جامعة نيويورك أبوظبي

NYU | ABU DHABI

PSYCH-UH 2218: Language Science

Class 5: Phonemes and allophones

Prof. Jon Sprouse
Psychology

How does the mind represent speech sounds: Acoustic properties or Articulatory features?

How does the mind represent speech sounds?

One of the driving questions in language science is what the mental representations of language are.

One straightforward possibility would be for the mental representations of segments to be the formant frequencies themselves.

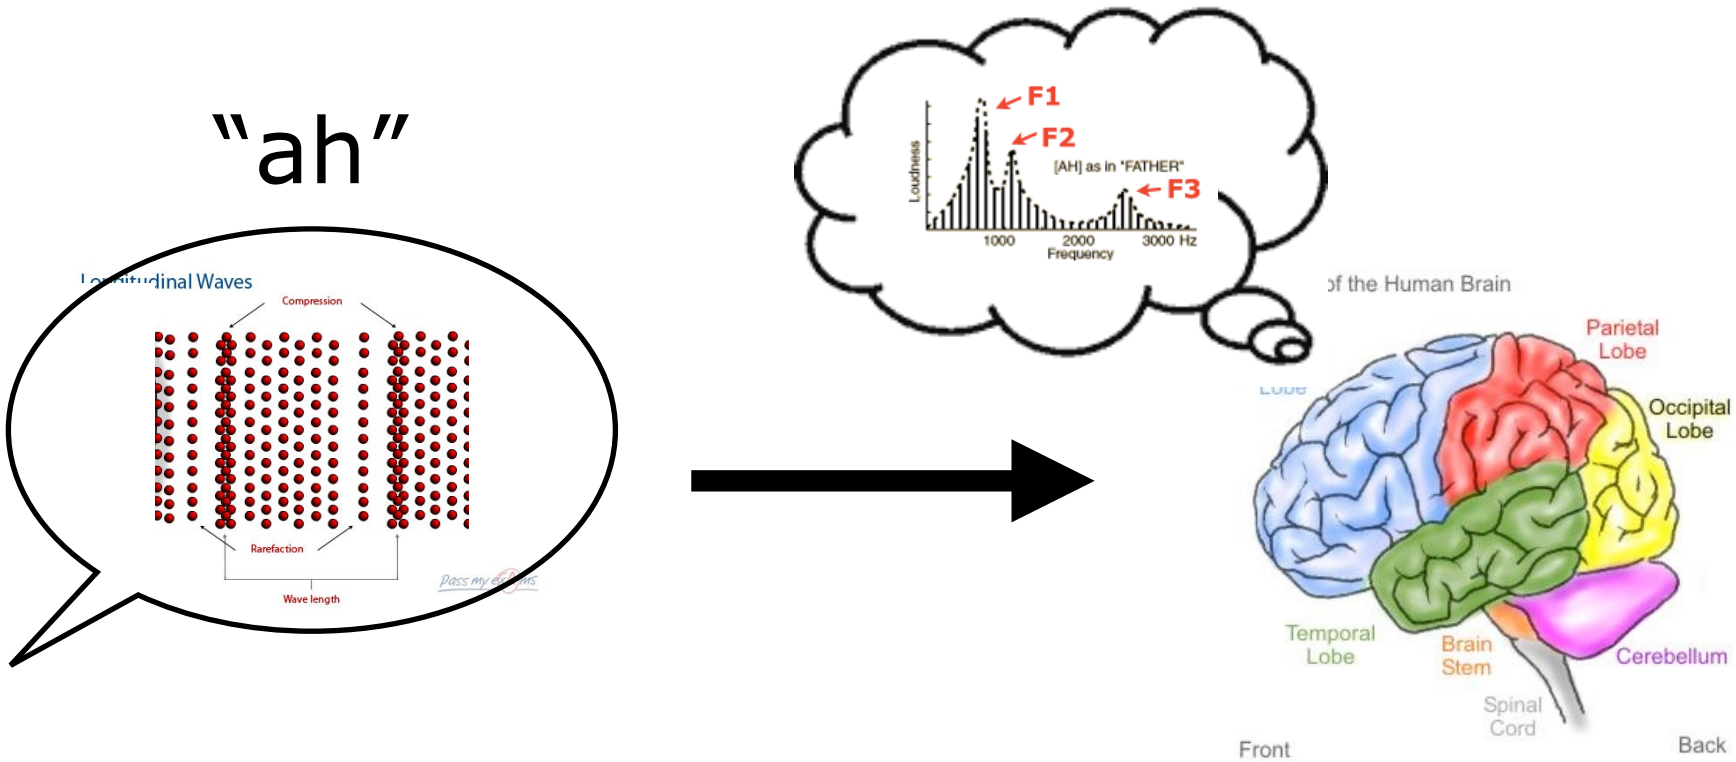

This would make lots of sense… but we have reason to believe the picture is more complicated than this!

Here is one challenge for that idea

To experience this effect, we are going to open and close our eyes throughout the video, to experience the auditory channel in isolation and combined with the video channel.

https://auditoryneuroscience.com/McGurkEffect

We call this the McGurk Effect

This is called the **McGurk Effect**.

The audio track is the syllable **BA.**

The visual track is the syllable **GA.**

The percept is either **DA** or **THA.**

The McGurk effect suggests that we fuse audio and visual information during the perception of speech.

This raises an interesting problem for the idea that segments are represented directly as formant frequencies.

If the mind simply represented segments as formants, why would visual information change the perception of the phoneme? The formants for BA are present in the auditory signal. No other formants are there. How could visual information affect an auditory representation? That would be very surprising.

Visual information is a large component of speech perception

If you haven't seen the McGurk effect before, it may have surprised you. But in reality, visual information is a large component of speech perception.

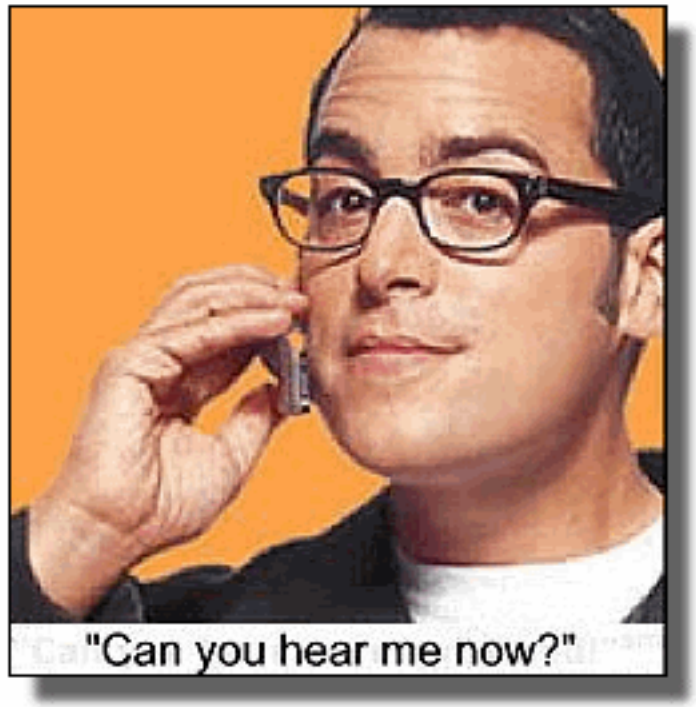

We've all experienced this when talking on the phone: sometimes there isn't enough information in the auditory signal to let us figure out the sound. Or when talking with masks!

