

MITOCW | 16. Numbers and math

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

All right, so this week we are going to talk about language, language being one of the really cool things that humans do that we haven't gotten to talk about yet language is different from everything else you've been talking about because it's pretty much uniquely human lots of animals communicate in one way or another. But they are not using language per se as a linguist would define it so let's start out by talking about some ways in which-- thank you.

OK.

Language is different from-- I'll snag Helen's, too. Language is different from animal communication. Can anyone think of any ways in which humans and animals communicate in this, really fundamentally differently? Think about it. I don't know how much you guys might know about animal communication. You've got everything from vervet monkeys who have specialized warning calls for different predators, versus humpback whales and their unique identifying songs, versus honey bees and dances and stuff.

Some of the ways in which human language is different from other kinds of communication-- one is that human language can talk about things that aren't here right now. You told me that Jen is in Harvard, trying to catch the bus to get over here.

[INAUDIBLE]

Hynes, trying to catch the bus to get over here. Not here, talking about a place that isn't here. So we're not just talking about things that we can actively perceive right now-- we're talking about stuff that's abstract or distant or temporally separated-- in time.

[INAUDIBLE] talk about it.

Good. We're going to have six things-- six, five, some number-- of things that languages are, and there aren't any forms of animal communication that meet all of these things. But you're right-- honeybees are talking about directions to a place that's not where they are right now. That's a really good point. Another thing about language is that language is arbitrary. What that means is that in English, you have a specific sound to represent something-- we've got a sound for chair, and a set of sounds for window, and a set of sounds for ceiling. And the nature of these sounds doesn't have any connection to the objects themselves-- they're just an arbitrary phoneme string that we've assigned a meaning to. Hey, Natasha.

Hi.

How are you? Feeling like it's Monday? We are talking about languages, and things that languages are that other forms of communication are not. What makes language different from other kinds of communications. Languages are arbitrary-- symbols and languages aren't necessarily related to the things they're representing. And again, there are some counterexamples to this-- every language has its share of onomatopoeia words. Words-- for example, the noises animals make-- which are related to what they're representing. But for the most part, the symbols involved in languages are arbitrary.

This is always someone's cue to say, but what about sign language? Sign language from people who aren't really familiar with it tend to assume that it's just gestures. It's not-- I mean, often if you're looking at a sign you can figure out kind of a gestural basis for it. For example, the ASL for give is-- like this, kind of like you're holding a hand out to somebody. On the other hand, the ASL for water looks like-- looks like that, which is a lot less directly derivative from some kind of gesture.

A lot of ASL has little mnemonics, where you can think about what might be a gestural basis for it, but it's a stretch-- so languages are arbitrary. Languages are generative-- people come up with new pieces for them, and add them in. I bet everybody in this room can think of at least one word that wasn't in the English language when you were born.

Google.

One.

[INAUDIBLE] What?

Google wasn't not in the language in 1990?

It wasn't the word for [INAUDIBLE] some number?

It was, but not spelled like that. Googol the number's spelled differently. No, I mean, Google the search engine is from--

It's a number?

Googol-- G-O-O-G-O-L.

10 to the 100.

Yes-- the search engine is named after the number. Yay!

Six! We're getting there.

[INAUDIBLE]

[SIDE CONVERSATION]

The company is named after the number, and the search engine [INAUDIBLE]

[INAUDIBLE]

Yeah. Well, more or less-- the company is the search engine-- the company, when it started, was the search engine. All of the other things Google does are more recent. What's a word that wasn't in the language when you were born, Zachariah?

[INAUDIBLE]

Sure, or as a noun for that matter. Hybrid, in the sense of car, for example-- probably wasn't. Wikipedia.

Facebook.

Facebook.

[INAUDIBLE] part of the language, though [INAUDIBLE]

I think once you start using them as a verb, they're definitely being expanded beyond just their name use. Like, you can definitely say, "I Facebook so and so," can't you?

[INTERPOSING VOICES]

Can you? I don't know-- I may not be hip and young enough to make these kinds of distinctions.

[INAUDIBLE]

Languages are generative-- they produce new words, they produce new uses of words-- so you'll see words that change their meaning, change their use. You'll see people who believe that language should not change get very unhappy about this from time to time. Who here has ever said, oh well, hopefully it won't rain-- we want to go to the Red Sox game tonight, or something along those lines. Hopefully in the sense of "I hope that this will not happen," or "that this will happen." Right. So some people will say, that's not what that means!

