

HOLLY BRANIGAN

Is the microphone working? No. Oh, here we go. How's that? Excellent. Good. Well, welcome, everybody. Thank you for coming this evening. And I'm delighted that you are able to join us for this very special occasion. Professor Jennifer Culbertson inaugural Lecture. And it's fantastic to see so many people here. It's a sell out event. I was saying to Jenny, I've never seen so many people in inaugural lecture before. So just to introduce our speaker this evening. Professor Jennifer Culbertson was born in Baltimore. I've been practising this Maryland. Thank you. I did actually check a YouTube video in the USA. Her first degree was in classics and linguistics from Brown University, and she then did an MA and PhD in cognitive Science at Johns Hopkins University. After a postdoc experience at the Department of Brain and Cognitive Sciences at the University of Rochester and a subsequent appointment at George Mason University, she came over the pond to join the University of Edinburgh as a Chancellor's fellow in 2014. Subsequently, she was appointed as lecturer in 2017 and then almost immediately, a year later promoted to reader in 2018 and appointed to the Chair of Experimental Linguistics in 2022. Hence why we're here today. So Professor Culbertson's research focuses on how languages are shaped by the human mind and in particular the tension that we see between the diversity of languages as we see them in the world. And yet what appears to be universal organising principles for language. So she uses experimental methods to investigate how typological patterns emerge, whether adults and child learners of language are biased in favour of common linguistic patterns, and whether language learning outcomes that we observe in the lab in the course of an experiment, mirrored patterns of language change that we see across time, and whether and how linguistic experience impacts language learning. And her contribution in these areas has been substantial. Her experiments, which are imaginative and I imagine highly enjoyable for participants to take part in, have been published in more than 35 papers in linguistics, psychology and cognitive science journals is that they're ingenious studies which include asking children and adults to learn many artificial languages, as well as studies in which people are asked to convey messages using only silent gesture. And when I look at these experiments, as I said, I think people must be excited to take part. And they certainly look interesting. And I'm sure participants ~~coming~~ come out wondering what it was all about, but also hoping to take part again. Chinese Her research has been highly influential among researchers who are interested in the relationship between language and cognition. Her career highlights include winning the Robert Glushko Prize for an outstanding doctoral dissertation in Cognitive Science in 2012, winning the Chancellor's Rising Star Award in 2019 and being elected as a fellow of the Young ~~Young~~ Academy of Europe also in 2019. Clearly a very busy year. The significance of her research program has been recognised in two very substantial grants, including most recently ~~won one~~ won one from the European Research Council, a five year grant looking at syntax shaped by cognition, transforming theories of syntax using laboratory experiments. And she's led and mentored a research team producing a series of exciting results that have moved the field substantially forward. She's also made a substantial contribution to the field in other ways, including as an ongoing member of the Program Committee for the annual Meeting of Cognitive Science Society. And this is since 2016. ~~in a~~ Anybody who's been to that meeting will attest to what a huge undertaking that is. It's a huge confidence and it's a vast amount of work to be involved in that. And also as associate editor of Glossy-Glossa Psycholinguistics since 2021. But her contribution also goes beyond her discipline. So I've said about how she has studied children and how they learn language, but children have also learnt from her. So she's reached around 600 schoolchildren in the Edinburgh area so far with her interactive language evolves workshop. So her impact is not just in academia but goes out into broader society and hopefully enticing children to get interested in these questions about language and the mind. Now, at this point, having heard all about Jenny's amazing achievements, you may wonder how she actually finds time to do anything

else. In fact, she does. She manages to find time to indulge in reading obscure books and gardening, as well as spending time with her daughter, Ada, who unfortunately can't be here tonight ~~with time~~and with her partner Alastair, and visiting her family in the US. And I'm particularly pleased to be able to welcome her parents who've come all the way from Baltimore to join us this evening. And ~~Alan~~Ally, ~~and~~ his mother as well. So please welcome to the stage, Professor Jennifer Culbertson.

~~UNKNOWN~~CULBERTSON

I got so_

SPEAKER 1

~~E~~xcited that I forgot to put a microphone on.

SPEAKER 2

So hold on a second.

SPEAKER 1

Okay is ~~I'm~~it working ~~or~~are people hearing me. Is it on?

SPEAKER 3

And now. Yes.

SPEAKER 1

Okay. ~~Okay~~. So this this talk is scheduled to last for an hour. But my dad told me yesterday that no one wants to hear anyone talk for an hour, So he's going to raise his hand when ~~when~~ he thinks it's I should stop. So who knows what'll happen? We might be going for drinks early. We'll see. Okay, so we all know that languages differ from one another. So I am an English speaker. When I hear a language like Persian or Malagasy or something like that. Despite being a linguist, this is very disappointing to many of my family members. But despite being a linguist, I don't really know anything other than that. Those languages are very different from English. So this map illustrates one particular dimension on which languages can vary. So each dot on the map is a language and the languages are coloured based on what kind of order they use in typical sentences. So in a sentence like Jenny loves bagels, in English, we've got the subject, then the verb, then the object. So English is a language that typically uses subject, verb, object order, but different languages use different orders. So Persian, for example, is a subject object Verb. Hawaiian uses verb, subject object. Malagasy uses verb object subject. There are languages that use object, verb, subject, and other languages that use object subject verb. So if you're paying very close attention, you will have noticed that any way of ordering subject, object and verb is possible in a language. So there will be a language that typically uses that order. So these kinds of maps that show distributions of languages

