Stage 1   Introduction

Meaning – What It Is and Where to Find It

> This is the primary commandment humanism has given us: create meaning for a meaningless world.
>
> (Harari 2017: 259)

1.1   How (Not) to Study Meaning

This book is about meaning – not the meaning of certain concepts such as the meaning of life or the meaning of freedom but about what it means for linguistic expressions to *mean* something. It is also a book about *meaning for humans*. The qualification 'human meaning' is not yet common in linguistics book titles but it is increasingly important to add it. We are already experiencing fundamental differences in *how*, and *what kind of*, meaning is approached in computational linguistics and computer science in general on the one hand, and in philosophy of language or sociopragmatics on the other. 'Semantics' sits in the middle and tends to be approached either (i) philosophically, through questions such as 'What are concepts and where can we find them?' and sociopragmatically, through questions such as 'What meaning does this utterance have in this context and what effect does it have on the addressee?'. This is the 'top-down' approach, from ideas and theories to the practice of use. Or it can be approached (ii) starting with databases of language use and distilling meanings, that is 'bottom-up'. Using (ii), we can distil patterns that can also be useful for producing algorithms and training data processors for human–machine or machine–machine interaction, as in deep learning in AI. That is how (ii) may prove to go beyond (i) and apply to other than human ways of carving and conceptualizing reality.

This book is about the first kind of meaning – human meaning, meaning that we want to better understand in order to understand the systems of natural languages we employ in communication and the language(s) in which we convey our thoughts (more about this soon). 'Bottom-up' and 'top-down' methods may meet in the middle but only if they prove to be about the same *kind* of meaning: meaning that humans make out of reality and externalize through language. We don't know if they do meet, and that is why the book is only about the 'top-down' methods and human meaning.

Non-biological intelligence that humans invent may develop our human feelings, thoughts, and ideas as emergent features of data processing, like, for example, in Ian McEwan's (2019) novel *Machines Like Me and People Like You*. But this is fantasy – we simply don't know yet. And until we do, this is to be read as a book about human meaning, in human natural languages and human discourse.

It focuses on language as a means of conveying and expressing meaning, although it will also on occasion touch upon non-linguistic vehicles such as actions, gestures, or employing context and shared assumptions. It will also focus on human communication – on people and linguistic communities (in that the latter are a source of some useful generalizations about meaning), almost totally to the exclusion of animal and human–machine communication. The field delimited by human languages is vast enough without including these as well.

'To mean' is a notoriously vague predicate. For example, it can figure in the following expressions:

(1.1)   Dark spots on the leaves mean a deficiency of iron.

(1.2)   The beep means that you are parking too close to the wall.

(1.3)   'Divertirsi' means 'to have fun'.

As mid-twentieth-century philosopher Paul Grice famously remarked in his seminal paper 'Meaning',[1] these meanings of 'meaning' are of little interest to a student of intentional communication.[2] Likewise, they will be of little interest to us. In (1.1), 'means' stands for a natural sign of a plant disease; in (1.2), it stands for a conventional symbol; and in (1.3), it functions to provide a translation from one natural language (Italian) to another (English). This book is about a much more nuanced connection between a sign that is used to perform the task of meaning something (a sign that an early-twentieth-century linguist, Ferdinand de Saussure, called a 'signifier'[3]) and the thing that is meant ('signified'). It is about meaning that is associated with, and arises from the use of, expressions of natural language – all of them, that is the lexical items that belong to the system of a given language, the sentences that speakers can produce by combining them, and the physical utterances that they can produce using such sentences. As such, it concerns the kind of meaning where the relation between the sign (a word) and the meant content is arbitrary, but the relation between the combinations of words (sentences or their fragments) is much less so. It strives to discuss explanations of this relation that account for the *productivity* of these linguistic devices (the power of the language system to produce expressions and meanings never produced before) and their *systematicity* (the property of being organized into rules, or a system,

---

[1]   See Grice 1989a [1957].

[2]   To reiterate, the book can be read either without the footnotes, as a conceptual journey through meaning – a journey through ideas – or with consulting the footnotes that provide detailed references to sources and recommendations for a more in-depth study. General recommendations for further reading follow at the end of each stage.

