LANGUAGE
UNLIMITED

the science behind our
most creative power

DAVID ADGER

DAVID ADGER

LANGUAGE
UNLIMITED

the science behind our
most creative power

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

CONTENTS

PREFACE

My fascination for language appeared when I was about ten years old. I'd been reading Ursula Le Guin's A *Wizard of Earthsea*, still one of my favourite books. In it, our hero Ged, is sent to a windy isolated tower on Roke, an island in the centre of Le Guin's world of Earthsea. The tower is the home of the Master Namer, Kurremkarmerruk, who teaches the core of the magical system of Earthsea: the true names of things. There, Ged learns name after name. Each plant and all its leaves, sepals, and stamens, each animal, and all their scales, feathers, and fangs. Kurremkarmerruk teaches his students that to work magic on something, you need to know the name of not just that thing, but all of its parts and their parts. To enchant the sea, Ged needed to know not just the name of the sea, but also the names of each gully and inlet, each reef and trench, each whirlpool, channel, shallows, and swell, down to the name of the foam that appears momentarily on a wave.

I found this thought fascinating even at the age of ten. I didn't really understand it, because it is paradoxical. How infinitesimal do you need to go before there are no more names? How particular do you need to be? A wave on the sea appears once, for a moment in time, and the foam on that wave is unique and fleeting. No language could have all the words to name every iota of existence. How could a language capture the numberless things and unending possibilities of the world?

I was captivated by this question. And I still am. For although Le Guin's Language of Making is mythical, human language does, in fact, have this almost mystical power. It can describe

the infinite particularity of the world as we perceive it. Language doesn't do this through words, giving a unique name to each individual thing. It does it through sentences, through the power to combine words, what linguists call syntax. Syntax is where the magic happens. It takes the words we use to slice up our reality, and puts them together in infinitely varied ways. It allows me to talk about the foam I saw on a wave, the first one that tickled my bare toes on a beach in Wemyss, in Fife, on my tenth birthday. It gave Le Guin the power to put Kurremkarmerruk's Isolate Tower into the mind's eye of that same ten year old. It both captures the world as it is, and gives us the power to create new worlds.

In this book, I explain how syntax gives language its infinitely creative power. The book is a dip into the sea of the syntax of human language. It is no more than a skimming of the foam on a single wave, but I hope it gives an idea of how important understanding syntax is to the broader project of understanding human language.

<div align="right">David Adger</div>

London
October 2018

1

CREATING LANGUAGE

I want to begin this book by asking you to make up a sentence. It should be more than a few words long. Make one up that, say, spans at least one line on the page. Now go to your favourite search engine and put in the sentence you've made up, in inverted commas, so that the search engine looks for an exact match. Now hit return.

Question: does your sentence exist anywhere else on the internet? I've tried this many times and each time, the answer is no. I'm guessing that that was your experience too. This isn't just a side effect of using the internet either. The British National Corpus is an online collection of texts, some from newspapers, some that have been transcribed from real conversations between people speaking English. There are over 100 million words in this collection. I took the following sentence from the corpus at random, and searched for it again, to see if it appeared elsewhere in the millions of sentences in the corpus. I then did the same on Google.

It's amazing how many people leave out one or more of those essential details.

There are no other examples. It seems crazy, but sentences almost never reoccur.

Think about your sense of familiarity with the sentences you hear or say. None of the sentences I've written so far feel new or strange. You aren't surprised when you read them. You just accept them and get on with it.

This is, if you think about it, quite remarkable. These sentences are new to you, in fact perhaps new to the human race. But they don't seem new.

The fact that sentences hardly reoccur shows us that we use our language in an incredibly rich, flexible, and creative way, while barely noticing that we are doing this. Virtually every sentence we utter is novel. New to ourselves, and, quite often, new to humanity. We come up with phrases and sentences as we need to, and we make them express what we need to express. We do this with incredible ease. We don't think about it, we just do it. We create language throughout our lives, and respond creatively to the language of others.

How can we do this? How can humans, who are finite creatures, with finite experiences, use language over such an apparently limitless range?

This book is an answer to that question. It is an explanation of what it is about human language that allows us to create sentences as we need them, and understand sentences we've never heard before.

The answer has three parts.

