PROFESSOR: Well, it looks like it's 7:00, more or less.

It looks like there's some people here.

Ooh, this is very-- I haven't lectured in here for years.

We did 900 in here one year.

I don't like it, except maybe they got the heat working this year.

Years ago, the heat didn't work in here.

It got kind of harsh in the winter.

This looks terrific.

I'm sure we'll have to erase some of it to explain stuff, but this looks so much-- some people are just better at drawing on the blackboard.

And, but I'll erase it as needed.

So what can I tell you?

I'm there.

Look at that.

STUDENT: [INAUDIBLE]. PROFESSOR: What was the-- it doesn't include tomorrow's lecture on animal language and stuff.

It does include the stuff on adult human language from last time.

STUDENT: Lecture [INAUDIBLE] language and acquisition.

PROFESSOR: Well, the 10 and 11 are both called language and acquisition, or something like that.

But anyway, last time's lecture is there.

Next time's lecture isn't.

So there won't be a question about whether or not apes can talk to you, and there will be a question on whether you can talk to you.

STUDENT: [INAUDIBLE]. PROFESSOR: Probably. Isn't-- I thought it's chapter 9.

What's the language chapter?

So the language chapter has a large chunk on regular adult language.

I would certainly think that would be a good thing to know something about because it's the same material as was in the lecture.

I wouldn't sweat the development stuff too much this time because it won't-- the exam must be in a slightly different position because usually this isn't where the break is.

But anyway, adult language, not kid language, not ape language. Yes, can I go home now? So as previously advertised, I will not stand up here and review the course.

I will respond to pithy questions from members of the audience.

In the absence of pithy questions, I will stand here for a short period of time, staring blankly at as you stare blankly at me.

And then, we'll all go home in a faintly disappointed kind of way.

So what do you got there for me?

STUDENT: [INAUDIBLE].

PROFESSOR: You've got to yell in here.

The acoustics are miserable.

STUDENT: Oh, can you go over the difference between negative reinforcement and punishment?

PROFESSOR: Negative reinforcement and punishment-- reinforcement is always a good thing, even if it's negative reinforcement.

That's the way to try to get this into your memory.

I think, if my memory serves correctly, I bungled this in describing it at least once in class.

But reinforcement is always a good thing.

So if something says, "negative reinforcement," there must be a good thing happening.

So a negative reinforcement is something like the baby stops crying.

If you're a baby minder, you'll do stuff to make that baby stop crying.

The cessation of something nasty is a negative reinforcer.

Punishment is common-sensically what you think it is.

If I whack you every time you use whatever that is, a purple pen, you'll stop using a purple pen because you're not stupid.

STUDENT: Yeah, can you talk about center and surround, and kind of visual perception, how that works. PROFESSOR: I like the first part of that question.

That sounded precise enough to-- a curved piece of chalk.

All right, so an object, separate energy, all right, so we're going to get rid of the spring.

STUDENT: Great eraser.

PROFESSOR: What?

STUDENT: Great eraser.

PROFESSOR: Great eraser?

STUDENT: Oh, white.

PROFESSOR: Well, it'll have to do.

So center/surround refers to receptive field organization.

And it's receptive field organization-- I'm going to leave that because it's pretty. So one place to see it, actually, is in the skin.

If you go and find a chunk of-- here's the brain, here's somatosensory cortex here-- if you go and grab one little cell there, what you might discover is that it's got a receptive field, let's say on the arm, such that if you touch the arm here, that cell gets active.

But characteristically, there will also be a region around that spot where if you touch around it, that cell will be inhibited.

So that's center-surround organization on the skin.

Do I have all sorts of pretty colored chalk here to play with?

Well, no, I have a color. So plus in the center, minus in the surround, and if you poke around out here, it doesn't care at all.

So that's a good way to localize stimulation on the skin.

In the visual system, you've got the same thing.

So now, this can be the retina with lots and lots of photoreceptors paving its surface. If you find some cell in the-- well, I suppose a nice place to look for center-surround cells would be the output cells of the retina, the ganglion cells of the retina-- you might find that that output cell-- so here's a cell hooked up to the retina and going off to the brain somewhere, and that would be ganglion cell.

And it's hooked up in such a way that if you stimulate, let's say, any of these photoreceptors, this cell would become more active.

But if you stimulate any of the ones immediately surrounding it, it would become less active.

These photoreceptors would have an inhibitory impact on this cell.

So it's got a center that's excitatory and a surround that's inhibitory.

So that means if you have light that just fills that center region, that's going to be the optimal stimulus for this cell.

It's going to be really happy to see a little spot, and that can be the basis for-- a set of things like this then become the basis for building more elaborate features later on in the central nervous system.

So if you go up into the brain-- so now we'll go to visual cortex, back here at the back of the brain.

That's the usual picture.

So if you go back here, you'd find cells that are organized, again, in a sort of a center-surround way, but now, it's elongated.

So their best stimulus would be a bar that just has this particular orientation and lies in one particular part of space.

And you could imagine building more and more elaborate things out of-- go from little photodetectors, to spots, to lines, and so on.

It's a little simplistic, it turns out, to think that you're just building up what you're seeing pixel by pixel.

But that's the thing that these receptive fields are doing for you.

Yeah.

STUDENT: So you're saying it starts very big when you're on the retina, and as you move up towards the brain, these little basic parts become more-- PROFESSOR: Right, if you look in the retina of a human or a monkey, you find cells that are interested in spots, and they also have a certain amount of interest in color.

And that's about it.

They don't care about motion, they don't care about orientation, none of that stuff.

Go up here into the brain, they start to care about orientation, and motion, and more sensitivity to size.

And their color processing gets fancier.

And then, you may recall that there are two broad paths coming out of visual cortex, a where path heading up into the parietal lobe.

And for present purposes, the more important one would be the what path heading down into the temporal lobe, where you get fancier and fancier cells-- cells with preferences for more and more elaborate stimuli.

So we've known for 30 years that there are cells down here in a monkey's brain that like things like monkey hands, and monkey faces, and stuff like that.

What we now know from fMRI studies in alert, awake humans is that there are little chunks of brain down in this pathway that are specialized for faces in humans.

There are little chunks of brain down here that seem to be specialized for places.

They like to see scenes.

They don't care, particularly, about faces.

They don't care about isolated objects.

But they would love a view of a room like this-- so spots, more elaborate features, heading on into ever more complicated stuff.

Does that help? All right, well, we'll let you do follow-up here.

STUDENT: How does [INAUDIBLE]? PROFESSOR: Oh, how do you build yourself a motion detector?

The simplest way to do it-- and almost exactly this is done in some lower organisms, and then variations on it are what's being done in your brain somewhere.

So imagine you have two little photodetectors, and let's have them both be connected to one cell in an excitatory-- well, let's see, how do I want to do this?

Now, we'll have it be one of these guys will be excitatory.

So this guy can be excitatory and this guy can be inhibitory.

So if you, just in a story like this, shine light on both of them, nothing much happens.

But what we'll do is we'll put a time delay on here, but just on one line.

So we'll build a little time delay in.

So that way, if you're coming from this direction, you hit this positive one.

A positive signal gets here.

This cell gets excited, and says, I saw something.

But if stuff comes from this direction, you hit this guy, it sets up some inhibition here.

You hit this guy, nothing gets through because the cell's been inhibited recently, and it doesn't see the motion.

So this would be the minimal little detector that would say, I got motion going this way and no motion-- and it likes motion going that way and not motion going the other way.

STUDENT: [INAUDIBLE].

PROFESSOR: It's in the brain.

It's in your brain.

If you were a rabbit, it would be-- or a frog-- it would be in your retina.

But in mammals, like in primates, it's up in the brain.

There's not much motion preference.

