

[SQUEAKING]

[RUSTLING]

[CLICKING]

JOSH All right. Fantastic. So last time we started talking about sound localization. And so what I'm going to do today is
MCDERMOTT: wrap that up. We're going to have a brief interlude to discuss psychophysical methods.

Everybody's going to learn how to measure a threshold. And then we're going to talk about pitch perception And then following on that, next time, we'll get into auditory scene analysis and then speech perception

So we talked about the fact that sound localization is this really interesting problem because location information is not made explicit in the sensory input. And we discussed the sources of information that we get about sound localization, these different kinds of cues that are informative about where something is coming from in the world. So there are two that are binaural, interaural intensity or level differences and interaural time differences. And then spectral cues that come from the filtering of the pinna that are monaural.

So the binaural cues, they originate from the fact that the ears are in different locations, and so there are differences in the path length. And then there's also an acoustic shadow that is cast by the head. We looked at the magnitude of these cues, in particular for the time differences that exist between the two ears, they're really, really small, but they're reliable.

We talked about the idea that time differences are arguably most useful for low frequencies, because for high frequencies, they are ambiguous, and potentially also limited by the upper cutoff of phase locking. Conversely, level differences are largest for high frequencies, and they get smaller for low frequencies, just due to the fact that the low frequencies pass over the head. And so there's this classical kind of cartoon story of binaural hearing called the duplex theory, where for low frequencies, we can rely on time differences, and for high frequencies, we can use intensity differences.

And then we talked about some evidence that was taken as being consistent with duplex theory, namely that localization is best for low frequencies and high frequencies, and gets a little bit worse for medium frequencies when you test people with pure tones. And then we talked a little bit about something that might be useful for the problem set.

We then discussed the kind of classical model for how you might detect timing differences between the two ears, namely the Jeffress model for ITD detection, where the idea is that you have neurons that can act as coincidence detectors. They get inputs from the two ears, but with a time delay-- a different time delay for the two ears, and that causes them to respond most when there's a particular time difference between the two ears. So really like possibly, I think the first model in the history of computational neuroscience, which is kind of cool.

And we talked about how this has been found in the brainstem of some species. And then how there's this kind of remarkable segregation of function, where two different parts of the brainstem have sensitivity to these two main binaural cues. So the intensity differences are believed to be detected in a different part of the superior olive. So you've got the segregation of function. The medial superior olive contains ITD detectors, and the lateral superior olive contains ILD detectors.

And we talked about how when you get to the cortex, the coding changes, and that the neurons tend to be cue independent and broadly tuned to one side or the other. And then got into the issue with these binaural cues, which is that they're informative, but they're also ambiguous. So there's the cone of confusion, which is this region of space that is consistent with particular binaural cues, just due to the symmetry of the situation. And so to resolve that ambiguity, we're believed to rely on these pinna cues.

And the pinna cue is derived from the fact that the ears have these folds such that depending on the direction that the sound is coming from, it will bounce off the ears in a particular pattern, and that causes it to be filtered differently. Everybody's ears filter the sound slightly differently, so this is a graph showing you two different ears. And the point is there are these transfer functions that changes a function of the direction from which the sound comes from, whether it's up or down.

And here's the concept where there's a transfer function in your ears. That's the filter of your ear. That colors the spectrum of the sound that comes in, and so you get some altered spectrum. And so in principle, you might be able to detect the transfer function and use that to localize sounds in elevation.

Now, we discussed some evidence that people actually use this. Namely, when you put plastic molds in somebody's ears, their ability to localize in elevation collapses. They preserve their ability to localize in azimuth.

And then we talked about evidence that people actually learn the cues from their pinna, because if you make people go home with these plastic molds and walk around with them for several weeks, they slowly regain the ability to localize in elevation. And I showed you very briefly an example of a computational model that is optimized to localize sounds that reproduces a lot of these effects. So systems that learn to localize sounds end up behaving a lot like humans.