This is also happens when older people (with hearing loss) claim that they "can't hear without their glasses" - they use visual (motor) information to compensate for the degraded auditory information due to hearing loss.

We think the representation is something like Articulatory Features

The McGurk effect (and other evidence we will see soon) leads linguists to believe that the representation of speech sounds is not in terms of formants at all but rather something like **articulatory features**..

The McGurk effect is cool, but it is not definitive. It is, in principle, possible to build theories based on formants that could accommodate fusion of auditory and visual information (the theories would just be fairly complex).

The deeper reason that we believe that articulatory features (or distinctive features) are the representation of segments it that they lead to a more explanatory theory of phonology — the theory that explains the patterns that we see in sequences of speech sounds. **And that is what we will start to learn today!**

Some puzzles to motivate our theory

A puzzle when we look at sounds

Here are some words that have the [t] sound in them. Go ahead and say these out loud. Let's really pay attention to how we say them.

top	stop
tick	stick

Do you notice the difference?

[t^h]	[t]

These are aspirated!

These are not.

What about these?

write	writer	written
[t]	[ɾ]	[ʔ]

[ɾ] is called a flap. It is a quick flick of the tongue against the alveolar ridge. It is like a trill - but with only one tap.

[ʔ] is called a glottal stop. It is a stop consonant articulated at the glottis.

Here's the question: Are these [t] or not?

One more [t] because it shows how the relationship between language and writing is complicated

This is a **TEE**
In IPA it is [ti]

But this is a **CHREE**
In IPA it is [tʃri]

This is like the previous examples - the “t” is pronounced differently. But with the others, there is no distinct writing symbol for them. But here, we do have a distinct writing symbol — “ch”. But we do not write it with a “ch”, we write it with a “t”.

This is another example where we have to be careful, and not allow ourselves to be distracted by writing.

But here’s the deep question: Is this a [t] or a [tʃ]?

An important fact - it is systematic

Is this a [tri] or a [tʃri]?

Is this a [treɪn] or a [tʃreɪn]?

Is this a [trʌk] or a [tʃrʌk]?

Is this a [trænzfɔrmər] or a [tʃrænzfɔrmər]?

What is going on here?

The sequence [tr] never occurs in English, we only ever see [tʃr]

Lots of consonants can appear before [r], but for some reason, [t] cannot. There is no word in English that has the sequence [tr] at the beginning of a syllable.

tʃr	sr
br	ʃr
fr	pr
gr	*tr
kr	vr

We use an asterisk to indicate that something never happens.

In fact, if you create novel words and ask English speakers to pronounce them, even if you spell them with a [t], they will pronounce it [tʃ]. **Try it!**

tronk	trimp
trallo	trulip

So this isn't an accident. It looks like English prohibits the sequence [tr], even for new words!

The same systematicity holds for the other [t] variants that we saw

[tʰ] appears at the beginning of words (actually the beginning of syllables).

top	[tʰɑp]
tick	[tʰɪk]
tone	[tʰon]

[ɾ] appears before [ɾ̩] (syllabic r) and [l̩] (syllabic l):

writer	[raɪɾɾ̩]
meter	[miɾɾ̩]
little	[lɪɾl̩]

[ʔ] appears before [n̩] (syllabic n):

written	[rɪʔn̩]
kitten	[kɪʔn̩]
mitten	[mɪʔn̩]

[tʃ] appears before [r]:

tree	[tʃri]

[t] appears everywhere else:

stop	[stɑp]
not	[nɑt]

Explaining the puzzle: A theory with two levels!

Is this a [t] or not? The answer is both!

To capture these facts, we need to postulate two levels of analysis of segments: an underlying representation made up of units that we call phonemes, and a surface representation, made up of units that we call allophones.

allophones: These are sometimes called "variants" of the phoneme. They are the sounds that are actually produced.

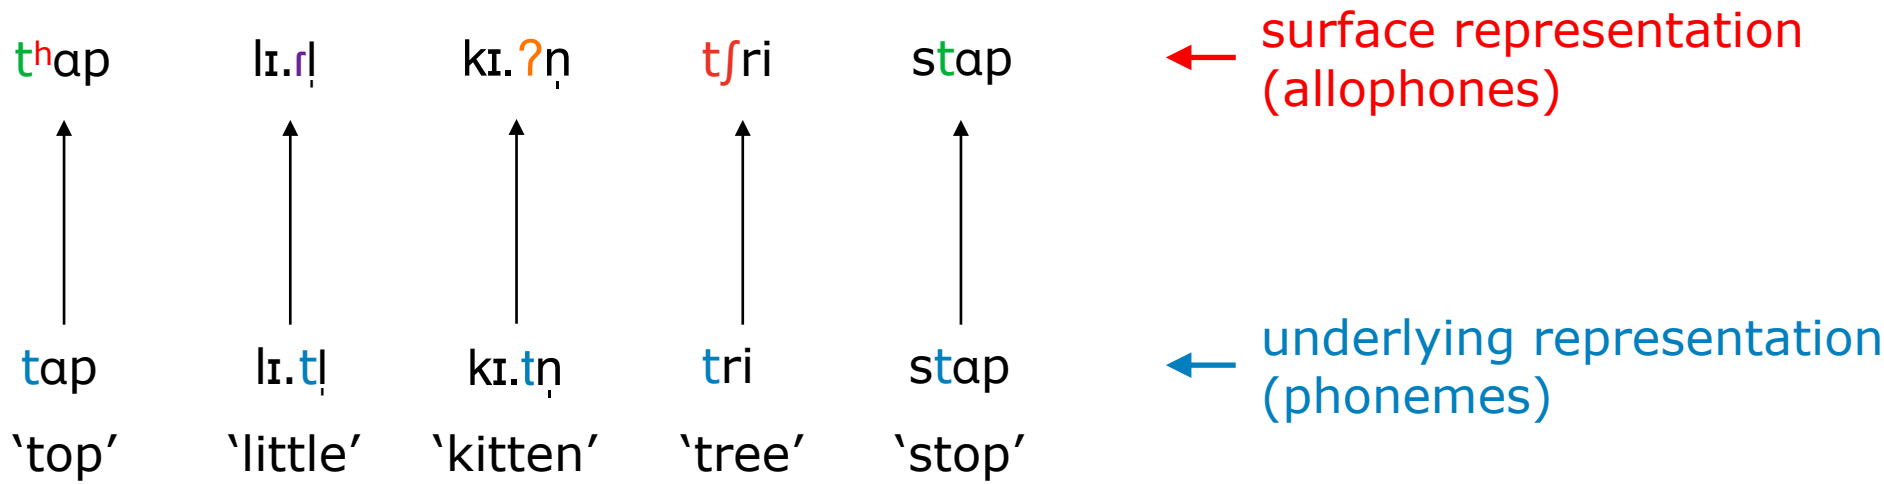

phonemes: The underlying form of sounds. (This is typically what we think of when we think of speech sounds.)