Cells in the retina of humans, near as we can tell, and certainly of monkeys, don't care about motion.

So you don't want me to give an hour-long lecture about what are known as Reichardt detectors because it's very cool, but it's not going to be a good use of your time here.

But the basic little circuit might be worth your time.

There was a hand.

Oh, there's the hand.

Yes.

STUDENT: Well, if you look ganglion, is that [INAUDIBLE] the word up there?

PROFESSOR: It was ganglion.

It was since erased.

STUDENT: [INAUDIBLE].

PROFESSOR: Well, quick, little retina tour-- so let's have some cone photoreceptors here.

Cone photoreceptors-- this would be, say, in the fovea-- would be connected to bipolar cells.

Bipolar cells, in turn, would be connected to ganglion cells going out to the brain.

So the through route would be photoreceptors here, cones, bipolar, ganglion.

And then, there are two layers of lateral connections.

Up here, you'd have horizontal cells.

They, at least, have the advantage of a name that makes some sense.

And down here, you have a second set of lateral interactions of these so-called amacrine cells. And this means, how are you going to build a center and a surround?

Well, you're going to have these lateral connections building the center and surround.

Otherwise, you just have pixel to cell to output cell.

It's these connections that give you, among other things, the center-surround kind of organization. Yes.

STUDENT: What's the difference between procedural knowledge and [INAUDIBLE] memory and then the clarity of the knowledge [INAUDIBLE]?

PROFESSOR: What's the difference?

STUDENT: Yeah.

PROFESSOR: Is that like I asked what the difference was?

STUDENT: I'm sorry?

PROFESSOR: I'm just wondering, looked like you were quoting to me from a prior exam, and I-- STUDENT: Oh, no.

PROFESSOR: Oh, OK.

So I would, for present purposes-- I mean, you can sit around and argue about subtle distinctions there, but really, the distinction that you want to make-- all right, we'll leave these cute, little vectors because they're pretty, too.

All the rest of this has to go. So you've got explicit versus implicit knowledge.

And explicit is the stuff that you're going to need on Thursday.

It's the stuff that you, the user of the brain, can go and somehow reach into your brain, ask yourself for, and get out.

That's explicit memory.

Implicit memory is a grab-bag term for a bunch of other stuff that's clearly in there, but that you don't have-- well, you don't have particularly explicit access to.

So let's see, under that-- and I'm forgetting all the various terms that you-- what were the other terms on there?

We'll categorize them.

STUDENT: Procedural and declarative knowledge.

PROFESSOR: Oh, procedural, no, declarative-- declarative sounds good-- so motor memories, how to write a bike, how to say a word, are often listed under implicit.

Procedural-- people fight about this thing, But procedural versus declarative is how you do it and how you talk about it, but is used by some people as a stand-in for explicit versus implicit.

The really important distinction-- well, and another version another way to break this apart is this is the stuff for which the hippocampus is vitally important at the encoding stage.

You don't get explicit memories, new explicit memories, without a hippocampus.

You do get a range of implicit memories without a hippocampus.

You learn to see that something looks familiar.

You can learn new associations, classical and operant conditioning.

So let's call this association learning are all things that might be under the general rubric of implicit learning.

So if you take HM, the guy with the hippocampal lesion, and every time-- I don't know, every time you show him a bunny, you do the little-- [VOCALIZES] baby Albert-- baby Albert?

Baby Albert experiment.

Every time you show him a bunny, play a loud, unpleasant noise.

HM would learn that there's something he doesn't like about seeing bunnies.

He wouldn't remember that you'd-- so that would be an implicit association that he had learned, classically conditioned association.

He would not have explicit memory of having learned it.

Does that help?

Are we getting-- OK.

STUDENT: So are they interchangeable?

PROFESSOR: Almost.

The major difference is that I would be likely to use the explicit/implicit language, and I will tend to stay away from the declarative/procedural language because I think it's hard to define.

I mean, I, as the guy who's writing the test.

This is a hint that if you've got a good handle on the distinction that I'm making between explicit and implicit, you'll be fine there.

And I can't quite say, we already asked one.

You already got one.

However, nobody else wants to ask anything.

So go ahead.

STUDENT: Can you talk about the difference between fixed and variable ratio interval in-- PROFESSOR: OK, so schedules of reinforcement in Skinnerian Pigeonland-- so ratio schedules refer to how many of whatever behavior the animal needs to produce in order to get rewarded.

So let's talk about pigeons pecking at a key.

A fixed ratio, 10, schedule of reinforcement would mean the pigeon's got to peck 10 times, exactly 10 times.

And after the 10th peck, the pigeon's going to get whatever it is that the pigeon wants, whatever's on offer in this particular cage.

A variable ratio schedule would be one where after an average of 10 pecks, the beastie gets reinforced.

Sometimes, it's nine, sometimes it's 12.

On average, the bird is going to need to peck 10 times.

So that's ratio.

"Interval" refers to time.

A fixed-interval schedule would be the bird gets fed for the first behavior for the first peck after a minute.

Doesn't matter what else he did during that period of time, but after a minute, if he pecks, he gets some birdseed.

And variable interval would be on average, he's got to wait a minute-- might be 50 seconds, might be 70 seconds, but on average, he's got to wait a minute. So the sorts of things that you want to know are what produces the most robust kinds of behavior, most robust, hard to extinguish, hard to get rid of behavior?

And variable ratio schedules are good for that.

The higher a ratio that you can establish-- in principle, if you could get the bird to work for a variable ratio 1 million schedule, that bird is going to work himself into the ground, but the bird is never going to learn that.

But a variable ratio 10 schedule will produce more behavior and be more resistant to extinction than, say, a fixed ratio of 1.

Why is that?

I think I used the example in class of the distinction between a slot machine and a Coke machine.

Coke machine better be a fixed ratio 1 reinforcer.

You put in your money, out comes the can.

You put in your money, nothing comes out, you don't put in any more money.

The slot machine is a variable-ratio reinforcer and with a long, with a high schedule on it.

And so you put in your money, you don't get anything out, and Las Vegas is busy hoping you'll just keep putting in your money, and getting reinforced on a slow-- and that will extinguish much more slowly.

Yep.

STUDENT: Do you think that one peck, and then all the [INAUDIBLE] would make a pigeon [INAUDIBLE] do it again and again [INAUDIBLE]?

PROFESSOR: No, if you peck once, and all the quarters fall out of the machine, you got all the money you need, at least for now, and you quit.

If you are-- this is why the behaviorist project was interesting in child-rearing land.

You can sit around and ask yourself, if you're a little kid, what's going to make you work harder?

If mommy gives you a cookie every time you do anything, how much work are you going to do?

Not that much.

How many cookies can you eat, anyway?

But if mom starts by giving you cookies for all sorts of stuff, and then starts to say, you've got to do more stuff and more stuff.

And now, I'm not going to give you cookies unless you go to MIT.

Then you work like a lunatic.

You stay up all night.

You see, if your mother had been giving you cookies all the time, you wouldn't be here.

Yes.

STUDENT: Isn't it the same amount of work per cookie?

Isn't the difference we expect [INAUDIBLE]??

PROFESSOR: Oh, is it-- no, certainly not the same amount of work per cookie.

STUDENT: On average the same amount, right?

PROFESSOR: No.

STUDENT: 10 pecks per cookie, not 9 or 11-- same work, really.

It's just-- PROFESSOR: Oh, you're just talking about the difference between fixed ratio and variable ratio.

Yes, that's really the same amount on average.

STUDENT: So when you say variable is more effective, you mean [INAUDIBLE]?

PROFESSOR: The uncertainty.

The uncertainty seems to produce both higher rates of performance and more resistance to-- look, if you had conscious access to this, and you're on a fixed ratio 10 schedule, and you emit 10 pecks-- peck, peck, peck, peck, peck, didn't get anything, you stop.