And then at the end of the lecture last time, we talked about stereo reproduction of audio, and how it's imperfect because it can't completely reproduce the spatial cues that actually happen in the world, but nonetheless, kind of does help you to distinguish different sounds in music. And that's where we left off.

And so the last topic that we'll talk about in sound localization is distance perception. So so far, we've been talking about how you tell the direction that a sound is coming from, whether it's here, or here, or here, or here, but obviously, another dimension to location in the world is how far away something is. And one of the main ways in which we determine how far away something is actually reverberation. And so here's-- you'll just be able to hear the effect here.

So here's somebody who's real close.

[AUDIO PLAYBACK]

- [INAUDIBLE]

JOSH And here's somebody who's far.

MCDERMOTT:

- [INAUDIBLE]

[END PLAYBACK]

JOSH It's supposed to sound further away. These reverb demos are always challenging in a classroom because they're

MCDERMOTT: played out of speakers and are subject to all the reverberation that happens in the room. But here's the concept.

So sound attenuates with distance. So if somebody is right next to you, there is sound-- so let's suppose you're this person, and there's this person in back of you who's talking. So there's sound that reaches you from that person directly to your ears. And that sound will typically be pretty high in intensity because they're close.

This person up here is much further away. And so the sound will attenuate over distance. And so the direct sound that reaches you straight from the sound source will be lower in amplitude when the person's further away compared to when the person's here.

Now, the reverberation will tend to be pretty constant. And the idea behind-- so reverberation, that's the term that we use to refer to all of the reflections that can occur in the world. So most surfaces in the world, they reflect sound. And so when a sound source is producing sound, the sound will reach you directly on a kind of a straight line. But also after reflecting off of all the different surfaces in the environment.

And so the sum total of all those reflections is what we call reverberation. And you can recognize that when you go into a big space. You walk out into the atrium, it's a big space. There's like lots of echoes and the path lengths are kind of long, whereas like if you're in your car, the path lengths are very short, so it sounds different.

But the key concept here is that the path lengths for reverberation, if you're in a space, they all kind average out. So some of them are short and some are long. And so the overall level of reverberation will be pretty similar irrespective of whether the person is close to you or far from you. But what will change is the amplitude of the direct sound.

So if the person's really close, that direct component of the sound will be very high in level, whereas if they're far away, it will be lower in level. So that's kind of one of the main cues to sound distance. So this is just kind of a diagram that shows that. So what it's attempting to depict is a situation where the source is near, that's on top, a situation where the source is more distant.

And this is like all of the different reflections that you would receive. So imagine you play a clip from the sound source. So this is the sound that reaches you from the direct path. And then these are all of the reflections, and they arrive later because the path length that they have to take in order to reach you is longer. So they arrive later, but their amplitude is also a bit lower.

And so the idea is that if something is close, that direct sound is high in amplitude. If it's further away, it's lower in amplitude, but the reverberation is pretty similar. And so the ratio direct to reverberant sound gives you information about how far away something is.

Also, another thing that you can use to tell how far away something is if the sound source is familiar and has a characteristic level. So for instance, like dogs barking, they typically bark at a certain SPL level of-- I don't know-- 70 dB or something, right? And if you're familiar with that sound source, then that can tell you how far away they are, potentially. So if the sound that arrives at your ears is low in level, and you realize that it's a sound that's normally high in level, that tells you it's far away. So those are two different kinds of cues that you can use.

So that's all that I got to tell you about sound localization. What questions do you have before we move on? Yeah.

AUDIENCE: So let's say it's an unfamiliar sound. If there were just like, no walls at all or [INAUDIBLE] would it just be hard to perceive distance?

JOSH
MCDERMOTT: Yeah. You would have no idea how far away the thing is. Yeah. Yeah. And in fact, yeah, you could go into an anechoic room and try to test this. Yeah. Yeah. Yeah.