[3]   See de Saussure's (1983) [1916] *Course in General Linguistics* that consists of a translation of his lectures delivered at the University of Geneva between 1906 and 1911.

that guide the derivation of meaning). The explanations we are going to consider tend to be *normative* – not in the sense of laying down the rules of how we ought to use language (that would be the long-abandoned prescriptivism) but rather in the descriptive sense of capturing regularities (or norms) of how the language works and how speakers use it. These explanations often have the status of *theories*. So, theories usually contain rules that capture the properties of the devices used for communication in a system of a particular language and/or conversational behaviour of the users of this system. I say 'usually' because this depends on the philosophical stance towards language: on some views, vagaries of language use are foregrounded, while regularities are downplayed. Most of this book will concern theories that do assume some regularity of linguistic behaviour, both in the behaviour of the linguistic devices and their users, and the ensuing normativity of theories. But I will also say a little more about 'the rebels' in Section 8.3.

1.2 Semantics, Pragmatics, and Philosophy (and Why They Are Best Done Together)

The four terms in the title of this book delimit what we are going to focus on. 'Meaning' in the sense just discussed is analysed from three perspectives: (a) devices of the language system and the primary output they express (semantics); (b) meaning-producing activities of the users of this system in their contextual settings (pragmatics); and (c) the 'overlay' of higher-level questions to do with how to define and delimit meaning, as well as related concepts such as referring, intending, inference, and semantics and pragmatics. The latter are the domain of philosophy – and within it, mostly, but not exclusively, philosophy of language. All of the above definitions are debatable, which reflects the state of the art in meaning research and, on a more positive note, testifies to the buoyancy of discussions on these matters. Where the boundary exactly lies between semantics and pragmatics has been a matter of fierce 'border wars' (to use the apt term by the American pragmaticist Laurence Horn[4]) since the late 1970s, and the disputes are not abating, although more and more linguists and philosophers opt to practise meaning theory 'without borders'. More recently, the debates have tended to focus on *metasemantics* and *metapragmatics* that ask foundational questions such as what facts endow a theory with meaning (I go into more detail in Stage 9), and, on the other hand, inroads into practical concerns that belong to the domains of ethics, sociology, social anthropology, or theory of social justice. Examples of the latter include the study of lying and misleading, language and humour, politeness and impoliteness, hate speech and within it the use of racial and ethnic slurs, to name but a few. A pertinent area here is research into *conceptual engineering* – attempts to affect the use of existing lexical items to express novel or updated concepts, for example concepts that better reflect attitudes to gender and ethnicity that are currently regarded as politically, socially, and morally correct. I introduce the semantics/

[4] See Horn 2006.

pragmatics boundary disputes in Stage 7 of the journey (although they are also implicit in the discussions throughout the book), and refer to some of these more practical, including ethical and social, concerns to do with language use in Stage 8.

The fluid and contentious boundary between semantics and pragmatics, and the associated meta-level inquiries into what the disciplines entail in the first place, demonstrate that separating them, even (or especially!) for the purpose of introducing the subject, would be a non-starter. Studying pragmatics without studying semantics is notoriously difficult if one wants to get to the bottom of the 'big questions' about meaning. One might argue that this is so because the study of what conversation interactants *do* with language presupposes understanding of how the linguistic tools they use actually work. But then, arguably, studying semantics without studying pragmatics is also very limiting, or even a non-starter to the same degree, in that it would somewhat dogmatically fix the boundary beyond which semantic concerns turn into pragmatic concerns and leave semantics helpless, unable to deal with the meaning of a vast bulk of natural-language sentences. As I try to demonstrate throughout this book, fixing such a boundary at the start takes away all the excitement and all the raison d'être from studying meaning. Language as a tool is extremely flexible – to the degree that many linguists and philosophers now argue whether words do indeed have any core meaning instead of merely functioning as 'pointers' to the meanings they help express in different contexts. Likewise, it is highly debatable whether the concept of *sentence meaning* stands up to scrutiny or, rather, whether it ought to be replaced with *utterance meaning*, in that both the lexical items and the structure in which they are used are so overwhelmingly underdetermined. Naturally, the debate concerning the feasibility of maintaining the traditional literal/non-literal meaning distinction is also under close scrutiny. I say more about the flexibility of word meaning in Stage 2 and in Sections 8.1–8.2.