But if you're on a variable ratio, you, huh, maybe this is an 11.

Oh, man, maybe it's 12.

Maybe it's exponentially distributed.

Maybe this one's come out 150 pecks later.

Anyway.

STUDENT: [INAUDIBLE] never any distinction, you always get a-- for the fixed or you always get it for the variable, which has a sharper learning curve?

PROFESSOR: Oh, sharper, well, in both cases, you're going to need-- you need to shape the behavior.

If you simply put a bird in a cage, and say, OK, 10 pecks before you get any bird food, you're not going to get much behavior out of the bird.

You've got to start with feeding them the cookie every time, and then slowly work them back to every 10.

The one thing that I think I can safely say is that the variable ratio-- and you can sit around and wonder about why this exactly is-- but variable ratios produce higher rates of emitting behavior than fixed for-- there's buckets of long, boring, theoretical books on distinctions like that.

And if you go on to the interval schedules, what the birds does or the rat does is show evidence of being able to time, particularly for the fixed interval ones.

For a fixed interval one, right after he gets reinforced, the pigeon simply doesn't do anything because he knows those pecks are valueless.

And then, he starts to think, have I waited long enough, and starts to peck.

And the rate of pecking will go up, up, up, up, blip, until you reinforce it.

And then it'll drop.

So you get this very scallop-y rate of behavior, whoop, boom, whoop boom.

That's how you know you've got a pigeon who's on an interval schedule.

There was a hand up.

There's a hand.

STUDENT: Can you talk about contingency and contiguity?

PROFESSOR: Can I-- who?

Contingency, was that contingency?

And contiguity, I bet, was the other.

STUDENT: Yes.

PROFESSOR: Yeah, so the important punch line there is that it's contingency that counts and not contiguity in association learning.

What the animal is learning is what predicts what.

So if you just go to, say, Pavlov's dog, if there's bells ringing randomly and food showing up randomly, even though the bell and the food may be contiguous to each other in time-- they may show up a lot at more or less the same time-- the animal never learns anything.

What they learn is if the bell shows up reliably before the food, and does not show up reliably before not food-- so if it's perfect contingency, it's bell, food, bell, food, bell, food, every time.

So that would be 100% contingency.

If you have 50% of the time after the bell rings, food shows up, the animal will learn that association, too.

It'll learn it more slowly and it will produce less saliva in this particular example.

But they'll still learn it.

The amount of learning is related to the contingency and not the contiguity.

So I can make all sorts of silly examples, but does that sufficiently define those?

STUDENT: What is contiguity?

PROFESSOR: Contiguity is just proximity in time or space.

So there are lights on.

Every time you have ever flunked a test, the lights have been on.

Yes, so?

That's not interesting.

Even though light and bad outcome or bad tests, those were contiguous events, you don't learn anything about that.

If every time I turned on a red light, you proceeded to lose 10 points on the exam, eventually you would come to-- even if you hadn't particularly specifically noticed this-- eventually, you would come to have a queasy feeling about red lights because there now be a contingent.

I don't know how I've done this contingent relationship, but it's red light syndrome, or something.

But now, there'd be a relationship between red light and bad things happening, and you would learn that. Ooh, there's a hand.

STUDENT: Could you explain feature integration theory? PROFESSOR: Oh, Lord, I could explain feature integration theory for the next 12 years because-- but I won't.

So the basic idea of the reason I can explain it for 12 years is because my major contribution to the literature is a follow-on to Ann Treisman's feature integration theory.

But the basic core idea of feature integration theory is that-- where'd my visual system go-- early in the visual system you have a bunch of simple features that are extracted from the visual image.

They would be things like the stuff we were talking about-- motion, size, orientation, color, and a bunch of other ones.

And Ann noted with interest-- Treisman noted with interest-- that there seemed to be specific chunks of the brain that were particularly interested in, say, motion.

And now, you can see this in fMRI.

It's here-ish on this picture.

There's a little chunk of brain that loves motion.

There's another chunk of brain that loves color and stuff like that.

But when you look at the world, you don't see that.

You don't see colors, and motions, and orientations.

You see objects that have associated color, orientation, motion, whatever.

How do you integrate those features?

And Treisman proposed that a critical role for selective attention-- that you couldn't do this without putting attention onto-- without selecting an object with your attentional processes.

She originally thought, actually, that somehow, the world was a soup of features beforehand, that they were completely just swimming around.

And then, you finally got your attention on something and bound them together.

But what seems to be the case, certainly, is that you can't recognize an object, you can't figure out that this orientation goes with this color goes with this motion, and so on, until you attend to that object.

And with luck, you all remember that I did a bunch of cool demos in one of those lectures, illustrating the sorts of problems you have keeping track of stuff outside of attention.

Oh, and since we met on this subject, Dan Simons, whose name I think I mentioned in class-- he's the guy-- I talked about the experiment where you're standing on a street in Ithaca, and somebody walks by.

You're giving directions to somebody, and these two guys walk through with-- yeah, yeah, oh yeah.

Well, that's Dan Simons.

And he also did the famous experiment where you don't notice a guy in a gorilla suit and stuff.

Anyway, he won an Ignoble Prize.

You know about the Ignoble Prizes?

Yeah, yeah, well, he won one.

That was good.

It's better to win a regular old Nobel Prize, but if you're not going to win one of those, an Ignoble is at least kind of fun.

So yes.

STUDENT: Also, do you need to be able to identify some of these structures of the brain [INAUDIBLE]??

PROFESSOR: Do I need to?

Well, no, I don't particularly, but you might.

Sure.

Yeah, this is the unendingly impossible question to answer which is, which bits do I actually need to know of all that blankety blank neuroanatomy?

Because the short answer is, you need to know the important stuff.

You don't need to know every little detail.

You just need to know the important stuff.

The problem is that I get to define what the important stuff is.

And anyway, very useful to know-- I certainly wouldn't want to not know where the four main lobes of the brain, of the cortex, were-- frontal lobe, parietal lobe up there.

This is a big fold, the central sulcus, big fold in the middle.

Occipital lobe at the back, temporal lobe down here.

I kind of want to know that visual cortex was here, somatosensory cortex up here on one side of the central sulcus, motor cortex on the other side.

I'd want to know bits and pieces of the neuron.

And then, I wouldn't go off and stay up all night worrying too much about it.

Because the worst thing that's going to happen is that it's going to turn out that I think that it's vitally important that where the whoosie foosie nucleus is, and you thought that was just so much trivia.

Well, there's 5 points shot.

That's sad, but it's not going to be the worst disaster.

If you can memorize every term in the book, great.

If you lose a few of them, it won't be a huge disaster.

Yeah, James.

STUDENT: What's the difference between the somatosensory cortex and the [INAUDIBLE]?

PROFESSOR: The somatosensory cortex is a sensory cortex.

It's there responding.

If I poke you, some chunk of somatosensory cortex lights up corresponding with the chunk of skin that I have poked.

If you then haul off and punch me, that's because some chunk of motor cortex has activated the relevant muscles that says, arm, pow, like that.

So motor is there for producing voluntary motor output, and this is for receiving sensory input.

That's the basic distinction.

STUDENT: Is that located on the left and the right?

PROFESSOR: Yeah, they're close neighbors to each other, probably not accidentally. And there's this nice map of the skin surface basically laid out over the somatosensory cortex and of the body, particularly the musculature of the body, laid out across motor cortex.

And they have a rough correspondence, with feet up here in both cases and head down here, as I recall.

But basically, the important thing is sensory input, motor output.

That would be the tidbit to remember there.

Yeah.

STUDENT: Can you explain [INAUDIBLE]??

PROFESSOR: No.

I mean, I could, but I failed to talk about it in mainstream at all.