The first is that human languages are organized in a special way. This organization is unique, as far as we know, to humans. Sentences look as though they consist of words in a sequence, but that is not how the human mind understands them. We sense, instead, a structure in every sentence of every language. We cannot consciously perceive this structure, but it contours and limits everything we say, and much of what we think. Our sense

of linguistic structure, like our other senses, channels particular aspects of our linguistic experience into our minds.

The second part of the answer is that linguistic structure builds meaning in a hierarchical way. Words cluster together and these clusters have special properties. A simple sentence, like *Lilly bit Anson*, is a complex weave of inaudible, invisible relationships. The words *bit* and *Anson* cluster together, creating a certain meaning. *Lilly* connects to that cluster, adding in a different kind of meaning. Laws of Language, universal to our species, govern the ways that this happens.

The final part of the answer tells us where this special structure comes from, and explains why we can use our languages with such flexibility and creativity. Throughout Nature, when life or matter is organized in a hierarchical way, we see smaller structures echoing the shape of the larger ones that contain them. We find this property of self-similarity everywhere. A fern frond contains within it smaller fronds, almost identical in shape, which in turn contain yet smaller ones. Lightning, when it forks from the sky, branches down to earth over and over, each new fork forming in the same way as higher forks, irrespective of scale. From slime mould to mountain ranges, from narwhale tusks to the spiraling of galaxies, Nature employs the same principle: larger shapes echo the structure of what they contain. I argue, in this book, that human language is also organized in this way. Phrases are built from smaller phrases and sentences from smaller sentences. Self-similarity immediately makes available an unending collection of structures to the speaker of a language. The infinite richness of languages is a side-effect of the simplest way Nature has of organizing hierarchies.

These three ideas, that we have a sense of linguistic structure, that that structure is governed by Laws of Language, and that it

emerges through self-similarity, provide a coherent explanation of creative powers that lie at the heart of human language.

∞

I wrote this book because I think that the three core positions it takes are deep explanations of how language works. Each of these ideas is about how our minds impose structure, of a particular sort, on our experiences of reality.

Over recent years, however, an alternative to these ideas, with an impressive pedigree, has emerged. This alternative focusses not on how the mind imposes structure on our linguistic experiences, but rather on how we humans have very general powerful learning abilities that extract structure from experience.

Language, from this perspective, is like many other aspects of human culture. It is learned from our experiences, not imposed upon them by limits of the mind.

This view goes back to Darwin in his book *The Descent of Man*. The idea is that our minds are powerful processors of the information in our environment, and language is just one kind of information. The way that language works depends totally on what language users have heard or seen throughout their lives. This idea places an emphasis not on the limits of the mind, but on the organization of the world we experience.

These two different perspectives on how the mind encounters the world are both important. This book is intended to show how the first approach is better suited to language in particular.

How would language look, from a perspective where its structure emerges from our experiences?

Language, Darwin said, should be thought of in the same way as all the other mental traits. Darwin gave examples of monkeys using different calls to signify different kinds of danger, and

argued that this was analogous to human language, just more limited. He argued that, since dogs may understand words like 'fetch!' and parrots might articulate 'Pretty Polly', the capacity to understand and imitate words does not distinguish us from other animals. The difference between humans and animals in language, as in everything else, is a matter of degree.

> ...the lower animals differ from man solely in his almost infinitely larger power of associating together the most diversified sounds and ideas; and this obviously depends on the high development of his mental powers.[1]

Darwin believed that humans have rich and complex language because we have highly developed, very flexible, and quite general, intellectual abilities. These allow us to pass on, augment, and refine what we do. They underpin our culture, traditions, religions, and languages. The vast range of diversity we see in culture and language is because our general mental powers are so flexible that they allow huge variation. Darwin argued that this cultural development of language augmented our ability to think and reason.

More concretely, the idea is that we can understand sentences we've never produced because we're powerful learners of patterns in general. We apply that talent to language. We hear sentences as we grow up, and we extract from these certain common themes. For example, we might hear certain words together over and over again, say, *give Mummy the toy*. We store this as a pattern, alongside *give me the banana*. As we develop, we generalize these into more abstract patterns, something like *give SOMEONE SOMETHING*, where the capitalized words stand in for lots of different things that have been heard.[2] Once this general pattern is in place, we can use it to make new sentences. The structure of our language emerges from what we experience of it as we

grow up, combined with very general skills we have to create and generalize patterns. The same skills we'd use in other complex activities, like learning to bake a cake, or tie shoelaces.