The great fun of studying meaning in this way is that one can question the bulk of long-established assumptions and fixed core concepts, such as 'word meaning', 'literal meaning', or 'semantics' and 'pragmatics', and keep asking oneself what meaning really is – meaning which the devices we have at our disposal (words and structures) allow us to communicate. As was said, the meaning of the devices themselves (which is the traditional understanding of the scope of semantics) and the meaning as it is intended and conveyed in the context of discourse (which is one of the core traditional understandings of pragmatics) are, arguably, inseparable, in view of the overwhelming flexibility of words and structures. And, of course, in view of the related fact that pragmatic devices of communication are devices on a par with the lexicon and structure.

No one who writes about these topics can ever produce an objective account. This is so for various reasons. First, it is because covering all existing approaches is impossible and one has to make informed but subjective choices as to which ones are at the forefront of research, or historically important and still produce valuable explanations – or which ones are promising and sufficiently advanced as theories in progress to be included. Second, it is because every linguist starts with some initial assumptions. Some of these assumptions foreground regularity, that is finding universal rules and

patterns on which human language is founded, others focus on linguistic diversity. If regularity is aimed at, then a theorist will strive to fit meaning into the mould of theories and then provide a separate account that explains the departures and exceptions. If diversity is foregrounded, then generalizations might be sought on the level of mental representations (conceptual structures) or/and social norms rather than a theory of systematic linguistic meaning. The importance of such desiderata will be attended to throughout this journey. Another important assumption concerns the relation between meaning and structure, where the generative tradition follows Noam Chomsky in treating semantics as a component of generative grammar[5] and the tradition of truth-conditional, model-theoretic, possible-world semantics follows the outlook popularized in linguistics by Richard Montague in conceptualizing semantics as a mirror of syntax.[6] I will have much more to say about the latter than the first, beginning with Stage 3. And third, every journey through meaning reflects a personal journey through meaning. This one is mine, but its aim is to inspire the reader's own experience with meaning that will be different from this one. There will be no exercises and no memorizing of rules – just a lot of thinking about meaning.

One may ask here: what about the *objective truth* about meaning? What about meaning as an objective reflection of reality? Or meaning that is stored and processed by the human mind? Let us begin with the latter: what about the cognitive reality of the theories, views about meaning that are in the minds of the users of the language? In my view, psychological reality ought to be one of the foremost concerns of a theory of meaning, but there are limitations. First, in spite of overwhelming progress in the neuroscience of meaning in the relatively newly emerged sub-disciplines of neurosemantics and neuropragmatics, we are still quite a long way from being able to correlate neuronal structures and images of brain activity with meanings and with processing meaning.[7] And if we can't map them onto meanings, they cannot act as reliable evidence for or against many core tenets of theories of meaning ('for' and 'against' because both corroboration and falsification of theories play a part).[8]

Next, we move to 'objective meaning' as reflection of reality. That there *is* a relation is a truism and as such does not require defending, in that language either *describes* the world or *acts on* the world: meanings are always grounded in reality, be it social or that of an individual. But to understand the form this relation takes, and justify calling it 'objective', is not an easy quest. I will say more about the reasons for this in Section 1.4, and about the solutions in the rest of the book.

Semantics and pragmatics strive to understand such mappings, so the excitement of critically assessing existing theories, and the even greater fun of proposing new, more explanatorily adequate theories (especially in view of growing empirical evidence) can

---

[5] One such introduction to semantic theory is aptly entitled *Semantics in Generative Grammar* (Heim and Kratzer 1998). See also the debate between Ludlow (2003) and Chomsky (2003) on the properties of such a semantics, and notably whether it can link meanings and the world (i.e. whether it can be *referential* – a label to be discussed shortly).

[6] See de Swart 1998, p. 26; Dowty, Wall, and Peters 1981, p. 41.

[7] See Pulvermüller 2010 and further discussion in Stage 2.

[8] The best source for the methodology of scientific inquiry and the importance of falsification is Popper 1959 [1934].

continue. The concerns of philosophical semantics and pragmatics are the richer for it – just as they are enriched (rather than made obsolete, as some might naively think) by various experimental approaches to meaning. This book is about such philosophical approaches to meaning and the semantic and pragmatic theories they endorse. It is also a guide to how to go about finding one's own stance on meaning: what to question and how, and what principles to follow in making one's own theoretical choices. As I said, every journey is different but this is as it should be: inspiration for new journeys follows whenever something makes you jump and think: 'This can't be right!'