Hopefully originally meant "in a hopeful manner," like he looked at the sky hopefully or something. It has been

established in this other sense that everybody uses it in for a very long time-- languages change, languages are generative, languages add new pieces. This is one way in which languages and honeybee dances are different, is that honeybee dances seem to be set, they don't seem to shift over time, change over time in the way that you'd see a human language doing. Yeah.

Would the establishment of grammar predict changes in [INAUDIBLE]. Because people will be corrected on what they think and say, [INAUDIBLE]

I don't know. Listen to your peer group, listen to your parents' peer groups-- do you guys talk the same?

Because English isn't my first language.

All right, so there's one reason. So to some extent, yes.

I can try to answer that question [INAUDIBLE]

Some things will slow it down a bit. One thing that's really happened is that spelling changes slowed down when English started having fixed spelling about 200 years ago, when people started really compiling dictionaries and having this idea that there was a right and a wrong way to spell things, and I think from the 1700s earlier, you'll see a whole assortment of spellings first off-- it's just whatever seemed good to that author at that time.

For grammatical structure, seems to change slowly and steadily no matter what anybody does about it-- but vocabulary definitely changes really fast. We all see this-- words that weren't in the language when I was born, for example. Cell phone was not in the language when I was born. Might be for you guys, early '90s I don't know.

[INAUDIBLE] Those big bricks. The big--

Cell phones were [INAUDIBLE]

Cell?

[INAUDIBLE]

Yup. All right, so we've got-- language can talk about things that are not immediately present, languages are arbitrary, languages are generative, languages are culturally transmitted. Language is not-- or a particular language is not something that you are genetically born with. It's something you've got to learn from being exposed to language, to other people who are speaking a language. You learn it, you learn it from your peers, you learn it from your family, you learn it from all of the language that you are exposed to as a smallish kid.

There's a more or less-- like with many things, there's a critical period for being exposed to language. I don't know if any of you guys have heard of the case of Genie who was a girl in the '70s who was found at the age of 11, and had basically lived her entire life locked in a closet. Her parents didn't interact with her, didn't talk to her, locked her in a closet, gave her food and water periodically, and was--

Why?

Because they were nuts-- like, completely batshit insane crazy. It's pretty much the only answer I have come up with for that. So as you might imagine, social services eventually found out about this situation, took the kid away, moved the kid into foster care. And what they found with Genie is that she eventually developed some vocabulary-- like, about 150 word vocabulary. To put that in perspective, most adults have something like a 40,000 word vocabulary. The other thing is that she developed vocabulary-- she didn't really develop the ability to use sentence structure in any kind of useful way.

Sentence structure lets us build all of these complicated, elaborate strings that we can all parse just fine, you know. This is the red chair in the corner of the classroom that we have Junction in, that I am in four days a week, as are most of you. One, perhaps not terribly well crafted, sentence-- but it's got a lot of subpieces that you could pick out and parse just fine. Genie can't do that. There's a critical period for language development-- this is why little kids can learn second, third, fourth languages just fine, and grownups have a really hard time with it.

The things that is often frustrating to people who study language is that most schools will start a second language in middle school, right? Usually it's where you'll start taking French or Spanish or Latin, or what have you-- right about when this critical language development period is closing, and it starts being really hard for you to learn new languages. Languages are culturally transmitted-- you get them from the people around you.

All right, slightly more complicated-- languages are what's called dual patterned. What this means is that you get-- the English language has a fixed number of sound units that make it up, of phonemes. If we talk about the word chair, for example, it's spelled C-H-A-I-R, but that's not terribly useful. There's three major sound units that go into it-- there's the "ch" at the beginning, and in "a" the middle and the "r" at the end-- chair. Each of those individual phonemes, those sound units, is not meaningful-- the sound unit alone is not a unit that has meaning.

There are a few single phoneme words in English, like "oh" and "eye" and "a," but there aren't many.

How about "ow?"

Ow.

Wouldn't those technically be made up of multiple phonemes [INAUDIBLE] diphthongs?

The one in the middle of "chair" is kind of a diphthong, yeah. Well, "oh" isn't-- "oh" is a straight up tense, fronted, rounded vowel-- no, backrounded vowel, "oh." "Ow" is, but a lot of these aren't. Single phonemes can be morphemes, units with meaning-- often, but usually aren't. Languages are dual patterned-- so the units that have meaning are made up of smaller units that don't have meaning, that can then be jiggled together and broken around in different ways to form these morphemes.