The big idea

The big idea is that there are two levels of analysis for speech sounds - the underlying form (phonemes) and the surface form (allophones). Any given phoneme could have multiple allophones!

allophone: These are sometimes called "variants" of the phoneme. They are the sounds that are actually produced.

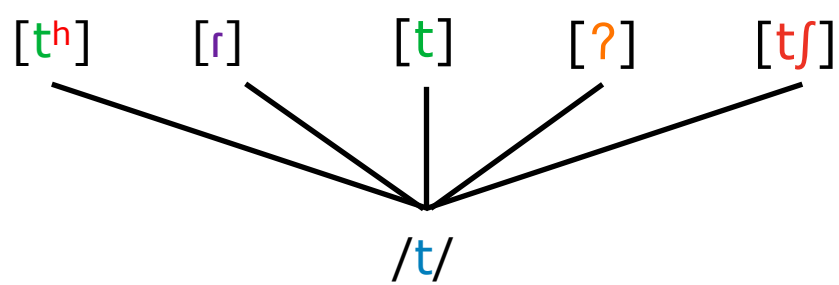

The phoneme /t/ has a quite a large number of allophones in English - at least 5!

phoneme: The underlying form of sounds. (This is typically what we think of when we think of speech sounds.)

Phonemes must have at least one allophone, and that could be the only one they have

Just to be absolutely clear, every phoneme (underlying form) must have at least one allophone (surface form). If it did not have at least one, we would never hear it! Phonemes can have more than one allophone, like /t/, which has 5. But phonemes can also only have one allophone (usually themselves).

allophone: These are sometimes called "variants" of the phoneme. They are the sounds that are actually produced.

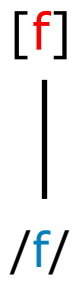

The phoneme /f/ only has one allophone in English - which means the underlying form and the surface form are identical.

phoneme: The underlying form of sounds. (This is typically what we think of when we think of speech sounds.)

It is worth sitting with this for a moment!

What we are saying here is that there are two mental representations for a given word: the phonemic (underlying) representation that is stored in our memory and the allophonic (surface) representation that we produce.

This is a metaphorically crazy idea! It is not obvious at all. The simplest hypothesis would be one representation! But we are led to this conclusion by the facts. That is what science does for us!

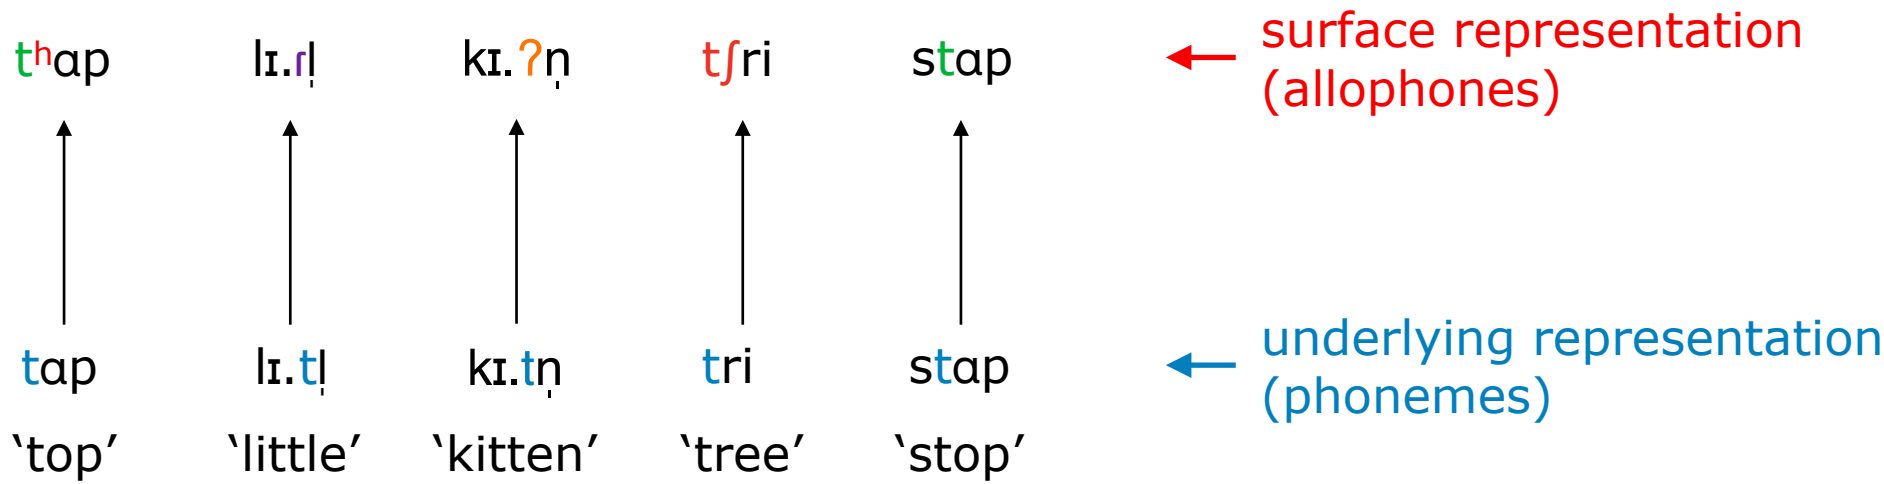

Phonemes and allophones in a language

Because our theory has two levels, we have two goals: to find all of the phonemes in the language (the underlying forms); and to find all of the allophones of each phoneme:

allophones: This is the surface form. This is what we hear in speech. We only ever hear allophones! Every phoneme has at least one allophone - typically itself. It may also have others, like /t/ in English.