Yeah. And I should say the-- so reverberation is one of these things that it-- I mean, it's everywhere in the world. It's also everywhere in music production. It's maybe the most commonly employed method in music production to change the way things sound. And so if you pay attention to recorded music, like the production engineer is often manipulating the reverberation to make certain things sound like they're in different places.

So sometimes a lead vocal will actually be pretty dry, so it kind of sounds like it's very close. Like the backing vocals will have more reverberation on them to make them sound like they're further away. So it's a very common kind of production trick. And you're leveraging the fact that your brain kind of understands the statistics of reverb from growing up in the world. Any other questions about sound localization or distance or reverb?

So let's talk about psychophysical methods. So oftentimes, when we are studying perception, we deal with thresholds. So we've talked about the idea of a just-noticeable difference. So that's kind of like the smallest amount that you can discriminate. So we encountered this when we were talking about intensity.

So remember like we talked about how one reason why it's often convenient to describe sounds with the decibel scale is that the just-noticeable difference is usually about one decibel. And that's true irrespective of whether the sound is pretty low in level, say, 40 dB SPL, or pretty high in level, say, 90 dB SPL. So that's just noticeable difference. It's a threshold.

And so what I want to talk about now is how you would go about-- is really what we mean by threshold and how you would go about measuring that, because it ends up being pretty broadly relevant. Historically, there's two types of thresholds that people talk about. One is what's called an absolute threshold. So in the context of sound, this is the minimum audible signal.

So this is what is measured when you get your audiogram measured. So remember, when we were talking about hearing impairment, we showed the audiogram. That's the detection threshold for different frequencies. It's the thing that kind of looks a U. Worse at low frequencies and high frequencies, best at medium frequencies. So that's an absolute threshold.

The other kind of threshold is a differential threshold or a discrimination threshold. And this is the minimum perceptible change. I like to refer to this as the Just-Noticeable Difference, or JND. I like it because it's self-explanatory.

The more classical term in hearing research is difference limen. I don't really like it because it's sort of inscrutable, but that's just what people classically call that, and so you'll often read papers that talk about the DL. It's just a discrimination threshold. So that's what they called it in the old days.

So naively, a threshold is a stimulus value that you can detect. So let's say we're talking about the absolute threshold. So I'm going to play you a 1-kilohertz tone, and I want to determine your threshold right. So we're going to vary the intensity of that 1-kilohertz tone.

And the naive notion is that-- let's suppose this is a graph where this is the stimulus level, and this is the proportion of time that you will detect things. Naively, the idea of the threshold is that there is this special value here, where if the stimulus is below the threshold, you don't hear it, so you're at 0%. And if the stimulus is above the threshold, you hear it, meaning 100% of the time you can detect it.

This is what is known as a psychometric function. So it plots performance. In this case, it's the percent of times you detect the signal as a function of a stimulus parameter, for instance, the stimulus level. So that's kind of like the naive idea of what might happen and what might be associated with the threshold.

Now, in practice, anytime you actually do an experiment to try to measure this, what you get is something that looks more like a sigmoid. So when the stimulus level is very, very low, people will never hear it. When you get to these intermediate values, sometimes people will say they hear it, sometimes they will say they don't. And then you'll get into this regime where people will always say that they hear it.

So in practice, when we are talking about a threshold, and you're dealing with real data, you assume some kind of arbitrary level of performance, and you just define that as the threshold. So maybe 75% correct, for instance. But the threshold is typically associated with the psychometric function.

So how would you measure this? So maybe like the simplest type of experiment that you might do is the one that they often do when they will measure an audiogram in a doctor's office. And that's what would be called a one-interval, two-alternative task, or a yes-no task. So we're always playing you a tone, and you have to say whether you hear the tone or not.

So is there a signal present? So we could be presenting tones and noise. That's what we were doing when we talked about masking, and we were trying to infer auditory filters. But there's lots of other kind of applications with this kind of thing. So for instance, you might want to measure how well are air traffic controllers can detect aircraft on a radar screen, or you might want to measure how well doctors can detect tumors on X-ray images.