Did I talk about signal detection theory in mainstream?

No, and I don't think I talked about it significantly in concourse, either.

It's a fascinating topic.

It's an endlessly complicated topic.

But because I didn't say enough about it, it's in the book, right?

STUDENT: Yeah.

PROFESSOR: Yeah.

All right, you get another benefit from having shown up to the review session.

Because I didn't get around to lecturing about it and because I know it confuses people, it ain't on the midterm.

Maybe it'll be on the final, just because I know it confuses people, but it's not on the midterm.

Good.

STUDENT: [INAUDIBLE]

PROFESSOR: Ah, he wasn't listening.

Should we tell him?

STUDENT: Signal detection theory.

PROFESSOR: Yeah, all right-- signal detection theory-- do you want to know about it.

If you're a course 6 major, you're probably going to take some course on it eventually and stuff, but wonderful stuff.

But I don't want to try giving a full blast lecture on it.

However, had I remembered that it was in the book, I would have answered that.

But never mind.

Too late.

Yes.

STUDENT: Can you go over Weber's law and Fechner's law?

PROFESSOR: Yeah, Vay-ber's law, since he's a good German guy-- Weber's law is basically the idea that comparisons in sensory systems actually broadly are made in terms of ratios and not in terms of absolute value.

So if I can tell the difference-- I'm going to make these numbers up, these are not the real numbers-- but if I can tell the difference between 1 pound and 1 pound 1 ounce.

Well, that's a lot.

Let's do it in metric land.

It's easier-- 1 kilogram and 1.1 kilograms, let's suppose I can tell the difference between those two weights.

And if it's less than that, I can't tell.

If I have 10 kilograms, I would need 11 kilograms to tell the difference, not 10.1.

So it's the ratio that's important, not the absolute value.

That's Weber's law.

And the Weber fraction, if the book talks about Weber fractions, the Weber fraction is-- well, in the example I just gave, the Weber fraction would be 1 over 10, and 1/10 would be the difference that you need in order to be able to tell the difference.

And it would differ for different-- all sorts of different discriminations would produce-- and now, no, you don't need to know the particular fraction.

Is there a table of fraction in [? Glickman? ?] Lovely.

You can memorize it.

It's wonderful for you.

No, no, you want to-- it's again, one of these things where if it turns out that I think some cool detail like that is really neat and I ask about it, you're not going to flunk the exam on that.

But no, I wouldn't go memorizing the Weber fraction.

But now, the Fechner's law basically says that sensation is log stimulus, log energy in some fashion.

So the brightness of a light-- and this gives you ratio laws.

The brightness of a light doesn't go up-- if you double the number of photons hitting you, don't double the apparent brightness.

The apparent brightness goes up with the log of the number of photons that are hitting you. In fact, does the book does Steven's law, too?

I forget.

STUDENT: No.

PROFESSOR: So Weber's law is almost right, but it's not quite right.

The story actually seems to be a little closer to-- you got the basic idea of Fechner's law here?

So Weber's law is ratio.

Fechner's law is the extension of that to saying that sensation goes up with log energy.

What really happens is sensation goes up with energy to some power.

It's a power law rather than a strictly logarithmic law.

They're very close for-- the predictions of those two are very close for many things.

But the interesting thing about having it be a power law is that if you do something like light, the exponent is around 0.33 or something, gives you this compressive function.

But if it's a power law, that gives you the possibility of a linear relationship if the exponent is 1, and an accelerating relationship if the exponent is greater than 1.

So can anybody think of a sensory domain where the exponent ought to be about 1? How about length?

If you have a 6-inch hunk of wood and a 12-inch hunk of wood, it would be really lame if the apparent length of an object, over short distances particularly, went up with log length.

Things would look really weird.

And that doesn't happen.

And then there's the entertaining cases of things where the sensation rises much more rapidly than the energy, the stimulus.

Where do you think you get these accelerating functions?

STUDENT: Pain.

PROFESSOR: Pain, yeah, great experiments.

[LAUGHTER]

They hurt to think about.

So the exponent for electric shock to the teeth turns out to be about 3.5, which is to say a tenfold increase in voltage gives, you what is that, a 50,000 times increase in perceived pain.

It's great stuff.

So watch what experiment you sign up for.

Oh, that reminds me to say, I know you're in the midst of studying for this midterm, but my lab begged me to say we really, really, really need subjects.

And we're not doing pain studies.

We'll pay you $10 to work out further details of Ann Treisman's feature integration theory.

Who knows, you might get an extra 5 points on the exam from just the experience, not because we'll give you extra credit.

But anyway, if you want to be a subject, send us email.

We'd love to have you.

Now, with that advertisement done, yes.

STUDENT: What's the difference between [INAUDIBLE]

rods and cones in your eyes?

PROFESSOR: Sure. Difference between rods and cones-- rods are working in dim illumination.

They are concentrated.

Let's make-- well, I'm going to not make that the fovea.

If this is the fovea, the rods are concentrated away from the fovea.

They are actually absent in the central fovea altogether.

And there's only one type of rod photoreceptor in terms of their photopigments.

So there's only one rod photopigment, with the result that you cannot see wavelength differences.

So you can't do color with just the rod.

So if you go outside tonight, when we're done here, and you're in dim illumination, the world will seem to be achromatic, not because something has changed about the physics of the world, but because you're looking at it with a single photopigment.

That's the so-called problem of univariance, by the way.

Cones operate in brighter light.

There are three different types of them in terms of their photopigments, one more sensitive to long, one to medium, and one to shorter wavelengths.

That gives you the possibility of color vision.

And they are most densely concentrated at the fovea, which is why you have in daylight vision-- your best acuity is when you foveate an item, when you point your fovea at it.

Closest packing of photoreceptors is here, but they're cone photoreceptors.

If you are a color-blind male or a color-deficient male, that's because rather than having three photopigments, you either have typically two photopigments or three photopigments, but two of them are so similar that they don't do any work for you, any differential color vision-type work.

It's very, very rare to be completely colorblind.

That would require that you only have one functional photopigment.

Those people are very rare.

But 8% of the male population-- it's a sex-linked trait-- 8% of the male population has some color deficiency.

It's only half of 1% of the female population.

And in fact, there are females, maybe even those some sitting amongst us tonight, who actually have four photopigments, which turns out to make almost no difference at all.

[LAUGHTER]

Because it turns out that while you've got four photopigments, the rest of the wiring that exploits this is built for a nervous system that comes with three photopigments.

And you lose most of the information.

There are very subtle differences, very subtle differences, between what a quadrachromatic woman and a trichromatic woman will say about color.

But if you want to think that your sophisticated color sense is because you actually have four photopigments, first of all, you are a woman.

And I think-- I think-- well, there's probably some guy out there somewhere-- but basically, you'd better be a woman.

And more power to you.

Yes.

STUDENT: Can you lose them?

PROFESSOR: Can you lose them?

Well, you'd have to be extremely clumsy.

[LAUGHTER]

Can you lose your photoreceptors?

Sure, there are all sorts of things you can do to lose your photoreceptors.

But the things you can actively do, if you're so inclined, are to do what your mother said not to do and stare at the sun for a while, which actually can destroy photoreceptors.

But there are a number of disorders, mostly genetic disorders rather than acquired disorders, that munch up photoreceptors, often in young adulthood, and lead to progressive blindness.

But basically, the answer is nah, you're fine.

If nothing traumatically bad happens to your eye and you don't happen to be a carrier for one of these diseases, your photoreceptors are there.

They'll stay there.

You won't get new ones.

You get new taste receptors because Nature knew that you were too dumb to wait for the pizza to get cool.

If you did same thing to your eyeballs that you do to your tongue, you'd have to be able to regenerate those receptors, too.