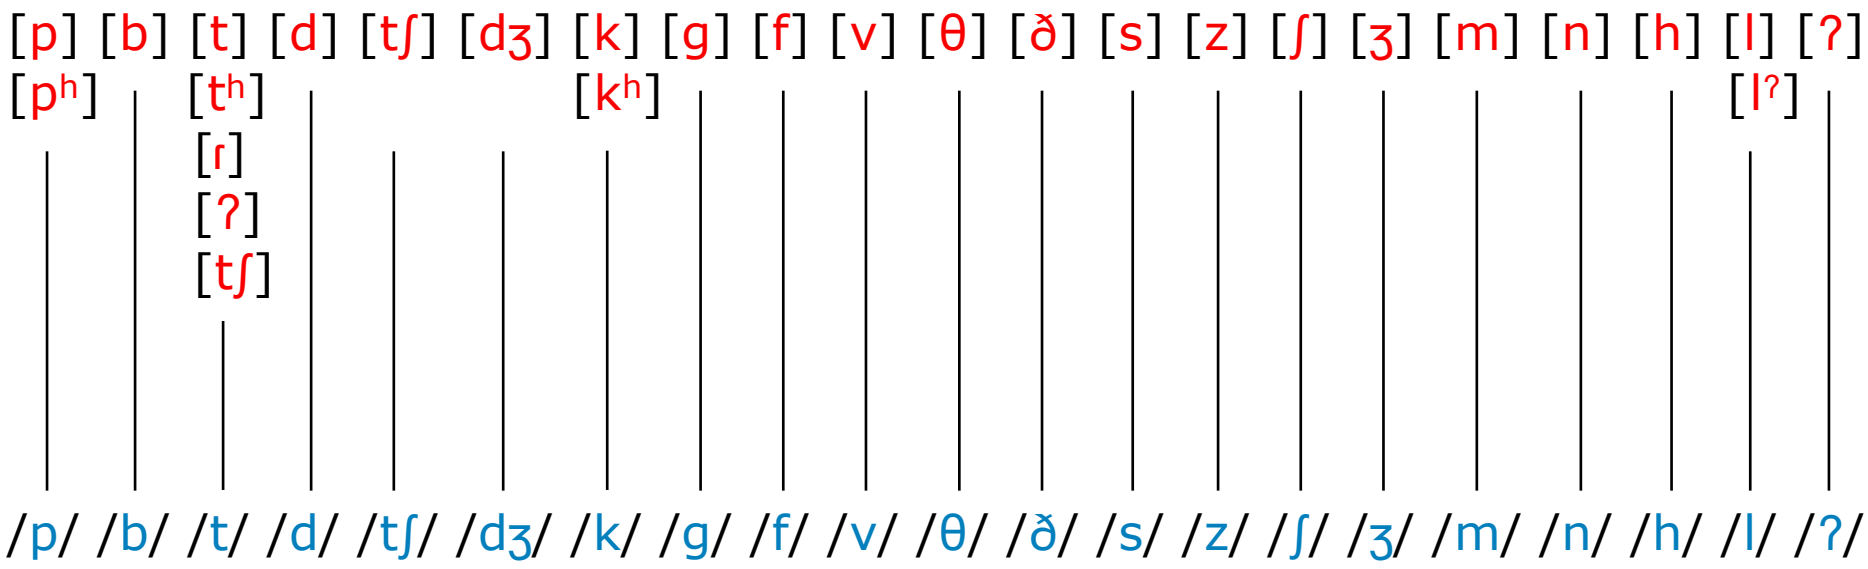

phonemes: This is the underlying form. You don't ever see these in speech. You have to infer them from the pattern of allophones that we see in the language. We will do this now!

A test to identify phonemes and allophones in your native language

The test is simple: minimal pairs

Here is the test we can use to identify phonemes and allophones:

Step 1: Pick a word in the language. s æ t

Step 2: Change one sound in the word. s æ d

Crucially, you have to keep every other sound the same. This makes the two words a **minimal pair** - a pair of stimuli that differ minimally.

And here is the logic of the test:

If the changed sound leads to a <u>different</u> word, then the two sounds (the original and the new one) are allophones of **<u>different</u> phonemes**.

If the changed sound leads to <u>no change</u> in the word, then the two sounds (the original and the new one) are allophones of the **<u>same</u> phoneme**.

Let's see this in action with examples!

Example 1: allophones of different phonemes

Step 1: Pick a word in the language. s æ t

Step 2: Change one segment in the word. s æ d

These are two distinct words in English! So [t] and [d] are allophones of distinct phonemes!

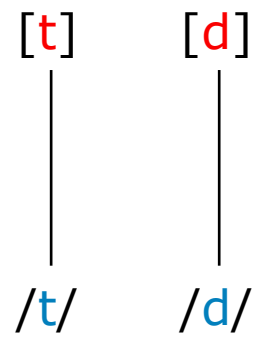

That is it. It is very simple. But notice that we don't know the full story yet. We don't know every allophone for each of these phonemes. And we also don't know what to call the phonemes (I am cheating by already labeling them). We just know that [t] and [d] are each one allophone of a distinct phoneme.

Example 2: allophones of the same phoneme

So let's try the "two distinct words" test with the two allophones of /t/ that we uncovered in the previous two examples. With the logic we have been building, we expect this to fail!

Step 1: Pick a word in the language. t^h ɑ p

Step 2: Change one segment in the word. t ɑ p

These **do not** lead to two different words. Yes, the second one sounds a little funny - this will almost always happen with allophones of the same phoneme (we will see why shortly). But the critical point is that the "funny sound" does not lead to a distinct word. So these are allophones of the same phoneme.

We already suspected that these would be allophones of the same phoneme from our earlier discussion. But this example shows us the rest of our logic. (But notice again that we don't know the full set of allophones, nor do we know the underlying phoneme. I am cheating again by already labeling it /t/.)

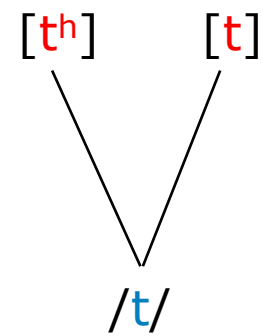

The words we use for this test do not have to be existing words, they can be novel words

One thing to keep in mind is that the existing words in a language are just a sort of accident of history. This means that a sequence of sounds could be missing accidentally. So, when we are doing these tests, we often just want to know whether the sequence of sounds could be a word if someone decided to make one up. So, we can create novel words (called "non-words") to test:

Step 1: Pick a word in the language. [tʰ] æ k

Step 2: Change one segment in the word. [d] æ k

[dæk] is not an actual word in English (as far as I know). But it could be a word. What we mean by that is that native speakers could imagine this being a word. (This contrasts with a word like [dlæk], which cannot be a word in English. Native speakers will say this word is impossible.)

So, from this test we conclude that [tʰ] and [d] are each allophones of distinct phonemes. (And we learn that the test does not need real words, just possible words!)

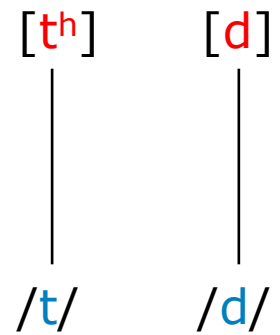

Why does the test work?