with a particular kind of feature like this one, they help us to see just how languages can vary from ~~from~~ one to another. So and that's a lot. They can vary a lot, but these maps also show us something else. So just by looking at them, you can see that not all the patterns are equally frequent. So if we count up all the languages that have a particular ~~colour_order~~ like this, we can actually see that two of these patterns, with subject object, verb and subject verb object, are much more common than the rest of them. And one intriguing possibility is that there's something about the human mind that makes these patterns more common, something about them that makes them cognitively preferred. And indeed, there's a rich tradition in linguistics and cognitive science of observing these kinds of skewed distributions and taking those as evidence of the link between the human mind and human languages. So we can imagine, for example, why might it be that ordering sentences using the subject the object than the verb or the subject, the verb and the object? Why might these be cognitively preferred? So you might have some guesses. So one thing they have in common is that they both have the subject first. So those are the two orders that have the subject first and none of the other ones do. So maybe there's something special about having subjects first. What could that be? We'll have a look at some kind of typical sentences of English. Jenny loves bagels. Always a good one. Laura drives tractors, Ada pushes prams, Ally reads books, Ron cooks dinners, just Jess about pets Scott's cats. Lindsay hugs Ruby, typical Typical sentences that one might hear in life. There are equivalents and other languages might look kind of similar. And maybe there's something you notice about the subjects in all of these sentences. They're humans. And of course, humans are salient and interesting and very important to other humans. And there's reason to believe that salient, interesting, important things ought to come first. So that's a kind of intuitive explanation for why it might be that the human mind prefers these particular orders. So have we solved the mystery? So can we sit back and say we figured it out? Humans are very salient, and important? We we should always come first, including in our language. Okay, We'll know well, no. and And know no for a couple of reasons. So first of all, just the premise itself, that subject object, verb and subject verb, object ~~or are~~ more common ~~or more common~~ than the others, it's more complicated than it might seem. And in particular, counting up languages like this and taking that as evidence for something is itself kind of problematic. And that's because languages are not independent from one another. So take a look at just these points here, this cluster of green dots here. So each one of those dots is counted as having subject verb object order. Those languages have that order. But it's not a coincidence that they all share that order. So these are languages that are related to one another. So they're part of the same language family. And it turns out to be a massive language family. It's very big. So what if some language families are just bigger than others for whatever reason, like guns, germs and steel? It might just be. So if that happens and if that family tends to have, you know, SVO order for whatever reason, then it's going to look like that order is more common, but not for any reason to do with the human mind. I mean, unless word order is much more exciting than ~~than~~ I think it is, I think this would just be an accident of history, right? The languages that happened to spread widely had these orders for whatever reason. So coming back to this map, what if we decided, well, let's count language families instead? So maybe that would be a better representation of what's really going on. But if we did that, ~~we would~~ we would quickly notice that we don't have very much data at all. And actually the relationships between languages are really deep. So they go back very far. And probably if we go back far enough, all languages are related to one another. And this brings up another issue. So consider English. English is an SVO language. It's spoken now in the present day, and English is a Germanic language, so it's related to other Germanic languages via a common ancestral source. So languages like German, Swedish, Dutch, Yiddish are all related to one another. We know that because when we look at them, we can see similarities. So, for example, the words for Mother are similar ~~and in~~ all of these languages. And it turns out that linguists agree, based on the records that we have, that the ancestral language, is