You're always doing stupid things to your tongue, and so those receptors do regenerate.

Olfactory receptors turn over but if you destroy them, they don't regenerate.

Auditory hair cells, the receptors in the ear, very active game is to try to figure out-- somebody, I think, over at Mass Eye and Ear now thinks they have a way to regenerate hair cells.

Because they don't regenerate.

If you lose them, the way-- oh, the way to lose-- photoreceptors, they're kind of hard to lose.

Receptors in the ear, that you can lose.

The way to lose that is crank those speakers.

Get those speakers up really loud because the music sounds really good if it hurts, right?

And if it hurts-- if you've ever had the experience, if you've ever played your music really loud, and then you hear a very pure, single, very pure tone, that's a hair cell in your inner ear saying, "Goodbye." [LAUGHTER]

"And I'm not coming back." If you hear that a lot, it's not good.

You want to turn the speaker down.

It used to be that the green line on the T was the place to hear that.

Back when I first came to town, they had these cars that-- you know the big turn from Boylston into Park Street Station?

It's a right-angle turn.

The cars going through there would produce this immensely loud, high-pitched scream.

And that was the one time where I heard the beautiful tone of a hair cell dying.

And ever after that I would go through Boylston Street like this.

I don't care that everybody else on the train thinks I look like a doofus because I'm going to have my hair cells when I'm 70 and they're-- anyway, they got new trains and it doesn't sound as bad anymore.

Is that a hand?

You want me to stop digressing and talk about something meaningful here?

STUDENT: What about drinking ethanol?

How does that make-- PROFESSOR: What about drinking ethanol?

That's not a good idea, either. There are all sorts of things that are neurotoxic in the world, and sufficient quantities of-- ethanol is what's in regular alcohol, right?

STUDENT: Yeah.

PROFESSOR: It's ethyl alcohol that's-- there's an expert out there somewhere-- ethyl alcohol is the stuff that's really bad for you.

It's a potent neurotoxin.

Yeah, you want to poison yourself, there are ways to do it.

And a of poisons have their primary effect on nervous tissues of various sorts, including receptors.

And actually, it was a problem with antibiotics.

There's a whole set of antibiotics that had the unfortunate effect of killing hair cells.

And so it cured you of the infection that was going to kill you, but you were deaf.

And obviously, they don't give those a lot at this point.

But that was a complication of a variety.

Which antibiotics is that?

I can't remember.

But yep, well, yeah, you've already asked one.

You haven't, though.

STUDENT: Can you talk about the difference between anterograde and retrograde amnesia?

PROFESSOR: Anterograde amnesia and retrograde amnesia-- all right, so I take this poor, young woman, and I hit her over the head.

Pow.

if she cannot remember what happened before that hit, which she probably wouldn't be able to, that's a retrograde amnesia, from before the point of injury.

Anterograde is from after the point of injury.

So typically, the hitting example is a good one.

Typically, if somebody gets knocked unconscious, they will not remember whatever it is that knocked them unconscious.

I mean, they may know what it-- they won't remember the point of impact or whatever that knocked them unconscious because that memory got lost in a retrograde amnesia.

And there may be some period of time after they've regained consciousness where they have no memory subsequently because basically, probably, because the hippocampus wasn't working properly and was failing to lay down new explicit memories during that period of time.

So that's the distinction.

STUDENT: Can you do that one more time?

[INAUDIBLE]

PROFESSOR: That's the distinction.

[LAUGHTER]

No, that wasn't it, huh?

Retrograde before the injury, anterograde after.

Right?

STUDENT: An-ter-o-grade?

PROFESSOR: An-TER-o-grade or AN-ter-o-grade-- I think I've seen it spelled both ways and pronounced both ways.

All right, how about way over there?

STUDENT: [INAUDIBLE] the book [INAUDIBLE]

a lot of times [INAUDIBLE]?

How does that work?

PROFESSOR: Oh, if you have a traumatic brain injury of some sort, it's like other injuries.

Stuff swells.

It's not working all that well.

And over time, as it heals itself, if you haven't actually destroyed tissue, you regain access to whatever it was that you didn't have access to.

Exactly how that works-- I mean, apart from saying things like the swelling goes down-- I know I can't give you a better account of the molecular level details.

I'm not actually sure anybody knows exactly what it means that you recover it.

And that recovery period can take a long, long time.

So after a stroke, it looks-- I take it-- well, my understanding is that it looks exponential.

If you have a stroke or some brain damage, you get most of what you're going to get pretty quickly, and then you get slow recovery.

It's not a nice, fixed function, the thing that makes doctors humble or ought to. Dr. So-and-So says, I'm really sorry, but this patient is simply never going to recover function.

And 90 times out of 100, the doctor is probably right, but the hundredth time, when this person comes out of a coma and talks to you again or whatever, that's what makes doctors humble.

STUDENT: Are these injuries localized in the brain?

Do they know-- PROFESSOR: It can be.

I mean, I would do it-- well-- STUDENT: [INAUDIBLE] specific amnesias and-- PROFESSOR: So how well is memory localized?

How well are specific memories localized in the brain?

Memories-- almost any reasonably complex memory is distributed fairly widely.

It was once upon a time thought that the whole brain was sort of a uniform mush, and it was just distributed in a giant net.

That doesn't seem to be true.

But only very, very precise, carefully trained memories in animals, to my knowledge, have ever been localized down to a cellular kind of a level.

But even in humans, you can get lesions, brain damage, in quite specific areas that produce very specific so-called agnosias.

Agnosia, a failure to know-- well, we'll put an agnosia in the frontal-- no, I better not put it in the brain because somebody's going to think it's a piece of brain.

That's going to be bad.

And then they're going to blame me when they get it wrong.

So we'll just erase something here. So agnosias are failures to know something.

And they can be quite specific.

You can have patients who now fail to name animals.

They're otherwise reasonably intact, but they have a specific deficit in naming animals.

If you have a lesion-- do I still have my little face area here-- if you have a specific lesion in the face area, you end up with a patient with a so-called prosopagnosia-- I'll just write the whole word, same agnosia on the end here-- where you have a specific problem recognizing faces.

So you can have specific losses, due to what boils down to memory with specific lesions.

But if you ask, where is my memory of MIT, it will turn out not to be localized here.

It's spread out all over the place because your memory of MIT is this very multifaceted thing.

It's got visual components, auditory components, emotional components, taste components, and so on.

And that will be quite widely distributed, even if it feels kind of unified when you ask about it.

Yeah.

STUDENT: What's the main function of the hippocampus?

PROFESSOR: The main function of the hippocampus appears to be enabling the consolidation of short-term, fragile memories into long-term, explicit memories. That's this beautiful-- I've erased half of it.

Not these guys, not these implicit memories-- it's not critical for that.

But for explicit memories, no hippocampus, no new explicit memories.

It is not the place where the memories are stored.

So if I go in and remove James's hippocampus-- and I would have to do it bilaterally, by the way.

So hippocampus, almost any structure I've talked about, is represented on both sides of the brain, both cerebral hemispheres.

So if I take the hippocampus out on both sides, James will still remember he's James.

James will still remember all sorts of stuff.

James just won't learn anything new, which could be a little awkward.

The great exception, there are-- actually, we're learning more and more about functions that are lateralized, that aren't represented on both sides of the brain.

But the great classic example is language.

If you are a right-handed individual, the primary organs, if you like, of language lie in the left hemisphere.

And if you're a left-handed individual, it's about 50/50 where language is located.

Just out of curiosity, how many of you are lefties?

Oh, it's pretty scarce.

I once upon a time had a notion that left-handedness was overrepresented at MIT, but it's not overrepresented here today.

The poor left-- so somebody, some poor left hander in this class wanted to know about left-handed desks in 10-250.

Was that you?

But you were a woman last time.