Phonemes are contrastive

An important fact you will notice from this test is that phonemes are contrastive with each other: they lead to a difference between words.

sat
sad

Replacing a t with a d changes the word!

pan
ban

Replacing a p with a b changes the word!

This is a deep fact of phonemes. Phonemes are the representations stored in our memory when know a word. That is why we call them the underlying representation.

Since phonemes are the representations of words, they will, by definition, be the way we make distinctions between different words.

The only tricky thing here to remember is that we never see phonemes. We only see allophones. This is because allophones are the surface representation. It is what we produce (and what we hear). We only know the existence of phonemes through our scientific logic! So, every sound we hear in a word when it is produced is an allophone. The question is just which phoneme each allophone comes from!

Allophones of the same phoneme are not contrastive

The other fact you will notice from this test is that allophones of the same phonemes are not contrastive with each other: though they are distinct (both in articulation and acoustics), and though it sounds odd to say them in certain words (we will see this fact next), they do not lead to a difference between words.

tʰɑp	vs	tɑp	Though it is odd to say this without aspiration, these are not different words the way tɑp and bɑp are.
stɑp	vs	stʰɑp	Again, though it is odd to say this with aspiration, these are not different words the way stɑp and swɑp are.

This is a deep fact of allophones. It didn't have to be this way, but it turns out that the surface forms that we produce are not necessarily contrastive the way that the underlying representations are.

A test to identify phonemes and allophones in a language you do not speak!

We look at the phonological environment

The phonological environment that an allophone shows up in is (in most cases) the sound immediately before and the sound immediately after it (so, on either side of it).

We can look at two of the allophones of the phoneme /t/ to see how this works. Here are some words with [tʰ] and with [t]. This is our data set:

word	IPA
tooth	tʰuwθ
tin	tʰɪn
truth	tʰɹuwθ
stool	stuwl
sat	sæt
klutz	klʌts

before	[t]	after
s	_	u
æ	_	#
ʌ	_	s

before	[tʰ]	after
#	_	u
#	_	ɪ
#	_	r

Now, what we can do is look at each relevant allophone in the words, and write down the sound immediately before and immediately after it. I like to do this with tables. We can use an underscore (_) to represent the location of the allophone. Then we look for a pattern!

All allophones but one will show up in what we call a predictable environment

A "predictable environment" just means that the allophone always shows up before or after a certain sound or a certain location in the word.

We can see this with the allophones of /t/. [tʰ] shows up in a predictable environment. It is always at the start of a word. The # symbol means the edge of a word.

word	IPA
tooth	tʰuwθ
tin	tʰɪn
truth	tʰɹuwθ
stool	stuwl
sat	sæt
klutz	klʌts

before	[t]	after
s	_	u
æ	_	#
ʌ	_	s

before	[tʰ]	after
#	_	u
#	_	ɪ
#	_	r

It is easy to see that [tʰ] shows up in a predictable environment because the symbol we wrote in the "before" column is always the same. But can we be more precise about what counts as a predictable environment and what does not? Yes — it involves articulatory features!

We say that sounds that share features form a natural class

Let's look at the IPA chart again. Recall that it is organized by articulatory features (for both consonants and vowels):

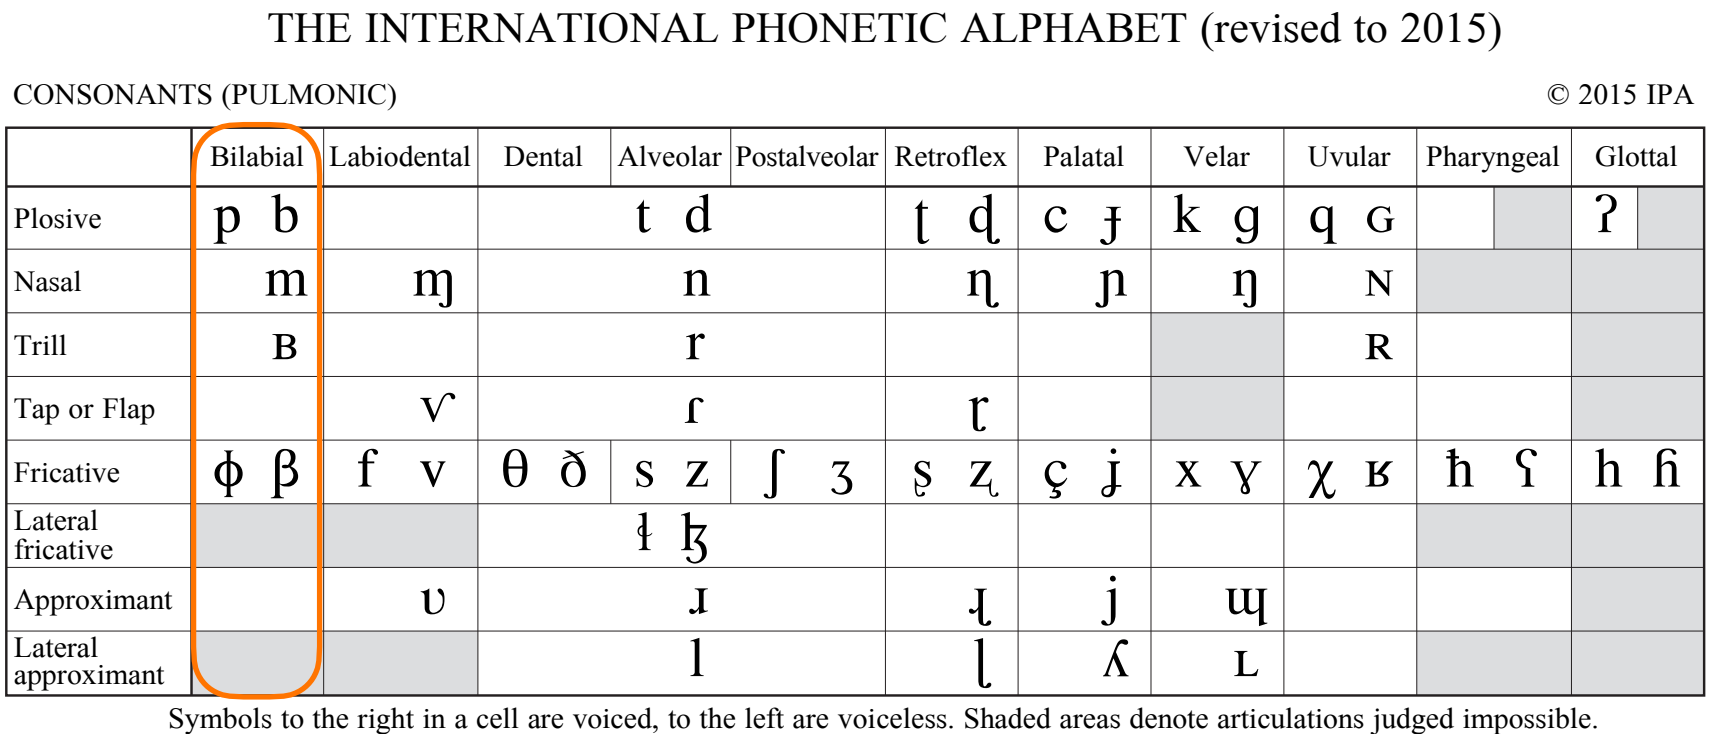

THE INTERNATIONAL PHONETIC ALPHABET (revised to 2015)

CONSONANTS (PULMONIC) © 2015 IPA

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap		ⱱ		ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Symbols to the right in a cell are voiced, to the left are voiceless. Shaded areas denote articulations judged impossible.