~~these all~~the language. ~~T~~these all came from, had subject object, verb order, and indeed, most of the common day Germanic languages still have that order. So this ancestral Germanic language is related to a bunch of other ancestral languages for other language families. And many of these are also SOV. Some of them have other patterns, but many of them are SOV. And if we go back far enough, these languages are also related by a common historical source. So the language ~~there~~that they come from is sometimes called proto Indo-European. So it's a reconstructed language. We don't have evidence for it. But many linguists argue that it also has subject object, verb order. So the fact that many of these modern day languages have this order is probably just down to the fact that their common ancestor also had that order, but it doesn't actually stop there. So it turns out that some linguists argue that the first human language also had subject object verb order. So the idea is that there was a language. It had this order. The speakers of that language spread all over the world, and they took that order with them. And the majority of languages just haven't changed from that. So it could be that maybe that's why that language is very common. It turns out languages often change from SOV to SVO, so maybe that's why those two are more common. Of course, we could ask, well, why did the first human language have that order? But you know, that could just be a flip of the coin. So maybe that's just a chance. Okay. So if we want to know what explains this distribution of patterns on this map, there's a bunch of potential reasons. So it could reflect something about our minds maybe to do with humans being very salient and therefore best to have them come first in a sentence. Or maybe it reflects accidents of history. So maybe some language families that did very well happen to have SVO and SOV, or maybe it reflects the history of languages themselves. Maybe the first human language had SOV in that sort of spread and stuck around. But these alternatives show us that we can't just look at a nice map with plots and draw the inference that there's a connection between how those plots look and something about our human minds. So if we ~~if we~~ think that there's some feature of cognition that has shaped language, we actually need to show it. And this has been the basic premise of my research since my PhD. So I've been harping on it for many years now, and I figured I might as well continue to harp on it in this talk. So for the rest, ~~for the rest~~ of the talk, I'm going to, to try to ~~to~~ show you that if we see a pattern like this, we can't just theorise that it reflects something about cognition. We actually have to show it. And to do that, we can probe the human mind directly. So in particular, we can subject the human mind to ~~to~~ experiments and see ~~does~~ do humans actually prefer some patterns over others? And can we see a link between that and these kinds of distributions? So the experiments that I'm going to tell you about today actually don't have anything to do with SOV and SVO. That was just a little taster. But they do have to do with the way languages ordered different pieces of different linguistic units. So let's start with pieces of words. So take an English word like happy so we can use that word to form other words. So for example, we can add a prefix like ~~UN~~un to make the word 'unhappy'. We can add a suffix to the word like ~~ness~~ness to make 'happiness', ~~we~~We can add lots of suffixes actually, so we can get 'happy', 'You're happy?', 'happier', 'happiest', 'Yes', 'happily'. And it turns out that in English, most words like this are formed by adding something at the end rather than adding something at the beginning. So take another word like 'walk'. It's easy to think of examples where we modify that by adding something to the end: 'walks', 'walking', 'walker', 'walked'. ~~h~~Harder to come up with an example of where we add something at the beginning. So again, this illustrates a property of English. We tend to use suffixes, we tend to modify words by adding stuff at the end. And it turns out that most languages also work that way. So most languages use suffixes more than prefixes. So here's another map for you. So here's English on that map. We have mostly suffixes, so all the dark blue points have mostly suffixes. The light blue points have more suffixes than prefixes, so a little less strong of a preference, but still there. ~~and~~And the magenta points and the red points ~~of~~have mostly prefixes. So it's a little bit a more subtle pattern, but it's pretty robust. So when you look across languages, if you just count individual languages, you will see that more of them use

prefixes than use suffixes. So as with all of these kinds of skewed distributions, when linguists see this kind of thing, we want to know why is it? And we tend to come up with explanations that have to do with our minds because those are the ones that are interesting to us. So there was a lot of seminal work done on this topic in the 70s, and linguists proposed that the prevalence of suffixes in the world's languages is driven by cognition. So just as we said, that subjects come first in a sentence because they express something important. The idea is that the beginnings of words are actually special. So because of this, we group words together based on their similarities at the beginning. And it's good to have this kind of grammatical information that's conveyed by things like -Ness-ness and -or-er and asked-and-Lee-est, and -ly. It's good to have that at the end rather than at the beginning. So this hypothesis has been more or less the accepted explanation for the prevalence of suffixes ever since this and about 15 years ago, this hypothesis was further supported by a series of fascinating experiments that I'm more or less going to conduct on you live. So in these experiments, participants are played auditory sequences like this. So sequences of syllables that don't mean anything like that'ta te'. And then they're played two other sequences. So one sequence, it's changed at the beginning, but the other sequence is changed at the end table-'ta te bo'. So you can see the first one looks kind of like a prefixed word. The second looks kind of like a suffixed word. So what they're asked to do is say, which one do they think is more similar to the base sequence? So is Tabo-'ta te' more similar to Boti-'bo-ta-te' or Tabo-'ta te bo'? Okay, so if you think that 'ta te'-Is more is more similar to Bo-Tate-'bo ta te'. Please raise your hand. All right. If you think Taytay-'ta te' is more similar to Tatay-Bo-'ta te bo', please raise your hand. Okay, great. Phew. Okay, so indeed, that is what participants in this experiment also also chose. Okay. And it turns out this is a kind of general perceptual bias. So you can see the same exact pattern with sequences of things like shapes or musical notes. So if I give you a base sequence of shapes and then I give you a sequence that's changed at the beginning and a sequence that's changed at the end. Participants in these experiments also found this suffixed sequence to be more similar to the base sequence. Okay. So these results suggest that indeed the prevalence of suffixes across language languages is driven by a general perceptual bias. Something like the beginnings of words are important and we weight them more heavily when we're determining similarity. And so it's better to have word formation processes that involve exploiting these similarities at the beginning and adding suffixes at the end, like happiness, rather than word formation processes that involve similarities at the end, adding things like prefixes to the beginning, like in the word unhappy. Okay, but as with the vast majority of all experiments, these experiments were done with English speakers, and English speakers have a lifetime of experience with suffixes. So many of you in the audience are probably English speakers. You have a wealth of experience with with-suffixes, and if you don't speak English, you probably speak another language that also has a lot of suffixes. So if if the words that participants in these experiments are used to hearing tend to have this oops, tend to have-this pattern of being similar at their beginnings rather than at their ends, then maybe they just transfer that to other kinds of sequential stimuli like sequences of nonsense syllables, sequences of shapes, sequences of musical tones. After all, language is a kind of sequential information that we get lots of. So it's not surprising that it would affect our perceptions. Okay. But as we saw on the map, not all languages are predominantly prefixing. So what would happen if we look at speakers of a language that has more prefixes than suffixes? So an example of such a language is DhakaKĩtharaka. So this is a language in the Bantu family that's the big family of languages I showed you in Africa before that, that big cluster of green dots. This is one of those languages. And in this language, prefixes are all over the place. So, for example, if you wanted to describe these these-monkeys here, you could say something like 'monkey, beautiful-monkeys, beautiful'. And if you want us-to do that, the word for monkey has a prefix on it, and the word for beautiful has a matching prefix on it. But if you were talking about something else that was beautiful, it would have a different prefix. Okay, so these words share a