These detection tasks are very widespread. And so the framework that we will talk about for thinking about them is very broadly relevant. And the framework is known as signal detection theory. And so the key idea is that this kind of very simple experiment that you might naturally think to do is affected by bias. And what I mean by that is that the results of the experiment will vary depending on how inclined the observer is to say yes, independent of how well they actually hear or see or whatever it is you're trying to measure.

So here, we have observer 1 and observer 2. We've done the classic experiment to try to measure an audiogram, which is we have all these trials where the signal is present. So let's say it's a 1-kilohertz tone at 10 dB SPL, pretty close to people's threshold. So observer 1 says yes to 90% of the tones. Observer 2 says yes only 65% of the time.

Looks like observer 1 is better. But the problem here is that observer 1 just really likes to say yes. And this is revealed when we also include trials where the signal is absent. So you can see that the response proportions for observer 1 are exactly the same when the signal is absent as when the signal is present. They're just biased.

Whereas observer 2, these response proportions change a lot. So when the signal is actually present, they say yes a lot more of the time than when the signal is absent. So it's kind of clear from looking at these numbers that you actually have no evidence that observer 1 can hear anything, whereas you have a fair bit of evidence that observer 2 is actually sensitive to whether the signal is there or not.

So the key concept here is that if you are doing an experiment like this, where there is a single interval and the person either says, yes, I hear it, or no, I don't, you have to consider both what are called hits and false alarms. So in this kind of experiment, there are two types of trials. The first type of trial is where the signal is present. So we play the beep.

The second type of trial is where there is no signal, so the beep doesn't get played. Now, there are two types of responses. The person can say, yes, I hear something, or no, I didn't hear anything. And so that gives us four types of trial-response combinations

A hit is a trial where the signal is present and the person correctly says, yes, I heard something. A miss is a trial where the signal is present and the person says, no, I didn't hear anything. And if we are computing proportions, if you know the hit rate-- so that's the proportion of the signal trials that the person said yes on-- you also know the miss rate because these two things have to add up to 1.

When there is not a signal present and the person incorrectly says, yes, I heard something, that's called a false alarm. And if the signal is not present and the person correctly says no, I didn't hear anything, that's called a correct rejection. So again, the proportion of the false alarms and the correct rejections have to add up to 1, so you really only need to worry about one of them.

And so the performance of a person in this experiment can be summarized by these two numbers, the hit rate and the false alarm rate. And so the key idea is that if you actually want to measure how sensitive the person is to the signal, you have to measure both hits and false alarms. One of them on their own is not going to tell you how sensitive they are. And this is because subjects or participants with the same sensitivity can have different biases. So one subject may err in the direction of saying the signal is there more often than it is, while another subject may err in the other direction.

So the framework that we use to understand this stuff, as I alluded to earlier, is what's called signal detection theory. And this picture here, this is kind of the classical picture that is associated with signal detection theory, and it's really worth spending some time to make sure that you understand this. So the central assumption of signal detection theory is that a stimulus elicits a response in the nervous system, and you base your decisions on that response.

And that response, it is a function of the stimulus, but it also contains noise. So the internal response x to a stimulus can be represented thus as a random variable. Now, typically we will assume that random variable has a Gaussian distribution because it kind of just makes the math simpler. And because of the central limit theorem, lots of things are often Gaussian, so it's probably not a terrible assumption.

But the key consequence of this assumption of internal noise is that two presentations of the same stimulus will not necessarily result in identical percepts. So especially when we get close to threshold, we can play the exact same stimulus, sometimes you'll feel like you hear it, sometimes you might not.

So the consequence of this is that all of the trials where the signal is absent are associated with some particular distribution of internal responses. So these are probability distributions. The y-axis here is probability, the x-axis is the internal response.