[LAUGHTER]

I remember that.

That actually puts you in a much smaller category of people who can switch that back and forth at will.

Anyway.

STUDENT: Actually, left handed and homosexuality, there isn't.

PROFESSOR: Is there a relationship?

STUDENT: No one knows me, but-- [LAUGHTER]

PROFESSOR: The late Norm Geschwind had a theory that there was a constellation of things that went together in males, not in females.

Homosexuality was not one of them, but it was left-handedness, certain brands of intellectual talent, autoimmune disease, which wasn't too great.

But it was a constellation of things that he thought had probably a single or common genetic path in there somewhere.

But we digress again.

Let's see, it's the same people.

You didn't get one in before or did you?

I don't remember.

Go ahead.

Whatever.

STUDENT: [INAUDIBLE] parallel distributed processing [INAUDIBLE]?

PROFESSOR: Parallel distributed processing-- when you go off and do AI and/or computer-- the Course 6 guys will know lots about this eventually, if they don't already-- but for present purposes, it's the notion that you can do complicated cognitive or mental things with simple, little units if you have them put together in a great, big, distributed network.

The details of exactly how to put this together are probably not desperately important for present purposes.

The bits that are important from your question, for present purposes, are semantic priming, a good probe for implicit memory.

So semantic priming are any a variety of ways of showing-- oh, let's see.

Let's do an example that I didn't do in class, just for the fun of it. If I say, what is the first word that comes to mind if you're completing this?

STUDENT: Monday.

STUDENT: Monkey.

PROFESSOR: Monday, monkey, money.

Those were the three I could think of off the top of my head.

That's why I needed to put something up that had more than one completion.

Mononucleosis-- there's a bunch.

There's a lot of things, but there are some that are more common than others.

An example of semantic priming-- so I can figure out which-- so let's suppose that-- whoops-- that the top winner was "Monday," and that "monkey" comes in second, if I went and I was running down a list of animals-- giraffe, lion, gorilla-- and then I put up this, you're much more likely to now complete with "monkey" than with "Monday." You're still going to do Monday?

What?

STUDENT: [INAUDIBLE]

PROFESSOR: Monte what?

STUDENT: [INAUDIBLE] like [INAUDIBLE] something, why is there [INAUDIBLE]?

PROFESSOR: I'm lost.

What?

Something about-- STUDENT: Yes, there were voles.

PROFESSOR: Voles, do they have-- were they mountain-- STUDENT: [INAUDIBLE] and some of them were montane voles.

PROFESSOR: Oh, oh, mon- something voles.

Well, you read too much about the voles.

And, well, in any case, we diverge a little from the topic at hand.

The basic point is that we could bias this by semantically priming you.

It's an example of semantic-- one of the interesting things about how do we know this is implicit, we can do this even if you didn't know anything about the prime.

For example, it turns out to be the-- I tell you about the anesthesia experiments?

No?

Very cool stuff.

I put you under anesthesia.

Well, I can't do this because the human subject committee won't let me, but you might need to go under anesthesia.

So I get a deal with the doctor that I can come in and whisper some words to you while you're underneath, while you're under anesthesia.

(WHISPERING) So one of the words I whisper is, "Monkey." I wake you up later.

Do you remember that I whispered words to you?

No, I don't remember any stupid words.

All right, what's the first thing that comes to mind if I put up these three letters?

Monkey.

So get out of my face.

You actually heard the words while you were anesthetized, stored them, have no explicit memory, but you have an implicit memory of it.

It turns out this came as a big surprise to the surgeons.

And it causes a change in surgical practice.

It's now considered to be bad form to bad-mouth your anesthetized patient.

[LAUGHTER]

Man, that is the ugliest kidney I have ever seen.

Or "I don't think she's going to make it," and those things.

Well, if she doesn't make it, it doesn't turn out to be a big implicit memory thing.

But it did turn out that people were reporting, when I see the surgeon, I feel angry or something, things that you didn't quite understand.

And some of it may be this semantic priming effect-- that the information got in.

It's activating chunks of this semantic network.

That's what the priming piece means.

But you have no explicit access to it.

You just somehow implicitly know that somebody's been dissing your kidneys.

There was one-- oh, Stroop effect.

The Stroop effect is a lovely demo.

Is there a demo of it in the book?

I should have done it as a demo in class. It's taken as, among other things, evidence for automatic processing.

Some things that you do over and over again become so ingrained that it's very hard not to do them.

For most of us, reading is one of those things.

And so the classic Stroop effect is that I write a bunch of words up, and I don't have colored chalk, sadly, but I just say, tell me-- I don't care what the stupid word is, just tell me what color it is.

So white.

OK. So what color is it?

Well, people say white, mostly.

Sometimes they make a mistake, but even people like you or me who know about the Stroop effects now-- this is the Stroop effect-- are slowed, substantially slowed, because the reading of "green" is automatic.

And it competes with the response.

You can't get that "white" response out because this thing is screaming, "It's green." And there's a whole lot of these Stroop effects, a very big literature on it, which mostly is there to tell you that there are processes that happen automatically, that you simply can't-- so one of the more interesting ones is, oh, what's the official term for it?

I'll talk about it later in the term.

But it turns out to be easier to associate young and good, for example, than old and good.

So just as you have an automatic reading ability, you also have these automatic biases in this giant semantic network that makes it-- that show up in Stroop-like effects.

So I'll talk about that later in the term.

It's cool stuff.

Yes.

STUDENT: [INAUDIBLE]? PROFESSOR: Depth cues and how we perceive motion-- well, all right.

Here, let's do this collaboratively. Give me some depth cues.

STUDENT: [INAUDIBLE].

PROFESSOR: No, give me some hands that are going to give me some depth cues.

Yeah.

STUDENT: Stereopsis.

PROFESSOR: Stereopsis, oh, start with the fancy stuff.

So stereopsis, there's the difference in the images on the two retinas because the two retinas are looking at a 3D world from different positions.

That so-called "binocular disparity" is a depth cue that you are extremely sensitive to if you've got two eyes that work together.

Particularly if you were a kid with one eye turned, that's a cue you don't have, typically.

You're so-called stereo blind.

OK, there's one good one.

What else we got?

Yep.

STUDENT: Overlapping PROFESSOR: Who?

STUDENT: Overlapping.

PROFESSOR: Overlapping or occlusion-- yep, if one thing is blocking another thing, it's probably in front of that other thing.

Yep.

STUDENT: Texture.

PROFESSOR: Texture-- well, here, we'll lump together for present purposes texture and size, all of which get to the same basic idea that we know that things get smaller, their retinal images get smaller, as they get further away.

And that is the depth cue.

Yep.

Oh, you're cheating.

You're reading them out of the book.

You're not-- STUDENT: [INAUDIBLE].

PROFESSOR: It's OK, you can cheat now.

If you do it on Thursday, you get in big trouble.

The best in 25 years of teaching Intro was, we're given-- I think it was actually the final, not the midterm, and over in Walker with all the cute, little desks.

And one poor student has her head down on the desk.

And it just looks like she probably didn't get her flu shot because none of us are going to get our flu shots, so we're all-- she's just looking miserable.

And so one of the TAs goes over to see if she's OK.

And she's not only OK, she's got the textbook open on the floor.

She's turning the pages with her foot.

[LAUGHTER]

That wasn't good.

But anyway, what do you got for us there?

Melville there, whatever, you, the green guy.

STUDENT: [INAUDIBLE]

PROFESSOR: Oh, linear perspective, lines going off to a vanishing-- parallel lines going in depth converge at a vanishing point.

So linear perspective, that's a good one.

What else we got?

Yep.

STUDENT: Color, like things far away look-- PROFESSOR: Yeah, haze, aerial perspective is-- I don't know why it's called "aerial perspective," but that's what it's called, haze, basically.