similarity at the ending. They differ at the beginning. That's exactly the same. I mean, that's exactly the opposite of what we see in a language like English. So what would happen if we brought these participants, speakers of this language into the lab and we asked them to do the same similarity judgement task that we asked English speakers to do. Will They they gave give us the same kinds of judgements or different kinds of judgements? So we ran We ran this study. So I'm going to show you data from individual participants. Each individual participant is going to be a dot on the graph. And if the dot is high up, that means that participants always are tended to choose the suffixed sequence. So they they thought that tat and tat 'ta te' and 'ta te bo' were more similar to one another. If it's down at the bottom, that means that they always chose the prefix sequence. So they thought tat 'ta te' and bowtie 'bo ta te' were more similar to one another. So here's what English speakers look like. So you can see that we just replicated the previous finding and the finding that you all so nicely replicated for me here. So English speakers tend to choose the suffix sequence, but speakers have the opposite preference, right? So they tend to think that the prefix sequence is more similar to the base sequence, and that exactly follows the patterns that these speakers have in their native language. And we can ask whether that is transferred to other kinds of sequential stimuli like shape sequences. So for example, if we run the same study with these sequences of shapes, we can ask, do these participants again show this kind of opposite pattern of preferences? And indeed, English speakers tend to think that the suffixed sequence is more similar. Speakers tend to think that the prefixed sequence is more similar. Okay. So what these results suggest is that in this case, it's language that shapes cognition, not the other way around. So, Kiitharaka Keith, speakers use prefixed words in their language, and these words share similar endings. English speakers are used to suffix words, and these share similar beginnings, and this affects how they perceive sequences of syllables and shapes. So that tells us something actually something quite interesting about how our language affects our perception. But it doesn't offer a possible explanation for why suffixes happen to be more common in the languages of the world. So to explain that, we might just need to look somewhere else. So here we have a case where there's a long-standing hypothesis about the effect of cognition on language, and this hypothesis turns out not to be supported. When we run experiments with diverse enough speakers so that they have the kinds of experiences with language that we need to actually draw a generalisation. So these speakers have different experiences with their native language. They have different perceptual patterns, and that suggests that it's not there's not a kind of universal perceptual bias that's underlying this particular distribution we see on the map here. So we can only do that when we compare when we compare a participants that have different kinds of backgrounds. Okay? So don't worry. It's not all doom and gloom. That experiment doesn't show a connection between between language and cognition. But there's more to come. So we're going to move to a different domain. So here it's again going to be about order, but now about how people order different kinds of information when they're talking about objects in the world. So if I were lucky enough to be standing in front of this tree full of monkeys, I wish I and I wanted to pick out a particular group of the monkeys, then I might say something like 'those two adorable monkeys'. So this phrase has a noun in it, but also three nominal, three types of nominal modifiers that give you additional information about the monkeys. So a demonstrative like 'those', a numeral like 'two' and an adjective like 'adorable'. And in English these words are ordered demonstrative numeral adjective noun: 'Now these two adorable monkeys'. But just like for ordering of elements in a sentence like subject verb and object, different languages use different ways of ordering these elements. So a language like tie Thai has noun adjective, numeral demonstrative, a language like Spanish demonstrative numeral Now an noun adjective, a language like Arabic depending on the dialect numeral now noun adjective demonstrative and so on. There are different patterns that languages can have. Here we've got four elements, and that means 24 different ways of ordering them rather than just six. So I can't show you a pretty map, but I can show you some information