1.3   Proposition: A Flexible Unit for Studying Meaning?

As I have just pointed out, in the light of the semantics/pragmatics boundary disputes that permeate pretty much every topic to do with meaning in language, the best bet seems to be to think about semantics as the study of relations between (i) linguistic units such as words and sentences on the one hand, and on the other (ii) some *correlates* that give them meaning. Although, strictly speaking, asking about the correlates belongs to a higher level of inquiry that we have just mentioned, that is metasemantics or the investigation about the *foundations* of meaning, it is important to think about these questions in tandem. These correlates could be concepts or objects *for words* and situations in the world, or complex mental representations *for sentences*. In other words, in semantics, we are interested in (a) how sentences of a natural language such as English reflect reality – be it objects, people, states, events, processes, as well as in (b) how sentences relate to our mental representations of reality. As we will see, some approaches are only concerned with (a) or (b) and some with both.

The units of the analysis of meaning in semantics are *propositions*. In pragmatics, to reiterate, we talk about *utterances* and, more and more often, *discourses*. The terms *proposition*, *sentence*, and *utterance* require closer attention. Utterances are concrete products of speech and writing that occur in discourse. They come with information as to who the speaker is, as well as information about the time, place, and other circumstances of the performed act of speaking. Sentences are abstract grammatical units that can be extracted from utterances. The meaning of a sentence is a proposition (at least on the standard understanding of 'proposition', but read on). Propositions are probably the most recalcitrant constructs to define. They can be thought of as descriptions of situations, or what philosophers refer to as states of affairs, or, on a more fine-grained conception, as the contents of beliefs and other mental states. They do not bear a one-to-one relation to sentence types. For example, (1.4) said by me and (1.5) said to me by someone else express the same proposition that Kasia Jaszczolt is happy.

(1.4)   I am happy writing about meaning.

(1.5)   You are happy writing about meaning.

(1.6) and the Italian translation in (1.7) also express the same proposition, insofar as they are accurate translations of each other.

(1.6)   I am happy.

(1.7)   Sono felice.

Similarly, (1.8) and (1.9) express the same proposition.

(1.8)   The dog has eaten the roast.

(1.9)   The roast has been eaten by the dog.

This is so in spite of the fact that a sentence in the active voice and its passive equivalent are not identical in meaning.

Now, when a proposition is thought of as a description of states of affairs, it does not exhaust the meaning of the sentence. It is even less successful in capturing the meaning of utterances of that sentence in their respective contexts. It constitutes the core, but there is more to meaning than the proposition. This concept is now known as a Russellian proposition. The more 'fine-grained' option is that of a Fregean proposition. Fregean proposition is sensitive to the *way* the speaker (or holder of the belief) can think about the entities referred to in the sentence. But Frege's proposition is not a thought in a speaker's mind either: it is an abstract concept, which, apparently confusingly, he called *Thought* (*Gedanke*) – apparently because this is in accord with his concept of Thought.[9] Such a proposition pertains to the way of thinking about the situation but it can be shared by different speakers and as such inhabits a Platonic 'third realm': neither a physical nor a mental world. While for Frege the way of thinking about an object is an important component, Russellian propositions, so to speak, 'have the objects in them'. The importance of the 'way of thinking' or the 'sense' will be made clear very shortly when we discuss belief reports and example (1.21). The third major stance on propositions is that they can be thought of as those possible worlds in which the sentence is true. This (Carnapian) proposition[10] will become relevant when we discuss the tools of possible-word truth-conditional semantics in Stage 3.

The main recent criticism of such theoretical constructs is that they are abstract: they do not reflect the meaning held by the actual individual speaker. So, Scott Soames, for example, proposes in his recent work a new concept of a *cognitive proposition*: propositions are cognitive event types.[11] The tokens of these types are *instances of speakers'/thinkers' representing things*. By representing he means here an act of ascribing properties to objects (act of predication that we will focus on in Stage 3). Such propositions come closer to accounting for semantic and pragmatic content that is conveyed in discourse in that they are generalizations over what speakers themselves do. But his overall approach to meaning in language is still Russellian in that, according to Soames, we predicate something of objects rather than of their 'mental equivalents'. As a result, what is true or false is not sensitive to such ways of thinking.