A morpheme is the smallest unit that carries meanings. Any word is a morpheme-- any base root word-- but also "d" you might add to something if you're saying, OK, I wanted. You've got two morphemes there-- the "want" part and the "ed" part.

How do you spell morpheme?

M-O-R-P-H-E-M-E.

E-N-E?

E-M-E-- morpheme.

[INAUDIBLE] I thought you said morphine.

[INTERPOSING VOICES]

Morpheme. Phoneme and morpheme. These are different units that linguists use to talk about what makes up a language. All right, six things languages are that other forms of communications aren't is that languages are recursive.

Wait, what's the other -neme? [INAUDIBLE]

Phoneme. P-H-O-N-E-M-E. Languages are recursive. You can take a sentence, "The chair is red," and you put it inside another sentence, "The girl couldn't tell that the chair is red," and you can put that inside another sentence and say, "The girl that I met last week in the park couldn't tell that the chair is red." And at least in theory, you can just keep doing this, and still be building perfectly valid English sentences. This recursive ability lets you take one piece and nest it inside of another piece. It means you use the same pattern over and over again. Questions about languages, about other kinds of communication, about critters who communicate?

So that last one is, the meaning can be used [INAUDIBLE]

Hm?

The last [INAUDIBLE]

It means that you can take sentences and nest them inside of other sentences, more or less. In practice, if you're actually speaking sentences you want other people to be able to follow, if you nest more than about four deep, it starts-- your listeners won't be able to do it. This is where normal human ability to keep track starts breaking off. But in theory, it would still be a perfectly good English sentence-- it would just be a pain in the butt to read. Or hear, or make sense out of. All right. These are things that languages are that other forms of communication are not.

Switching gears now-- let's talk a little bit about-- we're going to talk about auditory perception. We're going to start out, move into thinking about language by talking about how your hearing system works. We talked about vision a couple of weeks ago. I'm in S&P, kids. You're going to have to sit through this again. Basic organs for seeing things with are your eyes, basic organs for hearing things with are your--

Ears.

Ears, good. All right. So ears. Here's an ear. There's an ear canal, kind of a big one. Ears have all kinds of funky little wobbly bits-- take a look around, look at the ears of your classmates. They probably all look more or less the same-- they probably all have pretty much the same pattern of ridges and swirls and swooshes. This outer ear part is called the pinna. There you go-- it's from the Latin word for "wing." This pinna structure is pretty much uniquely mammalian-- other animals don't seem to have these pinnas.

Humans can't really move our pinnas very much-- some of us can wiggle them a little bit. Anyone here who can wiggle their ears? I can't. I know-- I always wish I could do it. I'm not that cool.

[INAUDIBLE] there a point to [INAUDIBLE]

It makes little kids laugh at you. Which may or may not be a good thing, depending on your life goals. The point of having a pinna, having this fancy ear structure that sticks out of your head, is that it collects sound waves and it funnels them into the ear canal, so there's a little bit of amplifying-- it also shapes the sound waves, so certain frequencies get amplified and certain frequencies get diminished. The pinna collects sound waves, and the sound waves run down an auditory canal, like the ear canal, right?

Sound waves go down here, and they hit the ear drum, which has a fancy name. Some people who study sound call it the tympanic membrane. Now we're into the middle ear. So here's-- outer, middle. There's a little opening in here. Here's our eardrum. So sound waves come through the air, are collected by the outer ear, transmitted down the auditory canal there, and they cause your eardrum to vibrate. Then what? Anyone know where the smallest bones in your body are?

In the ear?

In your ear, right. There's three little bones in here.

[INAUDIBLE]

Three little bones in your ear-- there's one that looks kind of like this, that's called the hammer, and the other one that looks kind of like this-- it's called the anvil. And then one that has a very distinctive shape, and it's called the stirrup. So hammer, anvil, and stirrup-- they have fancy Latin names that some people use as well. What happens is, as sound vibrates the ear drum-- vibrates the tympanic membrane here-- then that motion gets transferred to this chain of bones. The bones are rigid-- all of this is, of course, tucked right inside your skull-- right in here. The bone that actually surrounds this whole setup is the densest bone in your body-- the temporal bone there, on the sides of your head.

One reason is to isolate all of the vibration that's happening inside here from getting jostled around by whatever else might be going on at the same time. The vibrations get transferred to these little bones, and these little bones bonk into an organ called the cochlea. In particular, the stirrup here pushes up against the oval window of the cochlea. The cochlea here-- we're getting into the inner ear.