Things far away, particularly in a humid world, are bluer and less distinct.

Any more?

Yeah.

STUDENT: If I move my head back and forth, is that just stereopsis?

PROFESSOR: No, because if I do it with one eye, it'll work just fine, so better not be stereo anything.

But it's more or less the same thing.

That's motion parallax, another good one, motion parallax.

But it's the same basic idea.

So instead of having two eyes in my head that are in slightly different positions, I have one eye that I move to slightly different positions, and the difference is between the image-- the relative motion, in particular, I can read off as a clue to depth.

What else we got, any other good?

Yep?

STUDENT: So can you explain motion parallax because [INAUDIBLE] they're kind of confusing [INAUDIBLE].. PROFESSOR: Well, OK, so if-- this is the great motion cue that all kids discover sometime driving down the interstate highway.

Because if you're driving down the interstate highway, and you're fixating on one cow out in the field there-- oh, I better not look at anybody now because I've called them a cow out in the field.

[LAUGHTER]

We'll look at a blank chair.

That chair is in a fixed position on the retina.

The things in front of the point of fixation will move relatively in the opposite direction, and the things behind the point of fixation will move in the same direction as the viewer.

It's just geometry.

That's why the moon follows you around when you're driving down the highway.

What were you guys doing when you were driving at night?

So maybe it's only kids who are going to turn out to be vision researchers or something.

Anyway, just try it next time you're driving down the highway.

Look out the window, look at the trees.

The distance that you're fixating on is relatively stationary.

Everything in front moves against you.

Everything behind moves with you.

That's a depth cue.

You know about that regularity and you use it to give you a clue to depth.

Convergence and divergence of the eyes, basically a triangulation cue, is a weak cue to depth.

Accommodation, how you're focusing the eyes, is a weak cue to depth in humans.

Anybody think of any other ones that we've forgotten there of the-- yeah.

STUDENT: Light and shadow.

PROFESSOR: Oh, shadows, shadows are good clues to depth.

I showed those nice moony-face figures that had only shadow information.

And you can recover all sorts of information about a face from that.

STUDENT: What were the monocular cues?

PROFESSOR: Monocular ones are the ones that work with one eye.

STUDENT: Oh, that the [INAUDIBLE]..

PROFESSOR: So stereopsis is the only one here is that binocular depth cue.

All the other ones will work just fine with one eye.

And people tend to think that somehow, depth vision is vitally linked to having two eyes.

Having two eyes is a good thing, but the world doesn't go particularly flat if you just look with one eye.

So what are two eyes good for?

I mean, it's clearly not worth it to-- in evolutionary terms, it's very hard to imagine that it was worth it for stereopsis.

It's a lovely side benefit, but what else do you get from having two eyes?

Yeah.

STUDENT: If you lose one eye, you [INAUDIBLE]..

PROFESSOR: Replacement parts, good, very handy.

Probably an explanation for lots of bilateral stuff.

Why two kidneys?

It's just having two of things is a good thing.

What else?

What else do you get out of two eyes that you don't have from one?

Yeah.

STUDENT: Expanded field of vision.

PROFESSOR: Yeah, expanded field of vision.

So with one eye, if I'm looking at you, my visual field stops here.

With the two eyes, it goes out to a little more than 180 degrees.

If I was a rabbit, my visual field would go all the way around my head, it turns out.

They've got full 360 with eyes on the side.

They have lateral eyes.

They can also see straight up, apparently, with ears sticking out.

They don't see with their ears.

They got the ears up there.

But you don't.

Why do you have frontal eyes?

Why not put-- what do you-- what are you getting that the rabbit doesn't have by having frontal eyes?

Yeah.

STUDENT: [INAUDIBLE]. PROFESSOR: The lower the-- yeah, the distinction is-- it's not perfect, but it's prey animals who have the eyes on the side of their heads and predators who have eyes in the front of their heads.

And the logic for prey is clear enough.

I mean, the problem with being prey is that somebody wants to prey on you.

And they do unfair things like sneak up behind you.

So if you can see behind you, you have a certain advantage going for you.

If you've got two detectors looking at the same thing, particularly, for instance, in dim light, you can do better than if you have only one.

So predators have two forward-looking eyes to bring two detectors, basically, to bear on the same thing.

And you get stereo out of the deal, and all sorts of handy stuff like that.

But you should be happy to know that you're high on the food chain, at least by the placement of your eyes.

Now, I don't think you can go further with this and say that the people with narrowly placed eyes are way high up on the food chain.

[LAUGHTER]

I don't think that works. The question once upon a time also said something about motion, right?

Who asked that question?

What was the motion part of it?

STUDENT: I asked about motion parallax.

PROFESSOR: You asked about motion parallax, but the original question here was depth perception and motion.

How does motion work?

Stuff moves.

And the important thing about it-- well, there are a couple of important things to remember about motion.

And I can't remember which of them are actually discussed extensively in the book.

Well, I remember one bit that's discussed in the book.

One thing is that it's another case where you have to worry about ambiguities.

If you've got-- oh, yes, this is discussed in the book.

I just remembered.

Oh, I just remembered all sorts of things which I shouldn't remember.

Anyway, if you've got a line, and it's moving, so you're looking at this line.

So you've got a little-- imagine this is the receptive field of a cell.

It's behaving like a little aperture, basically, and this thing is moving.

Well, maybe it's moving like that, or maybe it's moving like this, or maybe it's moving like this.

If all you can see is this little bit of it, you can't tell which direction it's moving.

I mean, you can narrow it down to within 180 degrees.

It's not going that way, but you can't tell if this is a contour sliding up like that, sliding down like that, moving across like that.

You see the problem?

It's known as the "aperture problem." The example that is in the book is the so-called barber-pole illusion, where you've got-- suppose you've got a situation like that.

Oops, I need to make this exactly the same.

Here, let's see if we can manage to make them exactly the same.

[VOCALIZES]

Yeah, OK, they're all more or less the same orientation.

OK. And it is going to be a messy picture.

Oh, well.

Anyway, what I'm going to do is I'm going to have all these guys moving perpendicular to their orientation.

But what'll happen-- let's redraw this guy up here-- is that this thing, if these guys are actually moving like that behind a horizontal aperture, it'll look like it's moving along the aperture that way.

And these guys will look like they're moving down.

And the reason is that you're using the little endpoints to disambiguate.

How do we know how this thing is moving?

Well, if this thing is sliding, however it's sliding along, these little endpoints are going to all be moving like this.

And so I'm going to guess that the overall motion is along.

And these little endpoints are all going to be moving down.

So I'm going to guess that it's moving down.

The important point here, the general point-- if this doesn't make sense, go back and look at the example, the barber-pole illusion example that I remember looking at in the book.

The important point is that you're making an inference about motion.

Because little local bits of motion are inherently ambiguous.

You have to disambiguate it.

The other thing to remember about motion is that nothing actually needs to move.

The nothing needs to move part is the reason you can see movies.

And cartoons are the clearest example of this, but movies in general you know are a series of discrete frames.

And if I show you this, this, this, this, this, that's called "apparent motion." What you will see is a nice, smooth motion as you basically interpolate the motion in between.

In fact, if I flash a light here and a light here, you will see that light, you will swear up and down if I get the timing right, that you see the light here move through space to here, even though there's no-- there's never any light at this point.

You will see your inference in that case.

I just heard a nice talk on Friday by a woman whose fame is based in part on discovering that if you start making these sort of pictures with human bodies, that the system knows how human bodies work.

And so let's see, can I make a motion that works?

If I go from here-- oh, what's a good motion that doesn't work?

If I do this to something like that, the shortest path from here to here would be some weird path through my upper arm.

[VOCALIZES] But the more reasonable thing would be a swooping thing like that.