[9] See Frege 1956 [1918–19].

[10] See Szabó and Thomason 2019, Section 5.3 on controversies about propositions and for more background on the Fregean, Russellian, and Carnapian propositions.

[11] See Soames 2014 and 2019. See also a discussion in King, Soames, and Speaks 2014 and in Jaszczolt 2021a.

So, his cognitive propositions are the 'extra' layer, so to speak. They are 'cognitive', and 'naturalized', in the sense of not being abstract, but what is true or false applies only to their 'objective' equivalents. But even as cognitive propositions they are very sentence-bound: they are 'naturalized', in the sense of being about people and their meanings, but they cannot capture human patterns of reasoning when what is said is expressed indirectly or non-literally, for example.

So, the question to ask next is whether we want semantics to capture the ways speakers conceptualize states of affairs. We want to ask what kind of meaning semantic theory should be about in that different answers to this question will trigger different requirements of what a proposition ought to be. Linguists are happy to treat proposition as a functional, technical construct and choose the concept of it that does the best job for explaining meaning. If semantics is to include more nuanced meanings that the speaker, or the addressee, or both have in the mind(s), then proposition as the unit of meaning will follow suit and reflect it too.

And we have a lot to choose among – from proposition as (i) the meaning of a sentence, through (ii) the bearer of a truth value ('true' or 'false', on which more in Stage 3), to proposition as (iii) an object of beliefs and other mental attitudes. They pull us in different directions: should a proposition be the meaning of the sentence or a mental state the sentence is about? If the latter, do we construe it in terms of abstract, intersubjective *types* of mental states or in terms of *tokens*: the actual states? Stephen Schiffer[12] opts for propositions as objects of beliefs. He calls them *pleonastic propositions*, in that they are constructed, so to speak, 'out of nothing', simply in virtue of the general characteristics of our cognition and our use of language. They capture the nuances of meaning that a cognitively plausible semantic theory ought to capture but the pay-off is that they do not capture the *structure* of the natural-language *that*-clauses. This is important because *A believes that* ... construction linguistically represents the content of beliefs. His propositions are 'unstructured', as philosophers say. This reflects Schiffer's questioning of the compositionality of natural languages, that is the assumed property of languages that their meaning is given by the meanings of the lexical items and the structure of the sentence (more about this principle shortly). But as will be seen in Section 5.5 where I discuss the meaning of belief and other propositional attitude reports and Schiffer's contribution to the debate, his stance need not lead to such a drastic denial. For now, suffice it to say that our functional propositions introduced immediately below do better in this respect: they are also cognitive and naturalized, and they are sufficiently finely grained to capture the required nuances of thought. They are also paired with a semantic theory that allows for such nuances to be included. Details will have to wait until Section 7.2.3.

All in all, propositions are flexible and controversial beasts. No wonder that McGrath and Frank (2018: 1) call 'proposition' a 'quasi-technical word'. And yet, propositions have proven to be indispensable in most theories of meaning that strive for some formalization and predictive power. So, equipped with Fregean, Russellian, and various naturalized propositions, let us ask now how much content, and what kinds

---

[12] See Schiffer 2003.

of content, a proposition ought to contain to serve the purpose on this journey through meaning – that is, to serve a cognitively real theory of natural-language meaning. If semantics is to rely only on the words and structures of the sentences, we may have to do without a proposition. (1.10) does not correlate with specific situations, or even with situation types: who is 'ready', for what, and what time does 'is' refer to?

(1.10)   He is ready.

If we want to adopt propositions as units of meaning and if they are to describe states of affairs, don't they have to be embellished, filled in sometimes, when the sentence itself does not pick out a situation uniquely or when it picks out *no* situation without such an embellishment? Yes, they do, unless we opt for some very general propositions that would pick out a range of different situations or even different situation types and as such be of little use in explaining meaning. At the very least, the route from a sentence containing essentially context-dependent expressions (*indexicals*), such as *I*, *here*, and *now* in (1.11), to propositions has to involve filling in the referents (here, the speaker, the place, and the time of the utterance).

(1.11)   I am here now.