How do you spell that?

C-O-C-H-L-E-A. Like that. All right, so we started out with vibrations that were traveling through the air around us. They get funneled into the auditory canal by the pinna-- they vibrate the eardrum, that tympanic membrane, which in turn vibrates this chain of little bones, which in turn pushes on the membrane here, on the oval window of the cochlea. The cochlea-- up to this point, we've been going through air. Now, we've been moving through different kind of bits, solid pieces. The cochlea's actually filled with fluid. When the oval window membrane gets pushed on by these bones, then the fluid inside the cochlea sloshes back and forth a little bit.

One of the other interesting things the middle ear can do is that these bones that are transferring vibration from the eardrum to the cochlea are-- there's a couple of little muscles in there that can control them. When these muscles are relaxed, then the joints between these bones are really loose, and the vibrations from the eardrum can make big vibrations on the cochlea. They can go back and forth between them, and it all moves a lot. If, for example, you hear something really loud, then very quickly, it is triggered for a muscle in here to tighten up so that this whole construction is a lot less flexible, and so that the amplitude of the waves in the cochlea end up being smaller so that it can kind of dampen the strength of the input to your perceptual system.

This can protect your inner ear from really loud noises. It also, like when you talk or cough or sneeze, anytime

you're kind of making a noise of your own, this would be really-- this should be really, really loud. It's coming from right here, right? Very close, and it's being transferred through bone. It should all be very loud relative to everything else, but we don't hear it that way, and that's partly because when you talk or cough or anything, then again, the muscles around these bones tighten up so that the whole assembly is stiffer, not so much vibration is transmitted, or smaller vibrations are transmitted-- not such big vibrations.

The interesting part of what we want to talk about, the cochlea, happens inside all of this coiled up part. So let's uncoil it, and uncoiled, it looks something like this. It's wider on one end than on the other end, and most relevantly, if you look at it in-- so here's kind of an overall view of it uncoiled. This is the base, and this is the apex. This is the part closer to down here. The thing you should know about the cochlea is, it is actually made up of three-- how do I want to draw this? Three kind of parallel canals that are all full of this fluid, that all run together.

They're broken up by a couple of membranes. The oval window's here, so when the stirrup-- here's our stirrup-- pushes on the oval window, then the fluid sloshes back and forth down between all of these canals, and then eventually comes back around towards the beginning, and there is around window right next to the oval window. Round window. That actually kind of bulges out in time-- the fluid can't compress when the stirrup presses on it.

It just moves, so you've got to have kind of an outlet spot that can also change shape as the fluid moves around. This would then kind of bulge and go in and out a little bit, but that's just so that there's room for the fluid to move. What we're most interested in on the cochlea here is this membrane right here. This is called the basilar membrane, and all of these canals have names-- I need to hit my cheat sheet to remember them. Sorry. Not an auditory kid. I'm a vision kid. Where are they?

You've got a middle canal. I know-- that one I could have figured out. A tympanic, and a vestibular canal. What we care about is the basilar membrane-- the membrane between the middle and the tympanic canal. What's interesting, and how this works, is that as the fluid in here is sloshing around, one of the things that happens-- because this is narrow at one end and wide at the other-- is that, depending on the pitch, the frequency of the sound that's causing the sloshing, different regions of the basilar canal will get moved around, and it will move a little bit in most spots, and a lot at one spot-- at a spot that is kind of tuned to the pitch of the sound.

Let's say this is the-- for a given pitch, this is the spot we're looking at. Here's our basilar membrane. All along the basilar membrane are what are called hair cells. A hair cell looks kind of like that-- it's right on the basilar membrane. At its top, it's got these little-- they're called stereocilia hairs. They're an excitatory cell-- they're going to have a lot of properties that are a lot like neurons. There are neurons that come off of them and go and form the auditory nerve.

The way a hair cell works is that they work kind of like how a touch receptor works. They're a mechanoreceptor-- they respond to a pressure on them. They get all of these stereocilia, these little hairs. Stereocilia. All of these little hairs are connected by these fine protein threads called tip links. What happens is there's a kind of a secondary smaller membrane right over any given region of the basilar canal, and these are embedded into this secondary membrane. Where's my secondary membrane? There it is.

As the basilar canal moves relative to this other membrane-- as the basilar membrane moves relative to this other membrane, then the stereocilia-- these hairs-- will get bent by the way these two are moving relative to each other. This will bend up a little bit or down a little bit, and it will bend the stereocilia such that their tips end up being further apart than they are when it's at rest. At that point, these tip links, these little fibers, get pulled because the tips are becoming further apart when the hairs get bent.