about the frequencies of those different patterns. And you can see that most of these patterns are found as the typical pattern in a language, but it's another skewed distribution. And whenever there's a skewed distribution, again, there's a temptation to come up with ~~for an with an~~ explanation for ~~why~~ why it's skewed. And in this case, linguists have noticed that the most common patterns---actually the eight patterns I showed you ~~in the previous~~ on the previous slide that are highlighted here---have a similar structure. So in particular, all these orders involve putting the adjective closest to the noun, then the numeral, then the demonstrative. So in all of these orders, they all follow this kind of structure where the adjective is close to the noun and the demonstrative is peripheral. So you can think of it this way, like in terms of a schematic, it ~~does~~ not about whether the demonstratives and the numerals ~~in and~~ the adjectives are before or after the noun. It's about how they're ordered relative to one another. So the adjective stays close and the demonstrative farther, ~~farther~~ away. Okay, so but why do languages have this structure? So is ~~it something is~~ there something about this structure that's special to human minds? It could be. Or maybe it's just an accident. So maybe the earliest human languages had one of these orders. So maybe it's just chance. And in fact, if we use a method to control for relatedness, the pattern gets a lot more impressive, at least. So you can see now we have ~~a lots we have~~ a lot less data. It's still the case that these patterns are more common, but it looks much less impressive. So it's worth doing something else to actually check to see again whether these patterns are preferred by human minds. So how can we do this? So we want to know whether there's something special about this structure. So we could see, for example, whether these kinds of patterns, patterns that follow this structure are easier for people to learn. We could even look at how babies ~~are or~~ young kids acquire these patterns in their native language, so that would be something that would be a reasonable thing to do. But it turns out that ~~that~~ these languages are going to differ in lots of ways. So not just their patterns of noun phrase, word order, they'll differ in all sorts of ways and that that makes it basically a confounded experiment to do. So we won't know whether it's something about their noun phrase word order that makes them differentially learnable. It might be some other property of those languages. And besides that, I don't really know why people study babies. It's impossible, so I don't recommend it. So instead, we create languages in the lab and teach them to good old adults. So similar to what I showed you before, we're going to come up with kind of sequences of ~~of~~ nonsense syllables. But this time they're going to actually have meaning. So we're going to pare down the language a lot and see whether these patterns are easier for adults to learn in the lab. So in this case, we're not actually going to see whether they're easier to learn. We're going to do something even simpler or even dumber than that. So in particular, we're going to teach people just a really small part of a language, something equivalent to nouns with adjectives like 'red cup' or 'cup red' and nouns with demonstratives like 'this cup' or 'cup this'. So we're going to teach those kinds of phrases to people. And then we're going to just ask them how do they think the language should order words when more than one of those modifiers is present? So we won't tell them how that's done. We'll just ask them to take a guess. And the idea is that if people have this kind of representation in their minds, then they should use that to take a guess about what the order is in this new language. So, for example, if they learn that the modifiers come after the noun, then they can just read off the structure with the order should be, so ~~now a noun~~ adjective demonstrative because the adjective stays closest to the noun and the demonstrative is peripheral ~~peripheral~~. Okay. So to give you a sense for what it's like to be a participant in one of these studies, I'm now going to see what you do. Okay. So if you're a participant in this study, ~~you're~~ first going to learn about some novel nouns. So you might learn that a cup is called 'uhu'UU.

Excellent. A ball is_.

SPEAKER 1

Called a 'beautiful'. A feather is called Ettie'iti'. What's a feather? Good. You nailed it. Okay, so now let's learn some novel adjectives. So to say, to describe this picture here, you would say, Eat taka'iti taka', you're such sports. Eat poo-poo'iti puku'. How about some demonstratives? Et-je-me et'iti hono?'. Anyone? Get this one?

SPEAKER 5

Oh, that's too hard. It's too hard. It's okay. It's okay.

SPEAKER 1

That's okay. Okay. What about this one? So in this experiment, we only train people on the phrases you just heard. So things like 'eat-or-eatiti puku'? No, but this or 'iti hono'. This picture, if you want to describe it, it requires using both an adjective and a demonstrative. You need to put them both together. And we haven't told participants how to do that, so they just have to take a guess. So in this language, since the modifiers come after the noun, you could either say ate-poo-poo-no'iti puku hono' or you could say ate-jono-puke'iti hono puku'. So how many people prefer 'iti puku hono' eating buku-jono? Dang. And how many people prefer 'iti hono puku' Fuku? Basically nobody. Okay, So yet again, yet again. Thank God you have all validated this experiment, and I barely need to show you the data. So just to be clear, the prediction is that this is the order that you should prefer because it places the adjective closer to the noun and the demonstrative further away. This one should explode your brain. Okay, but now I'm going to show you actual human results. Not you. You're probably bots or something. So first I'm going to show you results from English speakers. Okay. So again, I'll show you dots that represent individual participants. If they're high, if the dots are high up, that means people always those people always produce the common structure. So something like 'iti hono puku', et-hono-noun adjective, demonstrative-, if-if they're down low, that means they always produce the uncommon structure. So 'iti puku hono'. et-. And as you can see, everyone in our experiments is like you-you---almost everyone. And they prefer this common structure. So that's great. This supports our hypothesis. English speakers assume that a new language that has its modifiers after the noun will have this common structure, this structure that we see reflected in most of the world's languages. But before we get too excited, there's another explanation for these results. So just like we saw with the experiments on prefixes and suffixes, English speakers obviously come to this task with a lot of experience with English, and English has already this kind of structure, right? So English doesn't have this same linear order, right? In fact, it doesn't have the same linear order of the modifiers. So it has the opposite linear order of the modifiers, right? So English has demonstrative adjective. This structure that people assume the language has involves adjective demonstrative, but it does have the same structural property. So if people if English speakers in this study are just modelling this language off of this feature of English, this structural distance between between the noun and the demonstrative versus the noun and the adjective, then maybe that explains why everyone in the room here has such a strong preference for that for that particular order.