Other animals have pattern matching abilities too, but, in Darwin's words, their 'mental powers' are less developed. The reason humans are the only species with syntax, from this viewpoint, is the huge gulf between us and other animals in our ability to generalize patterns. We have more oomph.

This book was written to make the argument that it's not a matter of more oomph, it's a matter of different oomph! We are not powerful pattern learners when it comes to language. We are limited—only really able use one kind of pattern for syntax, a hierarchical one. This is what I'll argue in the first half of the book. I'll also argue that patterns that depend on sequences of words are invisible to us, while syntactic hierarchy is unavailable to other animals. Though we do of course learn our languages as we grow up, what we can learn is constrained. Our limited minds are oblivious to the continuous in language, and to the sequential, and to many possible kinds of patterns that other animals can pick up on. The source of hierarchy in language is not creating patterns, storing them, and generalizing them. It's an inner sense that can't help but impose hierarchical structure, and it's the self-similarity of that structure that creates limitless sentences. That, rather than highly developed mental powers, underpins our incredible ability to use language creatively. That is our different oomph.

$$\infty$$

Unless you're an editor, or a teacher, you probably don't notice the hundreds or thousands of sentences you come across during your day. Most fly by you. In a sense you hear what they mean, without hearing what they are. But sometimes you might come

across someone writing or saying something and think 'That's a bit odd.' Maybe a verb is missing. Maybe the sentence starts but doesn't end. Maybe it doesn't mean what the speaker obviously wanted it to mean. You know certain things about the sentences of your language, though you usually don't stop to think about it.

Here are some examples. Which of them are clearly sentences of English, and which are 'a bit odd'?

> *Zfumkxqviestblwzzulnxdsorjj kwwapotud jjqltu ykualfzgixz, zfna ngu izyqr jgnsougdd.*
>
> *Sunglasses traumatize to likes that water by perplexed usually is tinnitus with amoeba an.*
>
> *An amoeba with tinnitus is usually perplexed by water that likes to traumatize sunglasses.*
>
> *A cat with dental disease is rarely treated by a vet who is unable to cure it.*

If you're a native English speaker—and probably even if you're not—you have probably judged that the first two are not good sentences of English but the latter two are. Of these, the last one is a completely normal English sentence, while the one about the amoeba is weird, but definitely English.

If I give you many more examples of this sort, your judgments about their oddness are likely to agree with mine, and with those of many other native English speakers. Not entirely, of course. There may be words that I don't know that you do, or vice versa. Our dialects might differ in some way. I might allow *the dog needs fed*, while you might think this should be *the dog needs feeding*. You might have learned at school that prepositions are not something that we end sentences with—or not! I might not care about what they taught at school. You might be a copy editor, armed with a red pen to swiftly excise every split infinitive. I might think that split infinitives have been part of English since Chaucer, and be

very happy with phrases like *to swiftly excise every split infinitive*. If we put these minor differences aside, however, we'd agree about most of it and we could agree to disagree about the rest.

How do we all do this? Why do we mostly agree?

Every speaker of every language has a store of linguistic information in their minds that allows them to create and to understand new sentences. Part of that store is a kind of mental dictionary. It grows over our lives, and sometimes shrinks as we forget words. It is a finite list of the basic bits of our language. But that's not enough. We also need something that will allow us to combine words to express ourselves, and to understand those combinations when we hear them.

Linguists call this the mental grammar. It is what is responsible for distinguishing between the first two examples and the latter two. As every speaker grows up, they learn words, but they also develop an ability that allows them to put words together to make sentences of their languages, to understand sentences, and to judge whether certain sentences are unremarkable or odd.

But do we really need a mental grammar? Maybe all we need is the mental dictionary, and we just put words together and figure out the meanings from there. Knowing what the words mean isn't, however, enough. The meanings of sentences depend on more than just the meanings of the words in them. Take a simple example like the following:

The flea bit the woman.

Using exactly the same words we can come up with a quite different meaning.

The woman bit the flea.

How we put words together matters for what a sentence means. Just knowing the meanings of words isn't sufficient. There's something more going on.

These two sentences also show us that how likely one word is to follow another makes no difference to whether we judge a sentence to be English or not. A bit of quick Googling gives about a million results for the phrase 'bit the woman' and just eight results for 'bit the flea'. This makes complete sense of course. We talk more about people being bitten than fleas being bitten. But the likelihood of these two sentences makes no difference as to whether they are both English or not. One is more probable than the other, but they are both perfect English.