What happens is each of these tip links is connected to an ion channel. The tip link gets pulled, the ion channel opens up, and we get a very familiar looking process whereby potassium flows in, depolarizes the membrane of this hair cell. Lower down, it's going to have calcium ion channels that are voltage gated. As it depolarizes, calcium comes in, and then just like our neurons, there's little vesicles of neurotransmitter hanging out down here. When the calcium enters the cell, what does calcium do to vesicles of neurotransmitter?

Makes them release [INAUDIBLE]

Yeah, it makes them release it into the synapse, so it stimulates the auditory nerve that's bringing this information back to the brain. Did that make any sense at all? All right. We've got the cochlea. All of this comes in, it vibrates the fluid inside the cochlea, makes these waves of fluid. Depending on the pitch of the sound that's happening, different parts of the cochlea-- of the basilar membrane inside the cochlea-- will vibrate to a different amount.

When that happens, the hair cells that are on that spot will get the hairs on their tips, their stereocilia, bent over, which opens up an ion channel right on the tip. These tip links pull it open, more or less. It's like a little trap door. Potassium flows into the hair cell, and then stimulates calcium flowing into the hair cell, which then stimulates neurotransmitter being released, information going, and then it stimulates cells that become part of the auditory nerve going back into the brain.

All right, let's talk a little bit more about this idea about how pitch gets sorted out. Actually, let's back way up and talk a little bit about what actually is sound. We skipped right over the physics part here. Sound is vibration-- is pressure waves in the air, more or less. Right? Something vibrates, usually. The classic example for this is a tuning fork, right? Hit it on something and it goes bing. It's vibrating, and as it vibrates, it's kind of pushing the air molecules in its immediate vicinity up against each other, and then away, and then up against each other and away, and it's causing there to be these little patches of dense air, less dense, dense, less dense, dense, less

dense, dense, less dense.

If you think of a sound wave-- what people are usually really graphing at that point is the density of the air molecules in a given spot-- so you'd have a high pressure spot and a low pressure spot, and a high pressure spot and a low pressure spot, that's caused by the vibration of this item banging into all the air around it. Two important measures here-- frequency. That's a cube, really it is. Frequency, and amplitude. I know you guys all know this, but skipping over it anyway just make sure we're all on the same page with terminology.

Frequency is how long it takes for one cycle of the wave to repeat.

Would that be wavelength?

That is technically-- that is wavelength, yeah. You're right. I'm being unclear in my diagrams. Frequency is the time it takes one cycle of the wavelength to repeat, more or less.

Isn't that period?

Frequency of cycles per second-- we'll just give it with cycles for a second, and we'll go with it, all right? You want Hertz, I'll give you Hertz. Cycles per second, it's just the number of waves you get in a particular time unit. Amplitude is basically how big the difference is between the high point and low point in the wave. Not qualified to teach physics-- I fully concede that I'm not qualified to teach physics. Amplitude is the bigness between the bottom and the top of the wave.

You can, of course, change one of these things without changing the other. You get a wave that's higher frequency, which would then be lower wavelength, and has the same amplitude. You can have a wave that has higher amplitude and then has lower frequency. When the muscles that stiffen up the bone link here are kicking in, they're kicking in usually in response to high amplitude.

Right? When something is louder it's got a higher amplitude wavelength. Higher amplitude sound wave. All right, what happens if you have a complicated waveform that you're trying to analyze? This is where we get into something called Fourier analysis and Fourier synthesis. Fourier, French guy. 100, 150 years ago, something like. Basically said that, for any complicated waveform, you can break it down into a combination of the simpler wave forms that are then added together into a sine wave and its harmonics, and then add them together.

Fourier said, if we were to take this wave-- we'll do it in reverse-- [INAUDIBLE] a Fourier synthesis, and I take a wave that is-- let's see, can I draw it out? And I take a wave that is about the same amplitude but a much higher frequency, and I try to combine them, what do I end up getting? I don't even know if I can do this-- let's see. About kind of like that. There'll be low, and a little bit-- that's kind of what I was going for. It's not lined up well. But you

can take two waveforms, add them together at each point, and get a more complicated waveform.

You can do the same thing in reverse. If you have the right tools, you can decompose something like this into its component waveforms. Take this idea, and look at what's happening on the basilar membrane. The basilar membrane is tuned, right? Different parts of it respond to different frequencies. Stuff down here near the base responds to higher frequencies. Stuff up here near the apex responds to lower frequencies, and it's all just lined up right along the middle.