All of the sounds that appear in a column share the same place of articulation feature. For example, all the sounds circled above share the feature bilabial. But notice that they have different manner of articulation features (the rows), and different voicing features (left or right in the cell).

We say that sounds that share features form a natural class

Let's look at the IPA chart again. Recall that it is organized by articulatory features (for both consonants and vowels):

THE INTERNATIONAL PHONETIC ALPHABET (revised to 2015)

CONSONANTS (PULMONIC) © 2015 IPA

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap		ⱱ		ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Symbols to the right in a cell are voiced, to the left are voiceless. Shaded areas denote articulations judged impossible.

Similarly, all of the sounds in the same row share a manner of articulation feature, but have different manner of articulation features (the columns) and different voicing features (left or right in the cell). The sounds circled above all share the feature stop (or "plosive").

We say that sounds that share features form a natural class

Let's look at the IPA chart again. Recall that it is organized by articulatory features (for both consonants and vowels):

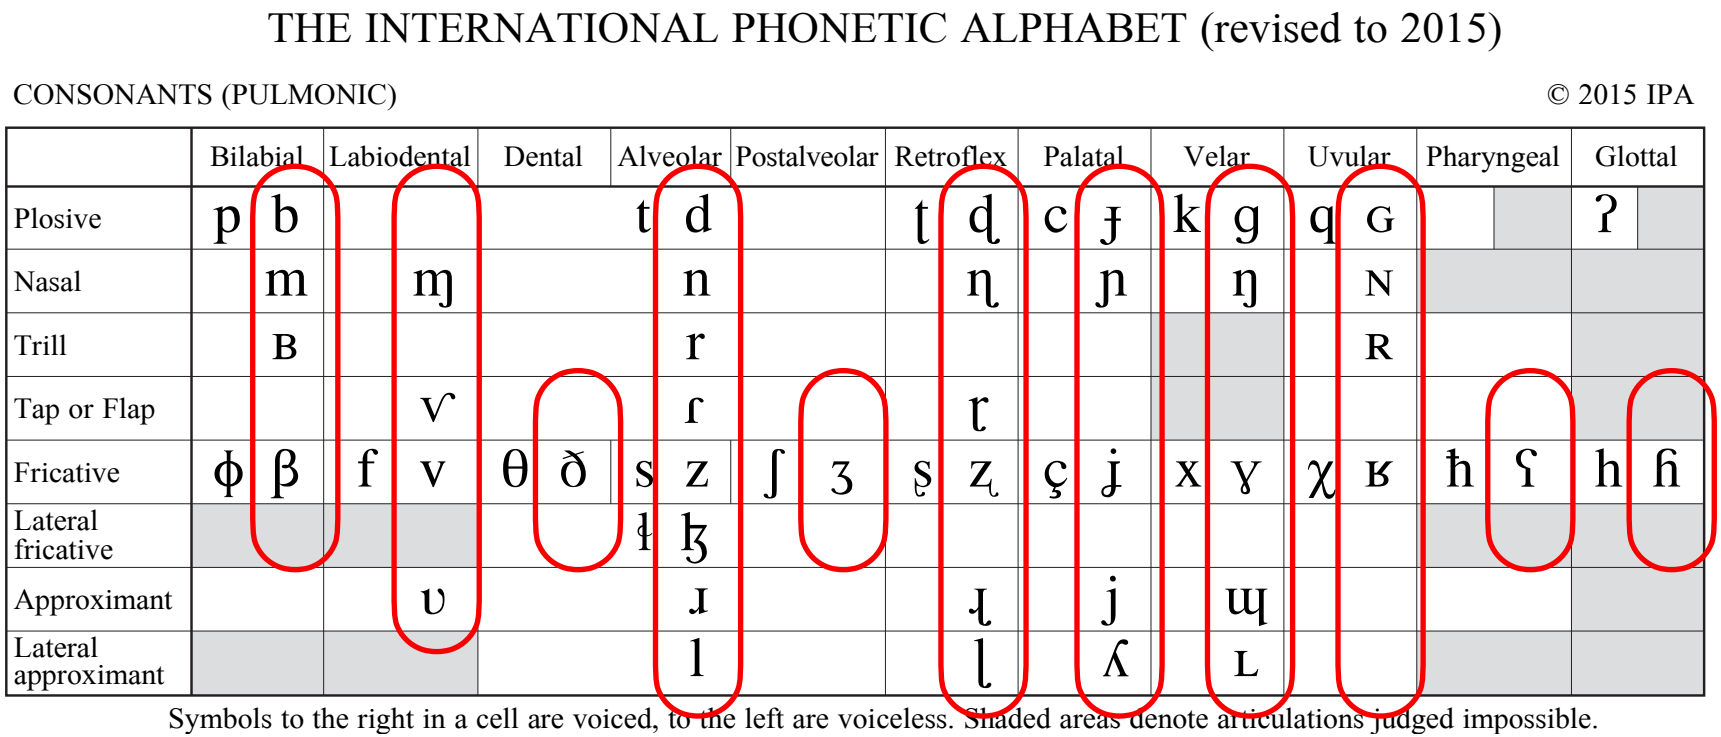

THE INTERNATIONAL PHONETIC ALPHABET (revised to 2015)

CONSONANTS (PULMONIC) © 2015 IPA

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap		ⱱ		ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Symbols to the right in a cell are voiced, to the left are voiceless. Shaded areas denote articulations judged impossible.

The voicing feature is more complicated to see because it is about being to the left or right in the cell. But, all the sounds on the same side of a cell share the feature — either voiceless (left) or voiced (right). Here, I have circled the class of voiced sounds.

A “predictable environment” will form a natural class.

Let’s look at the environments for [tʰ] and with [t] and see if we can identify which are natural classes, and which are not. We do this one column at a time. Let’s start with “before” for [t].

before	[t]	after
s	_	u
æ	_	#
ʌ	_	s

THE INTERNATIONAL PHONETIC ALPHABET (revised to 2015)

CONSONANTS (PULMONIC) © 2015 IPA

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap		ⱱ		ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Symbols to the right in a cell are voiced, to the left are voiceless. Shaded areas denote articulations judged impossible.