And then all the trials that are associated with the signal being present also give rise to a distribution of internal response. So these are these Gaussian bumps. So this is like the distribution of the internal response over lots of trials where the signal is present, and over lots of other trials where the signal is absent. And you can see that the distributions overlap, and that's just because of noise.

So you could think of the task of the observer. So let's suppose that we're just doing the detection task. You're trying to determine whether the signal was present or not. So your brain gives you a particular value on this axis. That's the internal response.

You don't know whether that came from the red distribution or the blue distribution. Only the experimenter knows that because they're the person who determined whether the stimulus was present or not. And so as the observer, as the participant in this experiment, you get a value on this axis, and you have to decide, did it come from the red distribution or the blue distribution? And because these distributions overlap, it will not be possible to do this perfectly. You're going to get some things wrong.

So a key quantity in signal detection theory is this measure of how separated these distributions are. And so the intuition here is that if a signal is very detectable-- so let's suppose we're presenting a tone that's like at 70 dB SPL, so way above threshold-- these distributions will be very far apart. So they won't overlap very much. And that means that you won't have to make very many errors.

But as you get close to threshold, those distributions will get closer and closer together. So the measure of how sensitive you are to the presence of the stimulus is given by d' , which is the difference in the means of these two distributions divided by the standard deviation. And for simplicity, we typically assume that the standard deviation is the same for the signal present and the signal absent distribution. So there's some constant level of internal noise that just gets added to all responses.

But the difficulty of distinguishing whether the stimulus is presented or not is determined by how much the distributions overlap. So here's a situation where they're separated by a fair bit. Here's a situation where they are closer.

So if you're doing an experiment, I said we've got to pay attention to both hits and false alarms. And so this quantity, d' , that is a measure of how separated these distributions are, can be estimated by taking the hit rate and the false alarm rate passing them through z , which is the inverse of the standard normal cumulative distribution function-- and that's because we assume that the internal noise is Gaussian-- and then taking the difference.

But conceptually, you can think of this as like this is a function of the hit rate and it's a function of the false alarm rate, and we're taking the difference. So what does that mean? Well, it means when the hit rate is a lot greater than the false alarm rate, you're more sensitive. And as those two things become more similar, you're less sensitive. And if they're equal, so if the hit rate is the same as the false alarm rate, then really, we have no evidence that you can detect the signal at all.

So if you're actually doing this task, though, you have to set a criterion. So remember, as an observer performing this experiment, your nervous system is giving you an internal response, so some value on this axis. And you got to make a decision. Is the signal there or not?

And that requires setting a criterion, C , where if the response exceeds the criterion, you say, yes, I heard something. And if the response is below the criterion you say, no, I didn't hear something.

And so the key concept here is that your sensitivity does not depend on where you put the criterion. So here's a situation where the criterion is up a lot higher. So that corresponds to somebody being very conservative. This means that you only say that you heard something when that internal response is very, very high. Whereas back here, you're more liberal, and you're going to say yes to everything that exceeds this value.

So how do the hit rate and the false alarm rate relate to this particular diagram? Well, the hit rate, remember, is the proportion of trials where the signal is actually present, where you say, yes, I heard something. Now, what we're positing here is that you're going to say that you heard something any time the internal response exceeds the criterion. So that's any time the response is above here.

Now, for the trials where the signal is present, we've got the red distribution here. And so the proportion of trials where the target is present for which you will respond yes is going to be going to be the integral of this red distribution from the criterion all the way up to positive infinity. So that's going to be the hit rate. It's the area under the red curve.

And similarly, the false alarm rate will be the area under the blue curve that exceeds that same criterion. And you can see here, that this area under the red curve is a lot bigger than this area under the blue curve. And so the person was going to have a positive value of d' . Yeah?

AUDIENCE: Is this criterion always constant or is this a Fourier line that it's been shown on?

JOSH That's a great question. So it could be constant. But really, this is a function of the observer. And the criterion
MCDERMOTT: could, in practice, drift around. And in practice, one of the things that is challenging about running an experiment like this, where there's just a single interval, is the fact that the observer has to figure out where to set their criterion.