What happens is that, for complicated noises-- things like speech that you're hearing, or things like even the sound of a musical instrument would make versus the sound that a tuning fork makes-- that pure tone versus that tone that has sort of a timber to it-- then what's happening is, different parts of the basilar membrane will be stimulated, different amounts depending on how much of that frequency is in the sound that you're hearing. Parts of the basilar membrane respond to low frequencies, parts of it respond to high frequencies, and different parts will be stimulated to different amounts.

One theory for how we hear stuff says-- it's called place theory. It says that only that we depend on where the hair cells that are stimulating the brain are located on the basilar membrane in order to figure out what pitch we're hearing. One theory of pitch discrimination is called-- the place theory of pitch discrimination says that it's all about where the hair cells are located. That's the tool the brain uses to figure out what's going on.

The other competing theory is what's called volley theory. This says that it's not so much about where hair cells are-- it has to do with how fast they're firing. We know that we can hear stuff that is down-- so how fast the neurons are firing is either 1 to 1 with or is integer fraction of the frequency of the sound we're hearing. For a really low, couple hundred Hertz-frequency sound, you might actually get a couple of hundred Hertz firing rate of the neurons that are responding to it from those hair cells. Volley theory says that it's all about how the speed and the patterns in which the action potentials are coming in, and not so much about which individual cells are providing the input.

As with many things in neuroscience, it looks like your brain does a little bit of both. For low frequency sounds up to about 4,000 Hertz, then you see there does seem to be some amount of volley coding going on, where for really low frequency sounds, it's a one-to-one correlation between frequency of the sound coming in and frequency with which the hair cells are firing. Towards the higher ends of that, once you're up into a couple thousand Hertz, you're looking at more like it'd be a multiple of some kind-- the sound coming in would be a multiple of the firing rate.

But there is a relationship between the firing rate and the frequency of the sound. For sounds higher than that, it really just stops being practical to use firing rate as a coding for the frequency. Firing rate ceiling's at on the order

of 1,000 action potentials per second. We can hear a lot higher than that-- we can hear up to 20,000 Hertz. It's probably a fairly reasonable upper limit. Higher for some people, lower for others. Individual variations. Higher than that, it seems to be pretty dependent on place theory, on place coding, on which neurons on the basilar membrane are being most stimulated by the input that's coming in.

The other thing about-- all right, let's take this model, let's make it a little bit more complicated. On each cross-section of the basilar membrane, you've got a bunch of hair cells. You've got one inner hair cell over here, and then you've got a bunch of three outer hair cells. Same but kind of sloppy inner. Inner hair cells are running on the side that is closer to the curl of the whole thing, and then outer are further out. So you end up with for each of these, you end up with a row of hair cells all the way up the basilar membrane.

One row of inner hair cells that runs closer to the edge that is curling, and three rows of outer hair cells that move further out. And it turns out that it's only these inner hair cells that actually are sending in auditory perception information-- all of the nerves-- 90%, 95% of the nerves that are actually part of your auditory nerve that go into your brain come from these inner hair cells. These inner hair cells will have lots of neurons coming off of them, and they'll be almost all afferent nerves. They'll be sensory nerves sending information into the brain.

Outer hair cells will actually share nerves between a couple of different cells, and they've got a lot of efferent nerves coming in too-- information coming from the brain. One thing that these outer cells do is, as we know the basilar membrane, different locations on it respond to different pitches. But the level of precision with which we can detect pitch differences and the level of precision the basilar membrane has in discriminating pitches are very different. We're much better at pitch discrimination than simply looking at the physics of the basilar membrane would suggest.

One theory that's what's going on here is that there's a lot of information that comes from the brain down to these outer hair cells. The outer hair cells apparently are not so much about responding, but they can do something cool-- they can stiffen themselves right up so that they can't bend back and forth anymore. They make themselves a little bit longer, and the result of this is that their particular chunk of the basilar membrane starts being more or less flexible. This is one way in which the basilar membrane can actually tune to be more sensitive to very specific frequencies-- by controlling the behavior of these outer hair cells. Questions?

So this suggests that the brain sends information that will then reshape our interpretation of [INAUDIBLE].

Yes.

[INAUDIBLE] consciously?