SPEAKER 6

But just like having-

SPEAKER 1 p

Prefixes, actually having an order like noun demonstrative adjective will not explode your brain. In fact, there are languages that have this order and they obviously do just fine. You don't have any problem. And it turns out that the language I told you about before, [Raqqa-Kiitharaka](#), is a language that has this kind of unusual order. So if you want if you wanted to describe these adorable monkeys, you would say something like this here. I'm not going to try to pronounce it because Patrick would kill me. But it uses the order noun₇ demonstrative numeral adjective. So in this language, demonstratives are more close to nouns. And then the numeral and then the adjective. So here you have the adjective, peripheral and the demonstrative close. So that's exactly the opposite kind of structure that we see in English and most other languages. So this would be the population to test, right? So these, these speakers will have a wealth of experience with this unusual structure. So if even they produce an order in this kind of experiment that accords with the common structure, that would be much, much stronger evidence than we got from English speakers. So just as we train English speakers, that individual modifiers came after the noun. So the opposite [that that of what](#) they have in their language. We train [Kiitharaka keep track of](#) speakers on a language where modifiers come before the noun, and if the structure they have in their minds is something like this₂, with demonstratives closest to now and adjectives farthest away, then we might expect them to produce an order like 'red this cup'. So that would follow that structure with the demonstrative [gloss close](#). If they produce an order like 'this red cup' instead, then that gives us really strong evidence that even those speakers, with all their experience with the other structure, still have this one lurking in their minds. They still assume that this structure will hold of a new language. So just to summarise, we're going to train them on the same novel words that you guys already saw, but the modifiers are going to be before the nouns. So something like [ET or Hono et pɛ 'hono iti' or 'puku iti'](#), and then they're going to produce sentences in the language that need to use both of the modifiers they could produce either [en-ET 'hono puku iti'](#) and that would follow the common structure, or they could produce this other one which we hope will explode their brains.

SPEAKER 6

Okay. But because the order you might have noticed₂.

SPEAKER 1

[L](#)ike the order that we're predicting that they use is actually the same order that's found in English: demonstrative adjective noun. So it's really important that our participants don't know any English, right? Otherwise that would really ruin the experiment. So to do that, we use [Kiitharaka Kitagaw-a](#) speaking participants who haven't had much schooling or who went to school before English education was mandated in Kenya. And that means that our participants are older than participants in most experiments. But that's cool because everything gets better with age. I heard. Okay. Right. So just to remind you, this is the critical thing that we're looking for. So here's this graph again. I'm

showing you the English participants. So remember, if the dots are high up on the plot, that means they produce the common structure for English speakers—, That's-that's noun adjective demonstrative. ~~for~~ For Kĩitharaka speakers that will be demonstrative adjective noun. If they're down low, they always produce the uncommon structure. So here's the results. What you can see is that Kĩitharaka speakers look just like English speakers. So they're not they're not statistically different from one another. Both of these populations assume that this new language is going to have the common structure. Okay, So-so these two experiments together suggest that this is the kind of representation that people have in their minds. This is reflected in the kinds of orders that we prefer, and it's there regardless of the surface word order pattern that people have in their language. And this preference in turn is reflected in the kinds of languages we see in the world today.

SPEAKER 6

Okay.

SPEAKER 1

So you might wonder, like, how is it that cognition actually comes to shape language? So what we've seen is preferences in individual humans. But how does that influence a population-level pattern? So how does that actually make its way into a language? And this happens really via cultural transmission of language. So it happens because languages are passed from person to person. So that happens when we learn a language and that happens when we communicate with each other. And each time this happens is an opportunity for our minds to shape language again. So every time language is transmitted is an opportunity for our minds to shape language. So that's how the things that are in our individual brains get into our languages. And it turns out that these preferences don't have to be super strong, they can be pretty weak and they can actually be amplified via this process of cultural transmission. So it doesn't have to be that every language has this pattern, right? It's possible. It's totally possible. And and many languages do not have this structure, but even a kind of weak preference can be can be amplified via cultural transmission and make its way into the typological distribution, this skewed distribution of languages we see. Okay. So you might also wonder where the heck does this representation come from? So why do we have this one and not some other one? And that is a fair question. And I'm going to share with you a kind of intuition about why I think this is the case, and then I'll provide you some more explicit evidence for it. So if I show you a picture of a banana, it's really easy for you to take a guess about what the colour of the banana is. Right? No problem. Apples. You got it. Wine. Oh, yeah. So these colours are very closely associated with these objects. And the same is true of most objects. So most objects, you do have some expectation about their properties, what they might have, sometimes even definitionally. So. So you expect Dalmatians to have spots, zebras to have stripes. These properties are very closely associated with these objects and that seems to be intuitively true in general. So things that we convey using adjectives like colour and texture are closely associated with objects in the world. What about numbers? So what if I asked you to guess? Like how many bananas should there be?

SPEAKER 4

One, two. Five? I don't know. What about wine?