The mental grammar can't be reduced to the mental dictionary plus meaning, or frequency. We need both the mental dictionary and the mental grammar to explain how each of us speaks and understands our language(s).

The question of whether we have mental grammars or not isn't really disputed. Whether we think of the human capacity for syntax as emerging from the structure of experience, or from the particular limits of our minds, we still need to say that the general rules of our particular languages are somehow stored in our minds.

But we can use the nature of what our mental grammars must be like to begin to dig into the question of the source of syntax. Is it part of our nature as human beings, or is it something we pick up from the world we experience?

∞

To ask a certain kind of question in English, you use a word like *what, who, where, when*. Take a scenario where someone is chatting away and mentions that my cat, Lilly, had caught something in the garden. I didn't quite hear the full details, so I ask:

What did you say that Lilly had caught?

Here, the word *what* is asking a question about the thing that Lilly caught. Although *what* is pronounced at the start of the sentence, it is really meant at the end. After all, we say *Lilly caught something*.

In many other languages, like Mandarin Chinese, Japanese, or Hindi, to ask a question like this you'd just leave the word for *what* right next to the word for *caught*, giving the equivalent of *You say Lilly caught what?* Here's how this looks in Mandarin Chinese:

Nǐ	shuō	Lìlì	zhuā	shénme
You	say	Lilly	catch	what

The Chinese word *shénme* corresponds to English *what*, and it comes after the verb *zhuā*, which means catch. That's the normal order of words in a Chinese sentence. In a question, nothing changes.

Let's think about how to capture this difference if what we have learned of our language, our mental grammars, develops through noticing and storing patterns from our experiences. Imagine a person, Pat, whose mental grammar grows and is refined over time in this way. Pat learns through noticing, and storing, patterns.

If Pat grew up speaking Mandarin Chinese, they would learn to treat question words no different from non-question words. If exposed to English, they would learn that a question word is placed at the start of the sentence. Pat's mental grammar in this latter case would contain a statement something like this:

> *If you want to ask a question about a thing, a time, a place etc., use a word like* what, when, where, *etc. or a phrase like* which X, *and place this at the start of the sentence.*

Pat doesn't consciously know this, but something about Pat's mind makes them behave according to this pattern. Pat has

unconsciously learned how to make and understand certain questions in English.

The word *that*, as we just saw in the example above, is used in English after words like *say, think, believe*, and so on, to introduce what is said, thought, or believed. When *that* introduces a sentence in this way, it is often optional in English. We see this in sentences like the following:

> *Anita said Lilly had caught a mouse.*
>
> *Anita said* **that** *Lilly had caught a mouse.*

We can put the word *that* in here, or leave it out. Both sentences are perfectly fine ways to express what we mean here.

It's not at all surprising, then, that we can leave out the word *that* when we ask a question too. Both of these next examples are perfectly fine ways of asking the same question:

> *What did Anita say Lilly had caught?*
>
> *What did Anita say* **that** *Lilly had caught?*

How would Pat's mental grammar look if they were an English speaker? They would have learned that the word *that* is optional after the verb *say*, and other verbs like it, so their mental grammar would contain something like this generalization:

> *Optionally put the word* that *after verbs like* say, believe, think ...

So far, so good. Pat's mental grammar contains these two patterns, and many more.

But now let's imagine I had had a different conversation. Imagine the discussion was about one of the neighbourhood cats catching a frog in my garden. If I want to identify the cat, I can ask:

> *Which cat did Anita say had caught a frog?*

A superficial difference between these two questions is whether we are asking about what was being caught, or who did the catching.

Given that the word *that* is optional after *say*, we expect Pat to think that the following sentence should also be fine:

*Which cat did Anita say **that** had caught a frog?*

For most speakers of English, though, this sentence is 'a bit odd'. It is much better without the *that*.

This poses a problem for Pat. They would be led to the wrong conclusion about this sentence. It is a question, using *which cat*, and, as expected, *which cat* occurs at the start. Pat, as we know, has learned a pattern which allows the word *that* to appear as an option after *say*. The sentence matches the pattern: we have taken the option to put in the word *that*. The trouble is that Pat, who is a good pattern learner, would think this sentence is perfectly fine. But most speakers of English think it's not fine—it's decidedly odd. This suggests that most speakers of English, unlike Pat, are not good pattern learners.