This resolution of indexicals is a pragmatic, situation-driven process, but it is pretty much uniformly accepted as part of semantics too: we need to resolve them to obtain a proposition. The semantic roles of *I*, *here*, and *now* are almost always fixed: they stand for, respectively, the speaker/writer, the place, and the time of the current act of speaking/writing (*pace* a handful of interesting cases that I discuss in Section 6.2.4). Indexicals that require a greater recourse to context, such as *he* or *then*, are the first tier on which opinions begin to diverge: is their resolution part of semantics as well, or does it belong to the pragmatic overlay? In other words, do the processes of assigning referents to such highly context-sensitive words contribute to the semantics of the sentences in which they occur, and as such to the proposition that is their meaning? Well, yes, if we pin down a particular utterance of this sentence. And when we do, we enter the territory of pragmatics. On the other hand, to reiterate, notice that when we don't fill the indexicals in with referents, we either have to relinquish propositions as objects of study or make propositions so general that *he* and *then* pick out, say, 'a conversationally salient male' or 'a contextually salient time', nothing more. So, it appears that we are still on the territory of semantics after all. To move to the next tier, if we want our semantics to be strongly guided by the criterion of cognitive reality, we can allow even more contextual resolution of referents: filling in slots provided by indexicals of any kind. For this purpose we can arguably include *ready* in (1.10) in this category, as it is in need of filling in 'ready for what?'[13] But we can also include expanding propositions where there are no such slots to be filled, as in (1.12), analysed as (1.13). Indexicals in (1.11) provided obvious 'slots' in the syntactic structure that require filling in. On the other hand, in (1.12), the quantifying expression 'everybody'

---

[13] The membership of the category of indexicals is a moot topic. I discuss it in Section 6.2.4. See also Kaplan 1989a; Cappelen and Lepore 2005; Borg 2004, 2012.

does not come with a syntactic slot but comes with a domain of quantification that requires a restriction in a given context (everybody *in what category/group of people*?), for example as in (1.13).

(1.12)   Everybody signed up for the *Semantics, Pragmatics and Philosophy* seminar.

(1.13)   Every linguistics student at Cambridge signed up for the *Semantics, Pragmatics and Philosophy* seminar.

Next, there are sentences that, in spite of being syntactically complete, do not provide enough information to identify a corresponding situation without help from context. Here, again, there is no syntactic slot to be filled, but, unlike in (1.12), there is no 'wrong proposition' ('Every person *in the world* signed up for the *Semantics, Pragmatics and Philosophy* seminar.') either: the hearer has to complete the speaker's thought in order to find it, as for example in (1.15) corresponding to the uttered (1.14). Or like (1.10) above.

(1.14)   Pyjamas are not interesting enough.

(1.15)   Pyjamas are not interesting enough to be given as a birthday present.

So, if the hearer has to 'find it', what does the semanticist do? They have to find it too, or else resort to the kind of semantics that has something to say only about the meaning of the tools we are using for expressing meaning in language (that is, words and structures of the language system) but rather little about the meaning itself – at least little of cognitive validity. If we allow the filling in of indexicals, then wouldn't there be tension in not allowing the resolution of meaning in (1.14) as well?

But we can go a bit further. As I said earlier, we can 'amend' Frege's fine-grained propositions and make them less abstract. And make them stand for the meaning that the holder of the belief actually has in mind – and, importantly for us, the meaning that the speaker and the addressee home in on as the primary message of the utterance. (I discuss joint construction of meaning by the interlocutors in Section 7.1.4.) But first, we can adopt the view that there is a need in theory of meaning for different kinds of propositions, corresponding to different aspects of meaning. Then, if we need propositions to account for conceptual structures, we can even go as far as adopting the stance that it is the proposition that corresponds to the *main message* (covertly or overtly expressed, no matter!) that is intended and conveyed by the speaker that constitutes the object of study of semantics. I call it 'intended' and 'conveyed' in agreement with the desideratum discussed earlier that the theory of meaning looks out for regularities and descriptive norms. Miscommunication then falls outside the scope of the theory of meaning but is firmly placed in the study of language processing in psycholinguistics and sociolinguistics. The kind of proposition that we will have to make use of then will have to have the status of what I call a *functional proposition*:[14] its content is delimited by the informative intentions of the speaker, not by the structure of what the speaker actually says. Depending on one's interests, say, in what the tools can do or what the

[14] See Jaszczolt 2021a.