Not consciously. I don't think anyone's managed to demonstrate anybody having conscious control over it. But for example, if you're trying to do a pitch discrimination task, or if you're trying to tune an instrument or something, where you really want to get it just right, then you might, for example, see your brain inhibiting one region of the basilar membrane in order to make stiffening one region of the basilar membrane to make other parts more sensitive. If you are a little bit off, then you'll hear it is a bigger difference-- because the next bit off won't respond at all, for example, by being able to tune it.

And it also seems like something similar is happening if you're trying to listen to one stimulus out of a whole bunch of stuff that's going on. Talking to your friend in a crowded room, listening to me when there's traffic on the street. Again, you're going to-- one of the things that happens is your basilar membrane gets tuned to be more sensitive to those frequencies-- the frequencies that are involved in that task-- and less sensitive to other stuff. It's actually almost a physiological attention mechanism, where you're-- even before it makes it into your brain, throwing out a certain amount of the potential sensory input that's out there.

Sound-- sound goes-- quick review. We've got our pinna, our outer ear. The funny shaping of the pinna causes it to select sounds, causes it to highlight certain frequencies of sound, so actually human ears are the funny shape they are in part because it means that sounds between 2000 and 5000 Hertz get amplified, and these are the frequencies that are important for speech. The pinna funnels sound down this auditory canal where the sound waves cause the ear drum, the tympanic membrane, to vibrate. That in turn vibrates these three little bones-- the ossicles-- so we've got our hammer, and our anvil, and our stirrup.

In turn, the stirrup presses on the oval window of the cochlea and causes the fluid inside the cochlea to slosh back and forth, which in turn causes this basilar membrane to wave up and down, and where it gets the most motion depends on the pitch of the sound that's coming in, the frequency of the sound. This in turn causes the hair cells that are on the basilar membrane to bend back and forth.

When they get bent, these tip links that connect to the stereocilia pull open these little ion channels. Potassium floods in, these hair cells are depolarized, and then that in turn triggers the opening of a voltage gated calcium channel. Calcium flows in, causes the vesicles of neurotransmitter to bind with the membrane, release their neurotransmitter, which in turn causes depolarization or hyperpolarization in these cells and the nerve that goes back to the brain.

Where the brain does all of this information go? Here's a brain. Looking at it kind of from the vertical view-- so here's an ear. We're going to disregard the part where those are completely different shapes. Here's an ear. Here's an ear-- I'm thinking we're kind of looking from the back of the head here, looking forward.

[INAUDIBLE]

We've got two parts of the brain, and then we've got the brainstem here, right? Coming down and going down into the spinal cord. Here's our ears-- we've got all of our kind of inner ear stuff. Here's our little cochleas. All right-- so nerve fibers come from the cochlea and they actually go to the brain stem first, and they go into what's called the cochlear nucleus right here. Cochlear nucleus. From there they actually-- most of the fibers actually cross and go to the superior olivary nucleus on the other side.

We've got a fair bit of processing happening right here. A few of these fibers go to the one on this side, but most of them cross over. Again, just like we're used to with all other aspects of the brain, stuff from the left goes to the right hemisphere, stuff from the right goes the left hemisphere. Here, it's happening in these auditory nuclei in the brain stem. There's a pathway that goes to the thalamus, and in this case, we're going to the medial geniculate nucleus of the thalamus-- remember, for vision, vision runs through the lateral geniculate nucleus. We're pretty much we're in a very similar piece of the thalamus-- slightly scooted over.

From there, we get processing that goes to auditory cortex out here in the temporal lobes.

[INAUDIBLE]

Yes. So we have our inner ear, right? Which is where transduction happens, where the mechanical energy of the sound wave turns into the electrical signal for a neuron. The nerve from that-- this is the vestibulocochlear nerve, which runs from the inner ear down into the brain stem. It goes to the cochlear nuclei down here. Of course, it's doing this on both sides, but I've only drawn this gentleman's-- or ladies, I don't know-- left, right ear. From the cochlear nuclei-- it goes to the cochlear nucleus on the same side, and then most of the signal crosses. It goes to the superior olivary nucleus on the other side. A little bit of it stays on the same side, but it mostly crosses.

Question?

[? No. ?]

OK. We're still in the brainstem here, so this is right in the back of your head, right? Right down here. And from there, the superior olivary nucleus sends cells that project to the thalamus. Remember, the thalamus is this gateway for sensory information-- it takes it, looks at it, and sends it off to the correct piece of cortex. Here we're looking at the medial geniculate nucleus of the thalamus. Remember for vision, we looked as the-- it went through the lateral geniculate nucleus, the LGN. In this case, we're kind of in the same region, but a different clump of cell bodies closer to the midline-- medial. Lateral of course is closer to the sides.