This argument doesn't prove that the pattern learning approach is wrong. Real English speakers could be more sophisticated than Pat is.

For example, it could be that children learning English do learn patterns like Pat does, and use those patterns to predict what they will hear. They expect to hear sentences like *What did you say **that** Lilly caught?*. But, the explanation goes, they never do. This means that what they experience doesn't match up with their expectations. The way that the children deal with this is to store an exception to the pattern they have learned. In this scenario, the children's experiences would contain enough structure to help them come to a more complex pattern.

This is an interesting idea, which we can test. In 2013, two linguists, Lisa Pearl and Jon Sprouse, did a careful study of the

speech directed at young children who are acquiring English. They looked at over 11,000 real examples where parents, or other caregivers, speak to their children.[3]

They found that parents, when they asked their children these kinds of questions, almost always dropped the word *that*. They did this whether they were asking a question about what had had something done to it, or what was doing something. It made no difference. The parents never took the option to put *that* after words like *say, believe*, etc. This means that the children didn't ever get the information they would need to learn that there was a difference between the two types of questions.

If we think about this from Pat's perspective, the syntax of English is completely mysterious. Pat's mental grammar consists of patterns they've learned from their experiences. If Pearl and Sprouse are right, Pat couldn't have learned the exception to the pattern that allows *that* to disappear. Pat's experiences, which we are assuming are just the experiences children learning English have, aren't rich enough to learn an exception to the generalization about when *that* appears. Adult English speakers' mental grammars, however, clearly have that exception in them. This seems like a strong argument that English speakers don't work like our imaginary friend Pat. They aren't simply good pattern learners.

Intriguingly, many other languages behave in the same way as English, even though these languages are not related to English or to each other. For example, Jason Kandybowicz studied the Nupe language, spoken in Nigeria, and found exactly the same pattern there. Here's how you say *What did Gana say that Musa cooked?* in Nupe, with a word by word translation:[4]

Ké	*Gana*	*gàn*	*gànán*	*Musa*	*du*	*o?*
What	*Gana*	*say*	*that*	*Musa*	*cook*	*o*

The order of words here is quite similar to English. The little word *o* at the end marks that a question is being asked and the word

gànán is the equivalent of English *that*. Just as in English, it is impossible to say the equivalent of *Who did Gana say that cooked the meat?* You can put the words together, but Nupe speakers don't judge it to be a sentence of Nupe:

Zě	*Gana*	*gàn*	*gànán*	*du*	*nakàn*	*o?*
Who	Gana	say	that	cook	meat	o

There are many other languages that work similarly (Russian, Wolof, French, Arabic, and some Mayan languages).[5]

It is a fascinating puzzle. Speakers end up with judgments about sentences of the languages they speak that don't depend on what they have heard as children. Certain ways of putting words together just aren't right, even though, logically, they should be. And these quite subtle patterns appear in unrelated languages over and over again. We humans seem to be biased against our languages working in perfectly reasonable ways!

There are many puzzles just like this in the syntax of human languages. Languages do have a logic, but that logic is not one that emerges from the patterns of language we experience. The linguist's task is to understand the special logic of language, what laws govern it, and how different languages find different ways to obey those laws. We'll find out in the rest of the book that it's the hierarchical structures that underlie sentences that are responsible for many of these quirks. Some are, without doubt, learned from experience, but others, as we've just seen, are not.

∞

Syntax is a deep source of human creativity. You constantly come across sentences that you've never heard before, but you have no trouble understanding them. My favourite headline of

2017 simply said *Deep in the belly of a gigantic fibreglass triceratops, eight rare bats have made a home.* Beautiful, crazy, and true. Syntax gives us the capacity to describe even the weirdest aspects of our existence, and, of course, allows us to create new worlds of the imagination.

The most basic units of language, words and parts of words, are limited. We can create new ones on the fly, if we need to, but we don't have a distinct word for every aspect of our existence, unlike the wizards of Earthsea. The number of words speakers know is a finite store, a kind of dictionary. We can add words to that store, and we can forget words. But the sentences we can create, or understand, are unlimited in number. There is no store of them.