VOWELS

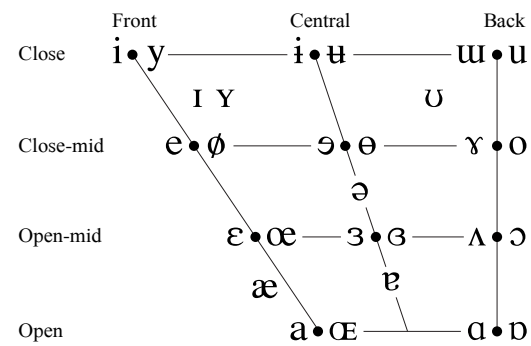

Where symbols appear in pairs, the one to the right represents a rounded vowel.

Do these three sounds share a common feature?

No! This is easy to see because [s] is a consonant and [æ] and [ʌ] are vowels. Consonants and vowels don’t share any features, so it cannot be a natural class.

A “predictable environment” will form a natural class.

Next let’s look at the “after” column for [t].

before	[t]	after
s	_	u
æ	_	#
ʌ	_	s

THE INTERNATIONAL PHONETIC ALPHABET (revised to 2015)

CONSONANTS (PULMONIC) © 2015 IPA

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap		ⱱ		ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Symbols to the right in a cell are voiced, to the left are voiceless. Shaded areas denote articulations judged impossible.

VOWELS

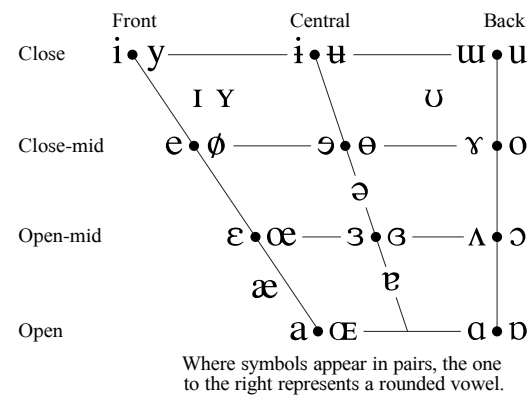

Where symbols appear in pairs, the one to the right represents a rounded vowel.

Do these three sounds share a common feature?

No! This is even easier — # means edge of the word. It has no features. So it can never form a natural class with sounds. (Also, [u] is a vowel and [s] is a consonant.)

A “predictable environment” will form a natural class.

Next let’s look at the “before” column for [tʰ].

before	[tʰ]	after
#	_	u
#	_	ɪ
#	_	r

THE INTERNATIONAL PHONETIC ALPHABET (revised to 2015)

CONSONANTS (PULMONIC) © 2015 IPA

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap		ⱱ		ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Symbols to the right in a cell are voiced, to the left are voiceless. Shaded areas denote articulations judged impossible.

VOWELS

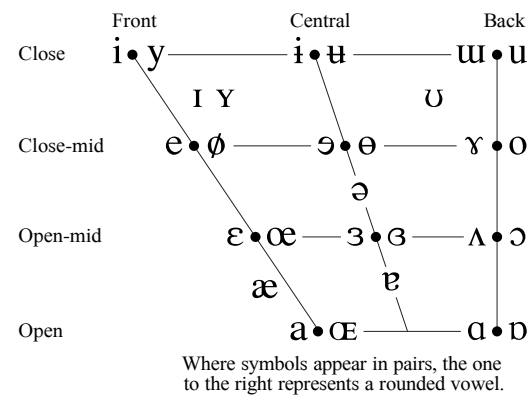

Where symbols appear in pairs, the one to the right represents a rounded vowel.

Do these three share a common feature?

Yes! This is easy because they are all identical, so we know they will form a natural class. In this case, there are no features, but they still form a class (the class of word boundaries).

A “predictable environment” will form a natural class.

Next let’s look at the “after” column for [tʰ].

before	[tʰ]	after
#	_	u
#	_	ɪ
#	_	r

THE INTERNATIONAL PHONETIC ALPHABET (revised to 2015)

CONSONANTS (PULMONIC) © 2015 IPA

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap		ⱱ		ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Symbols to the right in a cell are voiced, to the left are voiceless. Shaded areas denote articulations judged impossible.

VOWELS

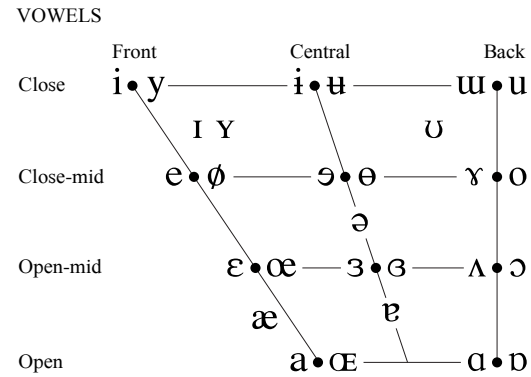

Where symbols appear in pairs, the one to the right represents a rounded vowel.

Do these three share a common feature?

No! [r] is a consonant and [u] and [ɪ] are vowels.
So they share no features.

A more complicated example for identifying natural classes

Another example: non-nasal and nasal vowels

Here is a set of words in English. The tilde over a vowel means it has the feature [nasal]. [i] and [ĩ] are allophones of the same phoneme /i/. They are not contrastive in English.

word	IPA
dean	dĩn
lean	lĩn
mean	mĩn
team	tĩm
seam	sĩm
deed	did
leap	lip
mere	miɹ
tea	ti
seek	sik

[i]

before		after
d	_	d
l	_	p
m	_	ɹ
t	_	#
s	_	k

[ĩ]

before		after
d	_	n
l	_	n
m	_	n
t	_	m
s	_	m

Step 1 in analyzing these is to list all of the environments for each allophone. I have already done that here. But be sure you can see how each row in these tables maps back to a word in the list!

Another example: non-nasal and nasal vowels

Here is a set of words in English. The tilde over a vowel means it has the feature [nasal]. [i] and [ĩ] are allophones of the same phoneme /i/. They are not contrastive in English.

word	IPA
dean	dĩn
lean	lĩn
mean	mĩn
team	tĩm
seam	sĩm
deed	did
leap	lip
mere	miɹ
tea	ti
seek	sik

[i]

before		after
d	_	d
l	_	p
m	_	ɹ
t	_	#
s	_	k

[ĩ]

before		after
d	_	n
l	_	n
m	_	n
t	_	m
s	_	m

Step 2 is to look at each column in the table, and see if the sounds in the column form a natural class based on articulatory features.

Another example: non-nasal and nasal vowels

Let's look at the "before" environment for [i]. I have circled each of the sounds in the consonant chart.

before	i	after
d	_	d
l	_	p
m	_	ɹ
t	_	#
s	_	k

THE INTERNATIONAL PHONETIC ALPHABET (revised to 2015)

CONSONANTS (PULMONIC) © 2015 IPA

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap		ⱱ		ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Symbols to the right in a cell are voiced, to the left are voiceless. Shaded areas denote articulations judged impossible.