From there, it spreads out and it goes to a bunch of auditory cortex. Where the cortical processing actually is in your temporal lobes right on the sides of your head here, but the signal gets there in a kind of roundabout

manner-- it's kind of easy to be like, hey, look, ears! Well, that's where my auditory processing is, but it actually goes to the back and across and around before getting back to the temporal lobe.

[INAUDIBLE]

Voila, the major auditory pathways.

So wait-- sound begins from one ear, goes to the other side, [INAUDIBLE]

Almost all of it crosses. Almost all the input from your left ear goes to your right brain, and the other way around. About 90% or 95% of the nerve fibers cross, and a little bit of it goes to the olivary nucleus on the same side. It turns out, in some animals, that the information about where a sound is coming from happens much later in the stage-- but in mammals, it happens right down here, in these early processing stuff in the superior olivary nucleus. Superior olivary nucleus gets mostly input from one side or the other, but remember, that we've got 5% or 10% of the signal didn't cross over, and that's the part that your brain uses to locate sounds.

We're pretty good at figuring out where sounds are coming from. Try it-- close your eyes. I'm going to-- put your stuff down. I'm going to walk around the room. I'm going to clap my hands. Point out where you think I am without looking.

Oh, sorry about that.

Don't take out your neighbors.

[INAUDIBLE]

All right, feeling fairly accurate? Did you feel like you had a pretty good sense of where sounds were coming from? It's not perfect, but fairly accurate. We're good at this. What kinds of information are we using to do that? How do you-- how would you figure out-- how could you figure out where a sound is located? How can you tell?

By which ear it's closer to.

By which ear it's closer to. How would you analyze the sound signal coming into your ears to figure out which ear it's closer to?

[INAUDIBLE]

Yeah?

Is there a time difference at which the sound wave reaches [INAUDIBLE]

Yep. Your ears are going to be six, eight inches apart, probably, on average. There's going to be a sound difference. One is latency between ears. Another is loudness between the ears. So latency and loudness difference let you figure out which ear things are closer to, especially for things that are not directly in between them. Right? Telling front from back is hard by latency and loudness, but you've got-- here's our hypothetical guy, right? If I've got a sound source over here-- sound source. Here's a tuning fork. Then we got sound waves coming off of it, and they're going to hit his right ear before they hit the left ear. So there's a latency difference.

This doesn't work for really low pitched sounds because the wavelength of the sound wave-- and I do mean wavelength this time-- is big enough that it actually would just go around your head. Your head isn't big enough to make a significant difference in how it'll get there, no. Wow, I can't teach this at all today. Never mind that-- that's true for the next thing I'm going to talk about, which is that you get-- which is loudness.

You get a sound shadow around your head. For a sound that's coming through, if it's coming to your left ear, it'll block-- your head will actually block some of that sound wave getting to your right ear for high pitched noises. For low pitched noises, the wavelength is big enough that it can just go around your head, and you don't get a loudness difference. Does that make sense, physics kids?

So [INAUDIBLE]

Latency, loudness, and some other kinds of cues too. There's a lot of cues-- there are some cues that your brain can use to localize sounds.

Like you're saying that for large wavelength and loud?

Not for loudness, but for-- latency is one thing that you get-- you get a difference in time to arrival. For loudness differences, if something is closer to your left ear, it's going to be louder in your left ear for reasonably close things, especially. But not for lower pitched larger wavelength sounds, because the larger wavelength allows it-- that means it is not getting blocked by your head, which isn't big enough relative to the sound to cast that kind of shadow.

Another thing that you'll see is you'll actually see spectral differences. What that means is that, because of the way our ears are shaped, sounds that are coming in from different angles will flow over and around the funky little crinkly shapes in our ears-- in our outer ears and our pinna-- differently. A sound that's coming from in front of you and a sound that's coming from behind you, even if they're the same sound, will have subtly different characteristics-- which frequencies are highlighted, which frequencies are decreased a little bit.

That information, we're not usually consciously aware of it. Our brain kind of goes out of its way to make us less aware of it for conscious processing, but you use it-- for localization, you become aware of whether something is behind or in front of you. That funny shaped outer ear not only highlights particular frequencies that we want to be able to pick out of our world-- it also allows us to get some localization information by how it affects things coming from different angles. Questions?