SPEAKER 1

One one is good. Two. Two is better. Five. Five is sloppy. No, not five. What about Dalmatians? Okay, well, I have an association with ~~With~~ Dalmatians, but I don't know if ~~I don't know if~~ that's a culturally shared one. So 101 Dalmatians. So we maybe have some, but mainly weak associations between particular objects and numerous ~~it-cities~~. What about this kind of thing? So what about the information that we convey in language using ~~Demonstratives demonstratives~~? So languages typically convey things like location or distance relative to the speaker or relative to the addressee. So do you have an intuition that these things are closely connected with the kinds of objects that they describe? Probably not. So almost by definition, these properties are flexible. They're not tied to objects. So maybe we have some expectations that some objects will be closer to us than other objects. Right? Maybe we expect cups that we're talking about to be close to us in trees that we're talking about, to be farther away, but ~~not to not~~ not to particular distances that we convey using words like this or that. So there's not something special about 'this cup' or 'that cup', 'this tree' or 'that tree'. So hopefully you share that intuition. Okay. So the idea is that properties that we convey in language using adjectives are closely associated with the objects that they modify or describe. ~~numerous-Numeroseities~~ are sometimes connected to particular objects, but generally the connection is weaker. ~~and-And Demonstratives demonstratives~~ really aren't closely associated with objects. So almost definitionally, so. So the differences in strength of association that these different properties have with objects match exactly the differences in their distance from the noun. Ooh. So is that is that just a coincidence? That seems like an interesting coincidence. So the first thing we should do is check that that's really true. So is our intuition actually right? And we can measure differences in strength of association to verify that we can do that using a measure called mutual information. So the pointwise mutual information between any two ~~things-things---~~ could be anything out there in the ~~world-world---~~ is given by taking the frequency of that pair of things together, a probability of that pair of things together and dividing it by the frequency of the individual elements. And this gives you a sense for whether those two things co-occur together more than you would expect based on the base rates, their base rates of frequency. So, for example, take bananas and the property yellow so we can get their mutual information by figuring out how often we see yellow bananas and dividing that by how often we see just bananas and how often we see the property yellow. So that will give us the mutual information of yellow banana. So say we look at a thousand pictures just to estimate mutual information. That's how we could do it. And yellow appears in 20 of them and banana appears in ten of them and yellow bananas appear in nine. So there are a few other yellow objects. You know, there's other things that are yellow other than bananas. And maybe one picture has a green banana. But nine out of a thousand is much higher than you would expect based on the base rates of frequency of bananas. So ten out of 1000 and yellow, 20 out of 1000. So that gives you a really high mutual information. So that's just a toy example. But we can do this for real. So sadly, we can't do it by looking at thousands and thousands of pictures. And now I have no doubt that given the non-stop progress of artificial intelligence, we will have this at some point. But for now, what we're going to do instead is use a corpus of language instead. So instead of a corpus of pictures, we'll use a corpus of language and we'll see whether different kinds of nominal modifiers have different mutual information with the nouns they modify. Okay, so let's start with English. So this graph is showing you average mutual information for these different kinds of modifiers. So if the average is high, that indicates that there's a strong association between that modifier and the noun that attends the nouns, it tends to modify. If it's low, that means there's a weak association. So we can see is that our intuition is validated, right? So adjectives have a high mutual information, strong association with ~~the nouns they~~ modify.

~~demonstratives~~ Demonstratives, tend not to have that kind of association, and numerals fall in between. So just some examples. Very high mutual information combinations are things like 'alcoholic beverage', 'dense vegetation', 'seven founders'. So these these combinations of words are very closely associated with one another. Low PMI, low mutual information pairs or things like 'that c-Child', 'New-'new fact', 'one program'. So those are kind of generic, right? There's nothing about the two words that that really connects them together. So importantly, we can also see this in speech to children. So this is data from English child-directed speech, and here we see the same pattern. So everything is lower, mutual information is lower. And that's simply because we say fewer types of different things to children. But the pattern is still there. So that suggests that the kinds of things that we talk to kids about, even those kinds of things, reflect these differences in the strength of association. And that suggests that it's something that kids could learn. So it's something about the world around us and things we talk about that that we could learn early on. And importantly, it doesn't reflect something just about English or what speakers of English talk about. So we can find this in corpora of other languages as well. So this is just a subset of languages we've looked at, but they all show this same pattern. So this is really something about the world and it seems to be directly reflected in languages. So we can think of it as the same kind of schematic that we saw before, except here. Objects are closer to properties. Locations are farther away. So if the patterns we see in language really are driven by how we represent the relationships between objects and their properties-properties---like bananas and colour, Dalmatians and spots---compared to numerous cities and locations, ~~Then~~ then actually we should see these same kinds of preferences, not just in language, but also outside of language in other. In other cases where we need to convey information sequentially, like charades, take charades, you need to decide.

SPEAKER 4

Oops, you-

SPEAKER 1

Wwanted to see that for longer. You need to decide what information you're going to convey and how you're going to order that information. What are you going to, which words are you going to do first? Okay, so let's try it. So please turn to your neighbour and convey this image using only gestures without speaking. I'm just kidding. That would be really hard. But. But imagine doing it with this picture instead. So imagine I swapped you out the feathers for two toothbrushes. Yeah, that's perfect. So you can imagine how you might do this. And we actually do this exact task with people in the lab. It looks a little bit different. So we have participants seated in front of two iPads. One is close to them, one is far. And pictures appear on these two iPads. The pictures of are of either toothbrushes or pencils. They have properties like being spotted or striped. They have numerous cities-. There's there's either 4 or 5 of them. And like I said, they're going to appear either on the ~~the~~ the iPad that is closer or-. ~~The~~ the iPad that is far away. So let's have a look at what this what this participant does. So you can all easily understand what image she's seeing, Right? So it's on the iPad, far away from her. It's a toothbrush. It's striped. There's four of them. Okay, let's get another. It's a pencil.

SPEAKER 4

Spotted. Four of them this iPad. Okay.

SPEAKER 1

So if you are looking carefully, you might have clocked that these gestures follow exactly this kind of schema that that we saw before. So they're different from one another. They're not, they don't, the gestures don't convey this information in the same order, but both of them involve the location peripheral to the rest. So the first order, the first participant gestures something like the location, the object, the property, then the number, ~~the~~ The other participant gestures, the object, then the property, then the number, then the location. So both of these follow the same organising principle. The same common structure that we see in conventional AIS-language is replicated in the kinds of gestures that these people produce. So we can look, for example, across all of the gestures that participants produce in this experiment, and they're coloured by whether they conform to this structure or they don't. And you can see the vast majority of them do. People produce all kinds of different patterns, but they all almost all conform to this pattern.