VOWELS

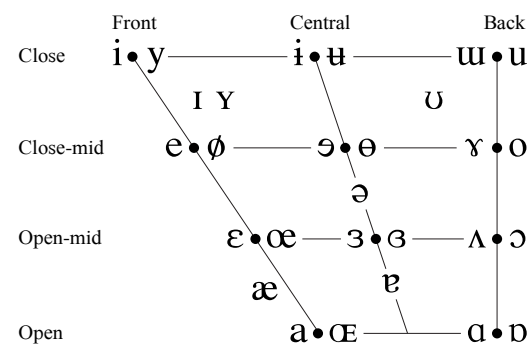

Where symbols appear in pairs, the one to the right represents a rounded vowel.

Do these three sounds share a common feature?

No! They show up on different rows, so they don't share a manner feature. They show up in different columns (m is the outlier!), so they don't share a place feature. And they show up on both sides of the cell, so they don't share a voicing feature.

Another example: non-nasal and nasal vowels

Now let's look at the "after" environment for [i]. I have circled each of the sounds in the consonant chart.

before	i	after
d	_	d
l	_	p
m	_	ɹ
t	_	#
s	_	k

THE INTERNATIONAL PHONETIC ALPHABET (revised to 2015)

CONSONANTS (PULMONIC) © 2015 IPA

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap		ⱱ		ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Symbols to the right in a cell are voiced, to the left are voiceless. Shaded areas denote articulations judged impossible.

VOWELS

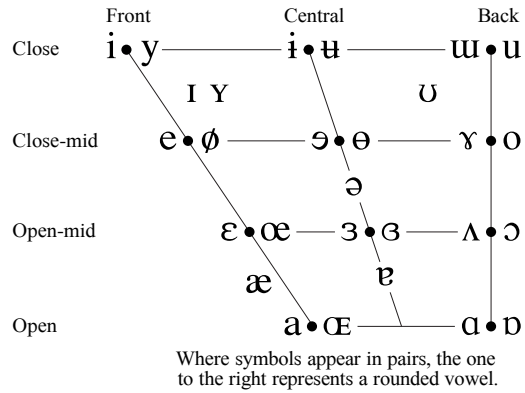

Where symbols appear in pairs, the one to the right represents a rounded vowel.

Do these three sounds share a common feature?

No! This is easy because # never forms a natural class with sounds.

Another example: non-nasal and nasal vowels

Now let's look at the "before" environment for [ĩ]. I have circled each of the sounds in the consonant chart.

before	ĩ	after
d	_	n
l	_	n
m	_	n
t	_	m
s	_	m

THE INTERNATIONAL PHONETIC ALPHABET (revised to 2015)

CONSONANTS (PULMONIC) © 2015 IPA

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap		ⱱ		ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Symbols to the right in a cell are voiced, to the left are voiceless. Shaded areas denote articulations judged impossible.

VOWELS

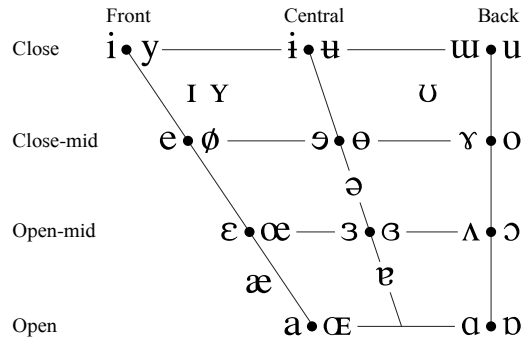

Where symbols appear in pairs, the one to the right represents a rounded vowel.

Do these three sounds share a common feature?

No! Also notice that these are exactly the same sounds that we saw for the "before" environment for [i]. This is another shortcut. When two allophones have an environment that overlaps (shows the same sounds), we know that is not the critical environment.

Another example: non-nasal and nasal vowels

Now let's look at the "after" environment for [ĩ]. I have circled each of the sounds in the consonant chart.

before	ĩ	after
d	_	n
l	_	n
m	_	n
t	_	m
s	_	m

THE INTERNATIONAL PHONETIC ALPHABET (revised to 2015)

CONSONANTS (PULMONIC) © 2015 IPA

	Bilabial	Labiodental	Dental	Alveolar	Postalveolar	Retroflex	Palatal	Velar	Uvular	Pharyngeal	Glottal
Plosive	p b			t d		ʈ ɖ	c ɟ	k g	q ɢ		ʔ
Nasal	m	ɱ		n		ɳ	ɲ	ŋ	ɴ		
Trill	ʙ			r					ʀ		
Tap or Flap		ⱱ		ɾ		ɽ					
Fricative	ɸ β	f v	θ ð	s z	ʃ ʒ	ʂ ʐ	ç ʝ	x ɣ	χ ʁ	ħ ʕ	h ɦ
Lateral fricative				ɬ ɮ							
Approximant		ʋ		ɹ		ɻ	j	ɰ			
Lateral approximant				l		ɭ	ʎ	ʟ			

Symbols to the right in a cell are voiced, to the left are voiceless. Shaded areas denote articulations judged impossible.

VOWELS

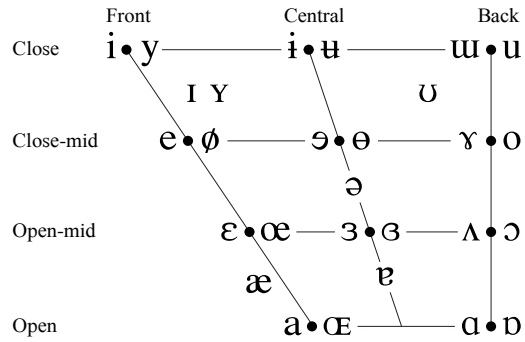

Do these three sounds share a common feature?

Yes! They are both on the same row, which means that they share a manner feature. In this case, it is [nasal].

Another example: non-nasal and nasal vowels

This is interesting. We can now say that [ĩ] is the allophone that appears when a nasal consonant comes after it, and [i] is the allophone that appears when any other sound (or word boundary) appears after it.

word	IPA
dean	dĩn
lean	lĩn
mean	mĩn
team	tĩm
seam	sĩm
deed	did
leap	lip
mere	miɹ
tea	ti
seek	sik

[i]

before		after
d	_	d
l	_	p
m	_	ɹ
t	_	#
s	_	k

[ĩ]

before		after
d	_	n
l	_	n
m	_	n
t	_	m
s	_	m

This makes some sense. [ĩ] is a nasal vowel and it appears next to nasal consonants! This is our first hint that there is something deeper happening! There is a logic dictating the choice between allophones. We want to uncover that logic!