur in the central region.

A. One is the "uh" sound. It is transcribed with a symbol called a *caret*, [ʌ], and pronounced in the mid region.

B. The final sound is one that the alphabet does not distinguish at all: *schwa*. This is the sound of *a* in *about* or *o* in *lemon*. It is very common in English, especially in unaccented syllables, and is written [ə]. It occurs slightly higher than [ʌ] but still in the mid zone.

C. One of the hardest sounds for an English speaker to master in Russian is the vowel ы in words like those for "was" (был) or for "mouse" (мышь), which is a high central vowel, indicated in the IPA with what is called a *barred i*, which is written [ɨ].

VI. Combinations of vowels and glides are called *diphthongs*. In English, they are [aj] (*rice*), [aw] (*house*), and [ɔj] (*boil*).

VII. The Great Vowel Shift.

A. The gulf between spelling and pronunciation in English is due to changes that occurred in the 1400s affecting where vowels are placed in the mouth. *Made* was once pronounced "mah-duh" ([madə]) but the vowel shifted from its low position, [a], to the mid position, [e].

B. At the same time, *feed* was once pronounced "fade" ([fed]), but the vowel [e] moved upward to be pronounced [i]. In the same way, *food* was once pronounced "fode" ([fod]) but the vowel [o] shifted upward to [u].

VIII. How a linguist transcribes a word and how we are used to seeing it spelled can differ considerably.

A. For example, many people use *a* not only before consonants but also before words that in writing begin with a vowel.

B. In fact, such people are using the "proper" rule, because words that we think of as beginning with a vowel are often pronounced with an initial glottal stop—a consonant: [ʔʌʔæpl̩], "a apple."

IX. There are a great many ways that a language can fill up its grid, with English about the middle in terms of the extremes.

A. Rotokas, spoken in Melanesia, is notorious among linguists for having startlingly few sounds, including only six consonants.

B. Polynesian languages such as Hawaiian and Samoan have unusually few sounds—though more than Rotokas.

C. The !Kung [!xũ] language, a click language spoken in the southern part of Africa, has one of the biggest consonant inventories of any language in the world.

Essential Reading:

Pullum and Ladusaw, *Phonetic Symbol Guide*.

Supplemental Reading:

Ladefoged and Maddieson, *Sounds of the World's Languages*.

O'Grady, Archibald, Aronoff, and Rees-Miller, eds. *Contemporary Linguistics*. I have given only a basic characterization of the IPA for the purposes of these lectures. For a fuller command, I suggest the presentation in this textbook.

Questions to Consider:

1. The goal of the past two lectures has been to wean you off of a natural tendency to think of words as composed of "letters," as opposed to sounds. In order to reinforce this, figure out what sentences the following transcriptions represent. Note: We do not pause between words when we speak, and thus neither does the IPA!

[ʃĩgaʔlɪdəl]

[gɛdæwdəðɛr]

[sʌʔmzʌp]

2. There have been many proposals over the years to write English "phonetically" due to our clumsy, arbitrary spelling system. Do you feel that it would be better if English (and maybe all languages) were written in the IPA? Or perhaps do you feel that this would be problematic or inappropriate in some fashion? Opinions differ considerably on this point.

Lecture Four
In the Head versus On the Lips

Scope: A fundamental insight about language's sounds has been that some are "real" sounds that distinguish meaning, like *b* and *p* in *pat* and *bat*, while some are just variations on other sounds, like the *p* that makes a puff of air on your hand in *pot* as opposed to the *p* that does not when you say, "spot." What is just a variation in one language is a "real sound" in others: In Korean, one can utter the *p* in *pul* either with a puff or without, and with the puff *pul* means *grass* while without the puff *pul* means *fire*. Languages differ in which sounds are phonemes ("real sounds") and which are allophones (variations of "real sounds").

Outline

I. You now understand that the actual sounds of a language are something quite different from what is indicated by its alphabet and that the sounds of a language are related to each other in ways that the alphabet gives no hint of, such that [p] and [b] are variations on the same operation (bilabial stop).

II. Armed with this, you are now in a position to understand one of the most basic, frequent, significant differences between the apparent and the actual in language. Namely, in this lecture we will see that we generate words on two levels, according to a basic contrast that linguists have studied over the past hundred years: the *phonemic* versus the *phonetic*.

III. The basic insight is that what comes out on the surface is often quite different from the way the brain originally generates it.

- **A.** In English, the words *bat* and *pat* are obviously different. What makes the difference is that *bat* begins with [b] while *pat* begins with [p].
- **B.** However, in Korean there are no words that begin with [b] nor any that end with it. To a Korean, there is no [b]. The sound does occur, but only as something that happens to [p] in a certain position.

C. That is, [p] turns into [b] when [p] is between two vowels. So, the word [pəp] means *law*. To make that word into *lawlessness*, you add a prefix *mu-*. So the result should be [mu-pəp].

D. But it isn't. Just like in English there is a rule that *a* is used before nouns starting with a consonant and *an* before nouns starting with a vowel (*a pig*, *an apple*), in Korean there is a rule that [p] turns into [b] between vowels. So instead of [mu-pəp], the word for *lawlessness* is [mu-bəp].

E. So [b] is just something that comes from a [p] sometimes. There is no such thing as a word [bəp]. That is, to a Korean, [b] is not a "real sound." The writing system indicates no *b* sound; where a *b* sound occurs like in [mu-bəp], Korean writing shows a *p*.

F. In the same way, say "pot" and then "spot." Notice that when you say the [p] in *pot*, there is a puff of air when you put your hand in front of your face. Then note that there is not that puff of air when you say "spot." The sound pronounced with the puff of air is called *aspirated*.