SPEAKER 4

Okay.

SPEAKER 1

So these results support the claim ~~t~~ That noun phrase word order reflects how we conceptualise and how we organise information about objects. So specifically, it reflects the fact that we represent properties like colour, size, texture --- properties that we convey in language using adjectives adjectives --- more closely together than other properties. And because we represent them close together, we also order them close together in our language. So properties like location relative to the speaker are not very closely associated with objects, and we don't order them very closely to nouns. And numerous ~~cities~~ fall kind of somewhere in between. And languages can conform to this structure in various ways. So there's not just one way to do it. There are eight ways to do it. All of them kind of in principle conform equally well to the structure. And we see all of those orders in languages of the world. You might have noticed if you were looking very closely that two of these orders involve putting all of the modifiers on the same side of the noun. So a language like English has all the modifiers before the noun. That's very common. A language like ~~the~~ Thai that has all the modifiers after that is even more common. So those are the two most common orders. And you could argue that they're most common because they perfectly reflect this underlying structure. So they very closely reflect the underlying structure and we can recover the whole structure from these patterns. But I've actually done a bunch of work using other experiments, and I'm not going to tell you about it today, that suggests that this property of having modifiers all before the noun or all after the noun is just something that people independently prefer. So it's important to note that all of these things are preferences and not requirements, right? So not all languages share this. This feature of having modifiers come before the noun or after the noun. Not all languages share this feature of conforming to this kind of common structure. These are possibilities not these are preferences, not not constraints on what a language can look like. But like I said before, even kind of weak preferences can be amplified over time by the process of cultural transmission. As languages are passed from person to person, they're shaped by these preferences, even weak preferences, and

they show up in these kinds of distributions. So here we see that our cognitive representations for objects which we can probably learn just by looking at the world around us, we can notice these things as children. These shape our language. So it's not always the case that these skewed distributions are driven by cognition. So it might seem obvious now, but we saw that, you know, earlier in the talk I told you about a case where for many years people assumed that these skewed distributions were reflecting something about human cognition. But when we did the right experiments, they don't support that hypothesis. It doesn't mean it's not the case. Maybe there are more experiments out there that one could do, but it shows you the value of doing these kinds of experiments to actually probe what the human mind finds easier to learn, easier to perceive, that kind of thing. So sometimes it's a no, sometimes it's a yes. We don't know until we actually probe human minds. And those human minds need to have diverse linguistic experience. So we wouldn't have had a good answer for either of these cases that I've talked to you about today---for either suffix-ing versus ~~prefix~~-prefixing, for this property that noun phrases have,---we wouldn't have had really the right answers unless we looked at the right speakers. So in this case, we were looking at English speakers and [Kĩĩtharaka](#) speakers. In both of those cases it turned out that English has a kind of common pattern. So suffixes are more common. English has suffixes. ~~having~~-Having adjectives closer to the noun than demonstratives is a common type of pattern. That's what English has. And [Kĩĩtharaka Kitagawa](#) had the unusual patterns, but the same. ~~The~~-the opposite could be true in another case. So English has many unusual properties and has many typical properties. So it really just depends ~~on~~ on the particular patterns that we're looking at, what's the right population of people to test. So when we do get a yes, when we do have evidence from experiments, that language is shaped by cognition, then we're in the great position that we can do further experiments and use further kinds of techniques to really probe what that means. So what is it about cognition that makes those patterns preferred? So is it something we learn? Is this something we're born with? These are all, as they say, empirical questions, and we can better understand both our minds and our languages when we actually ask them. Okay, so that's it. I want to thank my many collaborators that have really helped me over the years to understand better the fascinating ins and outs of noun phrase, word order, and more. And I want to single out in particular Patrick ~~Cannon~~[Kanampiu](#), who's here in the audience, and Alexander Martin, a post-doc of mine from a few years back. So none of this work on [Kĩĩtharaka Ketika](#) would have happened without them, especially Patrick. And also thanks to the CLIE who took me in when when kind of nobody else wanted me and my crazy ideas. So it's been a wonderful place to work all these years. And on a more personal note, thanks ~~to to~~ to Jess and Rob for being my surrogate parents in Scotland. It really breaks my heart that Rob's not here with us today. And to my actual parents, Ron and Laura, who I'm lucky enough to have in the audience with me tonight and who have been through the years in my corner always well, most of the time. And and thanks to to [Allyi](#) and [Ada](#), loves of my life. And thanks to you for listening. Thank you. And no questions will be asked.

[SPEAKER 7](#)

Sorry, Dad.

[SPEAKER 0HOLLY BRANIGAN](#)

Okay. So it just remains for me to say thank you again, everybody, for coming. And thank you for a fantastic lecture.

UNKNOWN

Thank you. Huh? Thank you. Huh?

SPEAKER 5

Did you enjoy it? I was.

SPEAKER 8

Right now, somewhere in the business school.

SPEAKER 5

I feel like I don't know where, but.

SPEAKER 6

Someone will show us, Right?

SPEAKER 4

Just.