And they may initially not be very sure, and it may change from trial to trial. And so that's one of the messy things that you sometimes have to deal with with this type of experiment. But this particular framework here, we are, for the moment, assuming that the criterion is a fixed value. So if you want to compute a hit rate and false alarm rate, you're assuming a fixed criterion. Yeah.

So there's an alternative way of doing the experiment where life gets a little bit simpler. And that's called a two-interval, two alternative forced choice experiment. And so in this type of experiment, on a given trial, there are two intervals, interval 1 and interval 2. This is an example experiment where the task is to detect a tone that is being played in noise. So this is like one of these masking experiments that we talked about a couple lectures ago.

So we've got the signal here and the noise. One of the intervals-- here, it's depicted as interval 1, but it would randomly be assigned to interval 1 or interval 2. One of the intervals will contain the tone. The other interval will just contain the noise. And now the task is just to determine which interval the signal is in. Is it the first or the second?

So the advantage of this type of task is that as a participant, you don't have to figure out where to set your criterion. You're just comparing the two intervals and trying to decide which one contains the tone. And that reduces the effects of subject bias, and it reduces this issue with the criterion possibly moving around and the person having to figure out where to set it. But you can map this kind of experiment onto the signal detection framework just as well.

So now the task is to decide which of the intervals actually contains your signal. And so the idea is that on every trial, you're getting one sample from the red distribution and one sample from the blue distribution, because one of the intervals has your signal and the other interval doesn't. And the decision rule is just going to be whether the internal response is greater in the first interval compared to the second interval.

Now, in this particular trial, you can see that the red sample-- so the internal response from the interval where the signal was present-- is greater than the blue sample, the interval where the signal was absent. And so here, you would get the trial correct because this is higher than this. But that doesn't always have to be the case because these responses are, to some extent, random.

So here would be a case where you got unlucky and your brain gave you a bigger response on the interval where the signal was absent. But in this particular case, over time, over lots of trials, you'd get things right more than you would get them wrong. And so your proportion correct would be above chance.

So you can use this 2AFC task, and the advantage here is you don't have to-- there's no hits, there's no false alarms. You just measure percent correct, and that's the measure of sensitivity. Whereas for the one-interval tasks, percent correct is vulnerable to bias, so you have to measure both hits and false alarms and compute d' . And then the person has to figure out where to put their criterion.

So we've got these two types of psychophysical tasks, one-interval tasks and two-interval, two alternative forced choice tasks. There are various advantages to the 2AFC task.

Now, how do you measure a threshold? So the simplest way to do this is to generate a psychometric function. And the way that you do that is what is called the method of constant stimuli. I've always found this to be a confusing way to describe this because it somehow makes me feel like it's an experiment where the stimuli never changed or something, which is not what it means.

What it means is that you have to decide beforehand what all the stimuli are going to be that you present in the experiment. And then you run a bunch of trials of each because this is what's going to be the outcome of the experiment. You're going to generate a psychometric function-- remember, the psychometric function plots performance. So if this is a 2AFC task, the performance here would be percent correct. If it was a one-interval task, this would be d' versus some stimulus parameter. In this case, the signal level if we're measuring a detection threshold.

So for a 2AFC task, chance performance is 50%. So when the signal level is below your threshold, you'll just be guessing, and you'll get 50% right on average. And then as the signal level increases, you get better and better. And so you get this kind of sigmoid-like curve. And then, typically we define the threshold as being the point that gives you 75% correct, for instance. And that would be about 57 dB in this particular case.

So it's called the method of constant stimuli because before you do the experiment, you have to decide what values on the x-axis you're going to present. So you're going to have a whole bunch of trials where the signal is 51 dB, maybe 100 of those. And then you measure how many of those people get correct. And then a whole bunch of trials where it's 52, 53, 54, so on and so forth, all the way up to 63.

