

Cambridge University Press
978-1-107-09674-5 — Introducing Syntax
Olaf Koenenman , Hedde Zeijlstra
Excerpt
[More Information](#)

We humans are surrounded by technology. We have machines for almost everything and computers allow us to achieve things that were long deemed impossible. We live in a day and age where dreams can become reality overnight due to technical innovation, and the amount of information that we have access to via a tiny machine in our pockets is simply astounding. There is talk right now about flying people to Mars and growing tomatoes there. Yes, we humans are a smart bunch.

Despite all this, there are some things we are still unable to do, and some machines we simply cannot build. And some of these failures have to do with language. A machine that can translate one language into another language perfectly? No, we don't have it (and please don't insult us by referring to Google Translate). Okay, how about something more modest, like a machine that can for any combinations of words in a single language (say, English) say whether it is a good sentence or not? It is perhaps hard to believe but even that is still out of our reach. Language, as it turns out, is an evasive and slippery creature.

At the same time, it is clear that such a machine, capable of stating for every English sentence whether it is grammatical or not, *does* exist. In fact, we have about 360 million of those machines on our planet. They are called *native speakers of English*. These speakers have at their disposal the knowledge of their mother tongue, English, and this knowledge can generate zillions of distinct combinations of English words and evaluate each of them, whether old or new, as being either a good sentence or not. Probably you have never heard someone say *Syntax is one of the most fascinating topics in linguistic theory*, but if you are a native speaker of English you know immediately that the sentence is grammatically correct (and hopefully after reading this book you will also find it to be correct content-wise). So these native speaker brains can do something that we cannot imitate with any man-made machine. The fact that we cannot mimic such everyday human language behaviour shows us that there is something worthwhile studying. There is something we apparently don't understand yet, namely the structure of English grammar. After all, if we already understood it, we would have no problem building some machine that imitates this behaviour. But, as said, we can't do that. This means we have to study the English language a bit harder. There is no other way.

But why would we want to know what makes some combination of English words a good English sentence? What is so interesting about knowing that *Maria drank some coffee* is good English, but *Maria some coffee drank* or *Maria drank some car* is not? They may just be facts of life. If so, asking these questions about English may sound like a good pastime for someone obsessed by the English language. Or an obsession for building the machine that we mentioned above. However, most theoretical linguists we know are not obsessed by a particular language. In fact, people obsessed by a particular language generally spend their time doing other things than comparing good and bad sentences. And most linguists we know don't really care that much about this language machine either.

They would be very happy doing their work realising that they will never come across such a machine.

So why are we doing this, if it is not some obsessive pastime over a language or a machine? Linguists are in this business for a different reason, namely the desire to understand the human brain itself. By studying human languages, linguists try to figure out not only the rule system – or the grammar – of the languages they study but also to understand the *nature* of these rules. Why are the rules the way they are? Why aren't they another way? If we know the answers to these questions, we think we will get closer to understanding how the human mind is used for thinking, talking, planning, fantasising, etc. and, therefore, ultimately how the human brain works. Linguistics, and therefore syntax, is a cognitive science, the science of the human mind. It is ultimately as much about the brain and the way the human brain works as it is about language. But why do we think that the study of a grammar is going to give us clues about the human brain? And even if that is true, how are we going to do it? Let us answer these questions in turn.

In this book, we are going to show you what the study of the English language can contribute to this bigger enterprise, understanding the human brain. But wait a minute. English is just English, right? And it differs from Japanese, or Swahili, or Russian. So to conclude after decades of research that English is the way it is because of properties of the human brain would be silly. Such a theory would predict that all humans speak English. And this is obviously not the case. Well, not so sure, actually. As stated above, we look at the rule system of English to discover what underlies these rules. To put it differently, what are the *principles* that underlie the rules of English? Other linguists at the same time look at the rule systems of Japanese, Swahili and Russian with the same goal in mind. What we discover is something quite extraordinary: *although the rules for English, Japanese, Swahili and Russian are quite distinct, the principles underlying these rules are the same for all these languages*. English, Japanese, Swahili and Russian look fundamentally different on the surface but when we look underneath that surface we discover that languages are much more alike than we would think by just looking at the surface.

The linguist Vivian Cook compared this kind of language variation with traffic rules. In the US, cars drive on the right-hand side of the road, whereas in Great Britain they drive on the left-hand side. On the surface, therefore, these traffic rules are completely different. However, they have one crucial principle in common: cars drive on one side of the road only, and drivers can't just pick their favourite side. All drivers, irrespective of where they are driving, therefore follow the same principle. Languages work the same: different on the surface but based on the same principles. So yes, in a deep sense all humans speak English, and all humans speak Japanese, Swahili and Russian, as well. Now, here is the point: these principles that languages have in common reveal properties of the human brain. Language is the way it is, and uses the principles that it does, because that is the way our human brain has evolved. The

brain wouldn't allow other ways of doing language. And since English is obviously one of the languages that the human brain can do, studying it will contribute to the discovery of those principles that tell us something about the human brain.