This book makes the argument that hierarchy and self-similarity underlie our creative use of language. On the way, we'll find out why language is not just communication, how we can sense linguistic structure without being aware of it, and how sentences are like gestures in the mind. We'll meet children who cannot experience the language spoken around them, and so they create new languages for themselves, languages that are taken up by communities and become fully-fledged ways of expressing thoughts. We'll see how human languages follow particular, limited, patterns; how scientists have invented languages that break these; and how they have used these languages to test the limits of the human brain. We'll invent languages to be spoken by imaginary beings, and imagine languages that could never be used. I'll show you how rats can pick up on linguistic structures humans cannot perceive, and how humans can discern ones invisible to our closest evolutionary cousins, the apes. I'll reveal the mysteries of how AIs understand sentences, and how different that is from what we do when we speak and understand language.

We'll also do a little linguistics. You'll learn about some of the Laws that limit how human languages work, and why these Laws can be Universal without being universal. You'll also meet some unusual languages, from Chechen to Gaelic, Korean to Passamaquoddy, and Yoruba to Zinacantán Sign Language. I'll gently introduce you to one of the most cutting-edge ideas in linguistics: Noam Chomsky's proposal that one linguistic rule creates all the innumerable structures of human language. This idea provides a foundation for understanding what underlies our ability to use language in the creative ways we do, but it also leaves open a space for understanding how that use is affected by our social nature, our identity, emotions, and personal style.

2

BEYOND SYMBOLS
AND SIGNALS

In 2011, an internet entrepreneur, Fred Benenson, crowd-sourced a translation of *Moby Dick* into emojis. The word *Emoji* comes from two Japanese words: *e*, meaning picture, and *moji*, meaning a written symbol, like a Chinese character, a hieroglyph, or even a letter of the alphabet. Emojis, then, are intended to be similar to written words: they convey meaning through a written form. Because emojis seem like words, people have talked about their use as the 'fastest growing language'. The initial set of about 180 emojis has grown to over 3,000. Over five billion emojis are used every day on Facebook.

Even more exciting is the idea that emojis are somehow universal. They are pictures, so we can understand them no matter what language we speak. But they are also like words, opening up the idea that emojis could be a universal way of communicating, a language for everyone.

Are emojis like words? When we string them together in our electronic communication is that a universal language?

The linguists Gretchen McCulloch and Lauren Gawne have argued that emojis, as we actually use them, are far more like

gestures than like words. They are a body language for our bod-iless internet selves. The thumbs-up, middle finger, or eye-roll emojis directly represent gestures, but the way we use other emo-jis is also gesture-like. McCulloch points out that we often repeat gestures three or four times to emphasize what we're saying, adding to speech by thumping a fist repeatedly on a table, or opening up our hands, entreatingly, in front of our bodies. The most common sequences of emojis are just repetitions: lots of smiley faces, love hearts, or thumbs-ups. We don't repeat most words in the same way—words have a place in our sentences, and few of them can be repeated without something going wrong.[1]

When we play charades, or watch mime artists, we're using and understanding a kind of pantomime. This, McCulloch argues, is very similar to the ways that you can use strings of emojis to tell stories. This is why Benenson's project was never going to work. It's the equivalent of miming the whole of *Moby Dick*.

Emojis are not really like words then. Though we use them to communicate, that communication is more like what happens with body language.

It's interesting to think about what we'd have to do to emojis to make them work more like words. Perhaps if we enriched emojis, they could work more like a universal language?

Unlike words in spoken or written language, emojis don't express sounds. Expressing sounds, though, is important for even something so simple as someone's name. Benenson's translations of characters from *Moby Dick*, like Ishmael or Queeqeg, are impenetrable. Ishmael becomes a boat, a whale, and an ok sign, signifying what roles he plays in the novel, not the sound of his name.

It is possible, though, to develop emojis so that they could express sounds. For example, you could associate certain emojis

with the sounds of the English words that those emojis make you think of. An emoji for a cat could be used for the syllable *cat*, so you could express *catatonic*, say, by using a cat emoji and an emoji of a gin and tonic. Each emoji would stand for a sound, rather than for what it pictures. This would allow us to express the sounds of names. My name, for example, could be a picture of a sun rising (day) and an old style video cassette (vid).