G. There is a rule in English: Stops occurring at the beginning of a word (or a syllable with the accent on it, like in *capacity*) are aspirated, but not when they are somewhere else in the word. Aspiration is written in the IPA with a superscripted h, and thus [p^h].

H. But we don't think of [p] and [p^h] as different sounds. We think of [p^h] as simply something that happens to [p] under certain conditions. This is because the difference between [p] and [p^h] cannot make the difference between two words' meanings. It's not that [p^hat] means something you cook with while [pat] means, say, to shiver.

I. But that's exactly how it is in Korean. [pul], for example, means *fire*. But [p^hul] means *grass*. The difference between an unaspirated stop and its aspirated equivalent does make the difference between words in Korean. To them, [p^h] is indeed a "real sound" in a way that it isn't to us.

IV. What this means is that there is more to a language's collection of sounds than just the fact that they are all there.

A. In each language, some sounds are "real sounds," making the difference between words' meanings, while other sounds are just

the product of something happening to a real sound depending on where it comes in the word and what kind of sounds are near it.

B. In linguistic terms, "real sounds" are *phonemes*. The other sounds are *allophones* of a phoneme. In English, [b] and [p] are phonemes; [p^h] is merely an allophone of [p]. In Korean, [p] and [p^h] are phonemes, while [b] is just an allophone of [p].

C. In transcription, linguists put phonemes in slashes (/p/) and allophones in brackets ([p^h]).

D. On the phonemic level, then, *spot* is /spat/ and *pot* is /pat/; with the original phonemes, before anything has happened to them based on what they are near or where they are in the word. This is the *underlying form*.

1. On the phonetic level, however, *spot* is [spat] while *pot* is now [p^hat]. Now, the *p* in *pot* is rendered as the aspirated *allophone* of phoneme /p/. This is the word's *surface form*.

2. In the same way, on the phonemic level, Korean for *lawlessness* is /mu-pəp/. On the phonetic level, /p/ is changed into its [b] allophone, and the result is how the word is actually pronounced, [mu-bəp].

E. The distinction between something going on underneath (phonemic) and something that happens on the surface (phonetic) occurs throughout linguistics, to the point that some linguists talk about just *emic* and *etic*.

V. Another example of underlying versus surface is the English plural.

A. Allophones—minor variations on a phoneme's basic nature—are not the only thing that happens in surface forms. Phonemes can also come up on the surface as sounds that themselves are "real sounds" in the language in general.

B. Another example of the difference between underlying and surface forms of words is the plural in English. Because of how we spell, we suppose that we make a plural by adding s. But note that, usually, the sound we actually make is [z]. *Cards*, *pigs*, *fans*, *hams*, *houses*.

C. In fact, the sound we make is [s] only with a small set of sounds: *caps*, *cats*, *hacks*, *coughs*, *moths*. From Lecture Two, you know that these sounds share a feature: They are *voiceless*.

D. This means that there is a rule: The actual *underlying* plural marker *for all words* is a [z] sound. There is a rule that if the [z] sound ends up after a voiceless sound, then it becomes voiceless as well. The voiceless equivalent (alveolar fricative) to [z] is [s].

E. The phonemic form of the plural marker, then, is /z/. *Pots*, underlyingly, begins as /pat/ + /z/. But in the surface form, /z/ becomes an [s] because the *t* in *pot* is voiceless. The result is the surface form [p^{h}ats].

VI. The study of how speakers generate words by working from underlying to surface forms is called *phonology*; the study of how speakers produce sound itself is called *phonetics*.

A. Allophones, although often represented by one letter of the alphabet, are quite central to how we express sounds.

B. For example, because there are many allophones of the phoneme /l/, the word *oil* actually sounds virtually identical when played backward.

Essential Reading:

Sapir, *Language*. Pages 42–56 on language sounds remains one of the most lucid introductions to the phonemic versus phonetic distinction ever written.

Supplemental Reading:

O'Grady, Archibald, Aronoff, and Rees-Miller, eds., *Contemporary Linguistics*. The chapter on phonology in this textbook gives more detail on the differences between words in phonemic transcription versus phonetic transcription.

Questions to Consider:

1. In Japanese, there is something quirky about the [s] sound. [s] never comes before the [i] sound. So, to use words familiar to Americans, there is *sake* wine, the karate *sensei*, the delicious *soba* noodles, and *sushi*. But then, there is the Shinto religion [ʃinto] rather than [sinto], there is sashimi ([saʃimi]) rather than [sasimi], and even when you feel like there "should" be a [si], it is a [ʃi] instead—*cigarette* is *shigaretto*. Based on what you have learned about underlying phonemic form and surface phonetic form, what rule do we see operating in Japanese regarding the [s] sound before the [i] sound? What is the phonemic, underlying form of the word *sushi*?

2. When we say "prove" (/pruv/), the [u] is longer than when we say "proof" (/pruf/). If we said *proof* with the vowel sound as long as in *prove*—"proo-oof"—we would still be saying "proof," just in an odd way. But in the language of people on the island of Yap near Guam, [pul] means *gather*, but if you make the vowel long like the one in English's *prove*, then it is a different word, *moon*. In English, the [u] sound gets long based on a certain rule—look at how the [i] in *dweeb* is longer than in *deep*, or how the [o] in *doze* is longer than the [o] in *dope*. We are dealing with what is termed *vowel length*. In which language, English or Yapese, is vowel length *phonemic*, and in which is it merely a surface phenomenon, *phonetic*?