This brings us to the next question: How do we study it? The grammar of English is present somewhere in brain tissue of a native speaker of English. However, just cutting away will not provide you with crucial information about what this grammar looks like. That would at best be a bloody affair. We could alternatively decide to do brain scans. Certainly less bloody, but this procedure has one serious thing in common with just cutting away: you first have to know what to look for. Without some prior understanding of the principles behind human grammars, it is like looking for the proverbial needle in the haystack. The best way to advance our knowledge about our linguistic capacities is to look very closely at the output of the grammar. What linguists do is just that. They 'reverse engineer' the grammar of English: if you can't look inside the machine, then make intelligent guesses about what it should look like inside based on what it can do. If you put coffee beans into a machine and press a button, hot coffee comes out. This means that there must be something in the machine that grinds the beans, heats the water and makes the coffee by combining the water and the ground coffee. You don't have to look inside the machine to know all this.

Doing linguistics works the same way. What is the output of the language machine? Well, it creates both correct and incorrect sentences, for instance, *Maria drank some coffee*, and *Maria some coffee drank*. *Maria drank some coffee* is a correct English sentence. We call that a grammatical sentence. And the sentence *Maria some coffee drank* is incorrect, or ungrammatical. Every native speaker of English can immediately tell you that. What we should do, therefore, is look carefully at grammatical and ungrammatical sentences of English and make educated guesses about how grammar in the brain of a native speaker makes the distinction. What components should grammar contain so that we get exactly the output that we see? Figuring this out is the job of a linguist. You formulate ideas that seem relevant to you, and test these against more (im-) possible sentences. And that is actually rather easy to do, since every native speaker of English knows whether a particular sentence is grammatical or not. Quite often, though, such first guesses about the internal workings of our language machine are not immediately correct, and more often than not you will have to go back to the drawing board. At some point, though, you will have an idea that works, and you will embark on the next problem. In this way, you try to slowly build a theory of English grammar. If all goes well, and we are pretty far ahead already, you end up with a plausible grammar for English. Not just a grammar of the surface rules, mind you, but a grammar that makes explicit which principles underlie these rules.

If you wished, you could then even build a machine that can generate only proper English sentences and never any bad ones. That is not where we will be at the end of this book (and no syntactician has actually ever

got there), but we will take you on the road towards that point. It won't be a highway on which we speed past syntactic discoveries. It will be more like a crooked path that we have to cut for ourselves, and we will occasionally walk into dead ends. The reason for going the scenic route is to show you not just *what* we discovered but also *how* we discovered it. In this way, you will get a feel for syntax as a science and develop a better understanding of why the theory we end up with looks the way it does, and of how one builds theories in general.

Unfortunately, what we cannot do is look at every aspect of English grammar. As the title of this book reveals, we are going to look at syntax. So what is syntax anyway? Basically what we described above: that part of grammar that distinguishes between grammatical and ungrammatical sentences. Sentences are, simply put, combinations of words, and words can be thought of as simple units that mean something and sound a particular way. We know that a sentence can be ungrammatical even if it is put together with English words that are all perfect. If the words are perfect but the combination is not, it shows us that something must have gone wrong in the building process. Syntax is about the rules that guide that building process. If you obey the syntactic rules, you end up with a grammatical sentence, and if you violate a particular rule, the result will be ungrammatical. Now, what is true for words is also true for these multi-word utterances: they sound a particular way and they mean something specific. *Maria underestimates Harold* sounds different from *Adrian loves yellow chairs*, and it means something different from *Harold underestimates Maria*. What a sentence sounds like and means, then, depends on what you combine and how you combine. And this is determined by syntax. This gives us the following picture of the language machine (1), which you can call a *model of grammar*:



Syntax combines words into bigger units, and these units are uttered (by producing sound) and interpreted (by assigning meaning to them), and the focus of this book is on this combinatorial process. However, in the final chapters of this book we will also look at the sound and meaning systems of the grammar: how are sentences expressed and what exactly do they mean? In this way, we can show what syntax does in the overall grammar, and how we can tell syntax apart from sound and meaning.

This syntactic study has two significant outcomes, which we will highlight from the start. First of all, many factors can contribute to the ungrammaticality of a sentence. *Mary have loved Ahmed* is ungrammatical because there is something wrong with *have* (it should be *has*). *Why you have seen her?* is ungrammatical because there is something wrong with *you have* (it should be *have you*). *We love myself* is ungrammatical because there is something wrong with *myself* (it should be *ourselves*). *I want he to leave* is ungrammatical because there is something wrong with *he* (it should be *him*). And we could go on and on. However, what we will discover in this book is something profound. Although these sentences seem to violate different, unrelated rules, they in fact turn out to violate one and the same principle. English syntax can therefore be thought of as a lot of rules but also as the output of a severely restricted number of principles. And it is these principles we care about, remember?

The second outcome has to do with the model of grammar presented above. Note that this model has different components, apart from the words. There is a syntactic component, a sound component and a meaning component. This has an important consequence. If a sentence is ungrammatical, there can in principle be three reasons for it. It can be syntactically ungrammatical ('syntax error', so to speak), it can be uttered in the wrong way ('sound error'), or it could mean something weird or not mean anything ('meaning error'). You cannot know beforehand (that is, before you do some proper studying) what the cause is of ungrammaticality. And once you have done some research, your conclusions may be different from your initial expectations. Now, recall the sentences that we presented at the beginning of this introduction: *Mary drank some car* and *Mary some coffee drank*. Both are bad English sentences. At the end of this book, however, you will understand that the first sentence is not syntactically ungrammatical at all. It has a weird meaning but structurally the sentence is fine. And what about the second sentence, *Mary some coffee drank*? Well, it will turn out to be syntactically correct, too. It is just expressed in the wrong way. If you find this surprising, please realise that we've barely begun.