The alphabetic system that English uses connects written letters to sounds, so it can easily represent how names are pronounced. The Chinese writing system works differently, and is similar to the original intent of emojis. It involves symbols for particular words as opposed to sounds. Because of this, it also faces challenges representing names, especially those that are not native Chinese. However, the users of this system have developed sophisticated ways of writing foreign names by using Chinese characters that have sounds similar to the syllables of the name. I was once given the Chinese name Ai Dao Fu. Surnames in Chinese come first, and usually consist of just one syllable. The Mandarin Chinese word *ài*, which means 'love', is close in sound to the first syllable of my surname (the 'a' in Adger). The words *dào* (meaning 'way', as in Daoism), and *fú* 'happiness' are close in sound, when put together, to David. Chinese has characters for the words 'love', 'way', and 'happiness', so you can use these characters with their associated sounds to write something that is pronounced a bit like my name: Ai Dao Fu—with some lovely meanings to go with the sounds.

A bit more abstract than this would be to use the cat emoji as a kind of shorthand for the sound *k*—often written in English as a *c*—that appears at the start of the word. Doing this connects the symbol to a sound and that's how many of us learned the alphabet. 'A' is for apple, 'B' is for book, 'C' is for cat, and 'D' is for

dog. This basic idea has appeared again and again in the history of writing systems. Pictures which are initially used to represent ideas end up being used to represent sounds.

Ancient Egyptian hieroglyphs worked like this. The word for 'mouth' in that language was pronounced something like *re*, and it could be written using a picture of a mouth:

This hieroglyph is actually usually used to convey the sound *r*. For example, the Ancient Egyptian god Ra, the sun god, was written as the sound *r* above another hieroglyph that was used for a sound that comes out a bit like what happens when you try to cough and swallow at the same time—linguists write this, in the international phonetic alphabet like this ʕ, and the Ancient Egyptian word for 'arm' started with it:
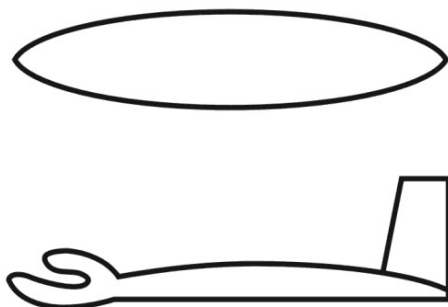
Adopting this idea would allow us to use emojis to write sounds. We could use a cat emoji for the *k* sound, an arm emoji for an *a* sound, and a cup of tea emoji for a *t* sound. We could then express the word for a furry purring animal as follows.[2]

It's rather hard to see how this would be an improvement on just texting though!

Much of the early hype around emojis was about how they were universal. Anyone who spoke any language would be able to understand them. This is certainly an exciting idea, but no symbol is truly universal to humankind.

A symbol is just some kind of a mark made on the world that stands in for something else, usually an idea in your head. This means that there are two parts to a symbol. There's a concept, something inside your mind. This is the meaning of the symbol. There's also something that is external to you, something which you can see (like an emoji), hear (like a spoken word), or feel (like Braille letters). This is called the form of the symbol. So a symbol is a connection between a mind-internal meaning and a mind-external form, between something abstract, and something concrete.

When the mind external part of the symbol, the part you can see, hear, or feel, resembles the symbol's content, then there's a direct psychological link between the two. In this case, the symbol is a bit like a computer icon, say one for a wastepaper basket. Symbols like this are called iconic symbols. Many emojis are iconic, like the ones we just saw for cat, arm, and tea.

There are also symbols without that direct link of resemblance—the relationship between the content and its

expression can be quite abstract or even arbitrary. A love heart is like this. A love heart doesn't look much like a real heart, and the association between the emotion of love and an internal organ is, at best, indirect.

Could we build a universal language built out of iconic symbols like emojis? Since they'd be iconic, people should understand them regardless of what language they speak or culture they come from.

In the early 1990s the US Government commissioned a report on nuclear waste. It had the rather dry title 'Expert judgment on markers to deter inadvertent human intrusion into the waste isolation pilot plant'. A team of experts was set up to figure out how to communicate to unknown people in the far distant future, that a particular plot of land in New Mexico was going to be dangerously radioactive for many millennia. The standard symbol for nuclear waste may not be recognizable in millennia. There may be many radical environmental or cultural changes for humanity.

Various ideas were considered as possible symbols, including 'menacing earthworks', 'forbidding blocks', 'horrifying facial icons' like Munch's *The Scream*. Carl Sagan, the astronomer, physicist, and novelist suggested a skull and crossbones. That didn't fly. The team reported that 'The lineage of the skull and crossbones . . . leads back to medieval alchemists, for whom the skull represented Adam's skull and the crossed bones the cross that promised resurrection. It is almost certainly a Western cultural artefact'. The fundamental problem is that all of the symbols of danger that the team could come up with simply might not mean danger to an unknown population in the future. As the report says,

> No symbol is certain to stay in use for the 10,000 year period. Future societies will probably create many of their own

symbols, and symbols from our time may have their meanings changed or distorted with the passage of time. Compare how the meaning of the swastika has changed in our own century, going from positive religious symbol of India to a hated emblem of the Nazis.[3]

The basic idea that emojis could be truly universal, then, could never get off the ground. Human symbols are always, in the end, deeply connected to our cultures.