So the experimenter decides that ahead of time, you run all those trials, you measure performance as a function of the stimulus level. You get your psychometric function, and you estimate your threshold from the psychometric function. And that can be a fine thing to do. Psychometric functions are beautiful. They look nice in papers. It can often be a good way to make sure that things make sense, that performance actually goes up as you increase the stimulus parameter.

On the other hand, you have to know beforehand what the relevant dynamic range is. If you want to hit the threshold, and you happen to choose stimulus values that are out here, or down here, you're not going to get a very useful result. And the other thing that's kind of suboptimal about it is that you spend quite a lot of trials on this region of the curve, and on this region of the curve, which really are not very relevant to constraining the threshold. The threshold is mostly determined by these kinds of values, so it's sort of inefficient.

So to address these issues there's an alternative way to measure thresholds, which is what is called an adaptive procedure, also sometimes referred to as a staircase. And so the idea of the adaptive procedure is that you start people off with kind of an easy version of the task. So you set the stimulus parameter to a value that you're pretty confident everybody's going to be able to detect.

And if they get things-- if they get things right, then you make the task harder. If they continue to get it right, you make it even harder. And then once they start to get things wrong, you make it a little bit easier. And what happens if you do this is, in theory, you should converge to a value that corresponds to the threshold.

So that's the kind of intuitive picture. One kind of classic way to do this is what's called an X-down Y-up adaptive procedure. So for instance, two down, one up. So what does that mean? That means if you get two in a row correct, you move the stimulus level down one step. If you get something wrong, you move the stimulus level up one step.

So here's how this would work. So you start out your experiment with some pretty high stimulus level. Person does two trials, they get both of them correct, so the stimulus level is lowered by the experimenter. The person gets another two correct, it's lowered. Get another two correct, it's lowered.

They get one correct, oh, they get one wrong. So that means that the stimulus level goes up. They get two correct, so it gets lowered. They get two correct, it gets lowered. They get one wrong, it goes up. They get one wrong, it goes up. Two correct, it goes down. Two correct, it goes down. One wrong, it goes up, and so on and so forth.

And so what happens is that you kind of, in theory, find your way to this regime where people are getting some of the trials right and some of them wrong. And so for instance, this particular adaptive procedure rule, the two down, one up rule will converge to a stimulus value that will give you 70.7% correct. And you can work out exactly why that is.

So what are the advantages of this? Well, in theory, it should be more efficient because you don't spend a ton of time with trials that are super easy or super hard. You find your way to the relevant part of the stimulus space. It also has the nice advantage that you can start out the experiment really easy, so the person can figure out what they're supposed to be doing. So you give them a bunch of easy trials, they build up their confidence and stuff, and then you move down into the more difficult regime.

The disadvantage of this is that you don't end up with the full psychometric function. You just get a threshold estimate to come out of this. And so sometimes you might want to know what the actual psychometric function kind of looks like.

So just to summarize, I talked about the distinction between one-interval yes/no tasks and two-interval 2AFC tasks. So key thing to take away from this is that with one-interval tasks, you have to measure both hits and false alarms and compute d' as a measure of sensitivity. You can't use percent correct.

And the challenge is that the observer has to set a criterion. And as we discussed, maybe that would move around, maybe they wouldn't know where to put it. If you use a 2AFC task, you can just measure percent correct, and that gives you a bias-free measure of sensitivity.

And then we discussed two ways to measure a threshold. One is the method of constant stimuli, where you decide on a bunch of stimulus values that you want to measure performance at, you get a psychometric function. The other is an adaptive procedure, sometimes called a staircase, where you converge on a stimulus value where people are getting some of them right and some of them wrong. That will just give you a threshold estimate as opposed to the psychometric function. And the threshold in both cases is defined as the stimulus parameter that is needed to achieve some criterion level of performance, like 75% correct.

That is the whirlwind guide to psychophysical methods. That's how you measure thresholds. Next time, we will talk about pitch perception and pitch discrimination thresholds. So I will see you then.