And you'll actually see the one that makes biomechanical sense is what she discovered, that you know how bodies move.

You do it with something where you don't know it, it'll go for the shortest path.

But if you use a human body where you know a lot about it, it will go for the shortest path. How are we doing here?

I see another hand.

Yes, that would be yours.

STUDENT: Can you explain overshadowing in conditioning?

PROFESSOR: Overshadowing-- so overshadowing is a conditioning paradigm where-- so if I ring a little bell, and I feed you, ring a bell, feed you, ring a bell, feed you, you'll salivate when I ring the bell.

That's just the Pavlov's dog thing.

If I have a compound event-- bright light, little bell, food, bright light, little bell, food-- even though that's the same little bell that we had before, now you'll only learn the association between the light and the food because the light has overshadowed the bell.

You're noticing the most salient relationship.

You don't necessarily notice the important one.

Now, an interesting example of this, it suddenly occurs to me is we talked about taste aversions, where if you eat something and then throw up, you won't eat that thing again.

But the definition of what that thing is is subject to great overshadowing effects.

And the great example is alcohol.

The sensory effects of alcohol are minimal.

It doesn't taste like much.

So if you have vodka mixed with orange juice, drink enough of it to get sick, what you will develop an aversion to is the orange juice because it's the salient sensory stimulus.

And the alcohol will be overshadowed.

So that's a nice practical example of overshadowing.

So who is sitting here monitoring the game on their something?

What's the score?

STUDENT: [INAUDIBLE] 4, 3, [INAUDIBLE]..

PROFESSOR: Oh, OK. Oh, well.

STUDENT: [INAUDIBLE].

PROFESSOR: Well, that's not good either, right?

4, 3, Yankees, bottom of the eighth. Yes.

STUDENT: [INAUDIBLE] difference between blocking and [INAUDIBLE]?

PROFESSOR: Oh, now blocking is-- I know the answer, and how am I going to not manage to confuse the issue?

If I go bell, tone, bell, tone, bell, tone, bell, tone, I establish the connection between bell and tone.

And now I introduce a light-- light and bell, tone, light and bell, tone-- you don't learn the light association.

It's been blocked by the preexisting connection to the tone.

So the critical difference is the way it's built up, that in the case of overshadowing, these two things show up at the same time for the first time.

And the salient one overshadows the second one.

In the blocking case, what you already know keeps you from learning something new.

Yes, that work?

OK. STUDENT: How can memories be distorted?

PROFESSOR: How can memories be what?

STUDENT: Distorted.

PROFESSOR: Oh, distorted-- all sorts of ways, many of which you might discover on the midterm.

So memories get distorted by things like your preexisting biases.

If you think something-- so remember that part of the way that things get into memory is by making association with stuff that's already in there.

That's the idea that you don't need to rehearse material, for instance, if you've got some way of chunking it automatically and getting it in there based on what you already know.

So if you already know something that's biased, or bogus, or something like that, that can distort a memory.

I can distort memory after the fact by giving you new information that interferes with it in some fashion.

So that's the example from the Elizabeth Loftus experiments of you see a car run a yield sign.

You're told later that it was a stop sign.

And you subsequently come to think it was a stop sign, even if that information is marked as false in some fashion. How you're asked about it-- that's another Loftus example-- how you're asked about a memory can distort the memory.

If I ask you, "Did the car smash into the truck," you'll have that car moving faster than if I say, "Did that car bump into the truck?" I know I should be able to generate a whole bunch of other great examples here.

Well, I mean, there are ways to distort it by knocking you on the head and things like that.

I suppose I could come up with a collection of physiological techniques for distorting memory.

What am I forgetting?

Yes.

Am I forgetting that I dropped the stick?

No?

STUDENT: I didn't say anything.

PROFESSOR: No, you didn't say anything.

You just waved your-- I thought you were bidding on that one.

Who can think of any other great memory distortion?

Oh, there are other distortions of memory.

I don't know if it's less active, but people forget the source of-- the source of memory comes disentangled from the actual memory.

So the one example of that is so-called cryptomnesia, where you think an idea was yours and it was really somebody else's.

Yeah.

STUDENT: You had a really bad day and someone asks [INAUDIBLE].

PROFESSOR: Oh, yes, so affective-- your affective state will influence it.

That's a special case of the bias example.

But if you're in a bad mood, what you recover tends to be different than if you were in a good mood. This is a beautiful example of the phenomenon that will undoubtedly afflict some of you at some point during the exam.

I can see my handout where it listed about five ways that memories can be-- five problems with retrieval or something like that.

And I can't quite read them off of that image.

So I've got this feeling I'm forgetting a couple of them in there.

STUDENT: Wasn't there one when a trusted parent will say, like, remember when you were little?

PROFESSOR: Oh, yeah, so you can have completely false memories that are generated by-- in the example that you're giving is yet another clever Liz Loftus experiment, where what you're doing is somebody, your brother, says to you, do you remember being lost in the shopping mall, you recover a memory that you think feels like a real memory but it's not a real memory.

It's a memory based on your schema-- a good term from one of the chapters in there-- your schema of what being lost in a mall would be like.

This is how most cognitive science sorts understand alien possession narratives.

There are some people who firmly believe that there are guys coming down from outer space, and taking people up into their spaceship, and probing them with interesting instruments, and stuff like that.

And maybe there are.

But more the mainstream cognitive science thought on this would be that this is an example of a false memory.

And the reason there's a degree of consistency there is that people know the story.

They know from-- I don't know how much supermarket tabloid viewing you do, but every so often, there's a picture of an alien coming down to talk to the current president.

I remember a particularly nice picture of an alien talking to Bill Clinton.

And they always look the same.

They've got the smooth, oblong kind of head with the big eyes, and stuff like that.

And people know flying saucers, smooth kind of aliens.

They poke you with stuff and probably do weird sexual kinds of things for whatever their weird alien purposes are.

And people firmly come to believe that they have had that experience.

If it ever turns out that aliens do look like that, we may have to revise the cognitive science take on this.

But in the meantime, it's a common shared story to all appearances.

Well, all right, that's a bunch of good ways that you can distort memory.

And there's that great handout there somewhere to give you the rest of them if we forgot one of them.

Think it's time to go home?

Is the game over? STUDENT: So can you explain drive reduction theory?

PROFESSOR: Drive reduction theory is a general class of motivation theories that say you've got a set of drives, some of them innate, basic drives, like hunger.

Hunger is unpleasant.

Why do you do food kind of things?

You do it to reduce your hunger drive.

It feels bad.

Sexual deprivation feels unpleasant.

Why do you engage in mating behavior?

Well, because you want to get rid of that.

You want to reduce that drive.

The counterexample to it is perhaps more interesting than-- your drive reduction theory, you get to a certain distance, but the things that it doesn't explain is why you would go and eat something when you weren't hungry, or why you might go out hunting for food-- if you're a leopard or whatever-- before you get hungry.

And there, you get the notion of appetitive behaviors, behaviors that are rewarding even though they are not reducing a drive.

In the case of appetitive hunting behavior, the logic is that you want to get the leopard out there hunting before he's really hungry because otherwise, he's going to run out of gas before he gets to catch the gazelle or something.

He's got to be out there while he's still got enough food reserves that he can go and-- it's a way of organizing more complicated behaviors.

So drive reduction is a good place to start.

And then, if you're really going to start explaining why complex organisms do things, you have to go beyond that to Yankee reduction theory.

Anything happening here?

We don't know.

No, no, no, it's still 4, 3 bottom of the eighth.

OK.

All right, well, everybody can make a run for it and catch the bottom of the ninth.

If you suddenly remember more questions, you can send email, post them to the website.

Good luck.

We'll see you.

I won't see you, but I'll hear all about it.