Words are the crème de la crème of arbitrary symbols. Aside from a few cases, like animal noises—did you know that the Mandarin Chinese word for 'cat' is *māo?*—they are associated with their meanings through a socially agreed convention. They don't resemble them in any psychological way. This is why the word for 'dog' is *dog* in English, *txakur* in Basque, and *inja* in Xhosa—the same concept expressed by quite different sounds.

We could, then, just as we do with words of spoken languages, or the signs in sign languages, link emojis to meaning using social conventions. The resemblance relationship between an emoji and its meaning would then be useful in guessing a meaning, but the meaning itself would be fixed by communities of emoji users. In fact, such conventions have arisen already through internet users interacting with each other. Sanjaya Wijeratne, while researching his PhD at Wright State University, discovered that gang members were using a gas-pump emoji in their tweets to signify marijuana. Other researchers have found that the meaning of emojis changes across cultures. In some cultures the handwave emoji is just a sign off, in others it's a snub.[4]

Emojis then could be developed to work more like words, though, if McCulloch's gesture idea is right, it's intriguing that that has not been what has happened naturally. Such an emoji language wouldn't, however, be universal.

There is more to a language than just words, though.

If someone texted you the stream of emojis you see here, what would it mean?
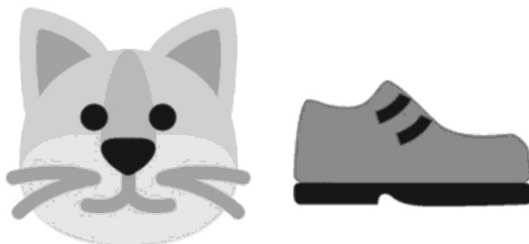
Does it mean a cat is kicking something? Or someone is kicking a cat? Or is it about the story of *Puss in Boots?* Or maybe your friend wants you to get a pair of boots with a cat on them? And how would you even go about clarifying which of these you meant? In spoken or written English (or Cantonese, or Swahili), it's easy to express what you mean with a fair level of precision—in fact, I just did. When you are using emojis, the context might make the message clear. Perhaps you've already been talking about one of these topics with your friend. But in the absence of context, emojis are far too vague to work like a language.

Or think of this the other way around, in terms of expressing yourself, rather than understanding what someone else is trying to convey. How would you express, in emojis, that something has happened in the past? Or the thought that, if some-thing were to happen, so would something else? Or that something didn't happen? Or how would you express that every cat was kicked, not lots of cats, every cat. These concepts, so easily expressed in a few words using a language like English or any other human language, are completely beyond the capacity of emojis, at least without changing what emojis are: a simple connection of a picture and an idea, obvious to everyone when they see it.

The failure of emojis to express past time, events not happening, possibilities of events taking place, quantifying objects, and hundreds of other purely grammatical ideas, gives us a clue to why emojis are different from a natural human language, even if we let emojis include arbitrary symbols. Emojis do communicate ideas using symbols, but human language goes beyond symbols and, as we will see, beyond communication.

To see this, let's go back to our cat and boot emojis. Is the cat kicking or walking? Or is it being kicked? Let's add one more emoji:

We've got a cat, a boot, and a boy. What message is being expressed?

If you speak a language like English, you might be tempted to assume that the order of the emojis is linked to the order of the corresponding English words 'cat' 'kick' 'boy' in the sentence *The cat kicked the boy.* This would give you the meaning that the cat kicked the boy. But isn't it more likely that the boy kicked the cat—after all, boys wear footwear, but cats generally don't? That would be a more sensible and likely message, so maybe you should ignore the order and just go for what is the most probable message that's being communicated.

But maybe you speak a language like Malagasy. The order of words in Malagasy is quite different to that in English. In Malagasy you'd say something like 'kicked the cat the boy', to express that the boy kicked the cat. The person doing the kicking comes last in the Malagasy sentence. This might tempt a speaker of Malagasy