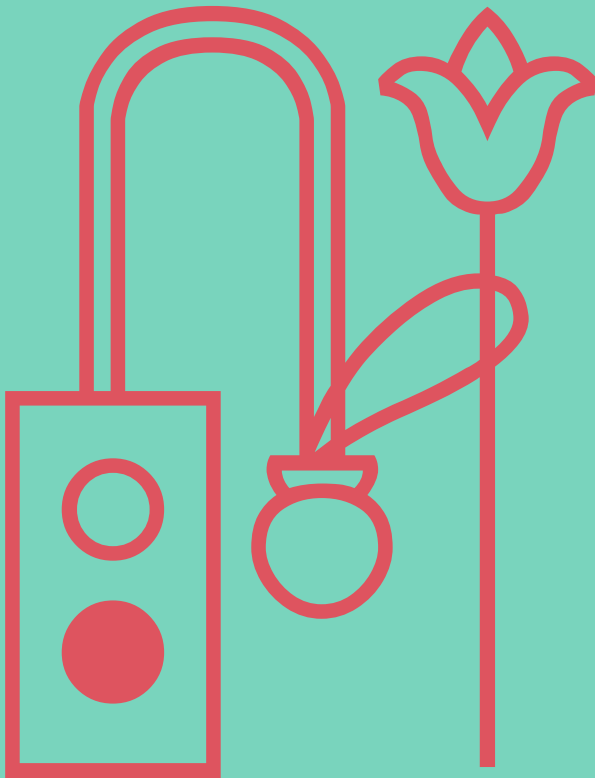Richard Sproat

Symbols

An Evolutionary History from the Stone Age to the Future

Symbols

Richard Sproat

Symbols

An Evolutionary History from the Stone
Age to the Future

Richard Sproat (ORCID)
Google Japan
Shibuya City, Tokyo, Japan

Cover illustration: The image on the cover is based on an idea of the author and represents the Ancient Egyptian hieroglyph for "scribe". It depicts the scribe's equipment consisting of a tube for holding reeds, a leather bag for holding ink and a palette for mixing ink.

Todo lenguaje es un alfabeto de símbolos cuyo ejercicio presupone un pasado que los interlocutores comparten.
                    Jorge Luis Borges, "El Aleph"

Preface

Nearly 20 years ago I was contacted by Steve Farmer, a comparative pre-modern historian. I had recently published a book (Sproat, 2000), where I developed a formal computational theory of writing and its relation to the language and speech it encoded. He had found that and a previous book of mine useful in helping him understand computational linguistics and formal models of language, and he had a question for me. His question was deceptively simple: was it "possible to distinguish statistically any linguistic character string represented by a fairly large corpus of texts from non-linguistic symbol chains?"

Steve's interest in this question stemmed from his collaboration with Harvard Indologist Michael Witzel. They had recently published an article (Witzel and Farmer, 2000), which debunked a recent much heralded "decipherment" of the so-called Indus Valley Script, the to-date largely uninterpretable collection of short cryptic "texts" from the Bronze Age Indus Valley Civilization (3rd Millennium BCE). After his work on that article, where he and Witzel largely accepted the standard assumption that the Indus Valley had a full-blown writing system, Steve started to consider various aspects of the Indus Valley culture and was coming to suspect that the symbols were not writing at all, but some sort of non-linguistic system or at best a "proto-writing system." For one thing, the only surviving texts (on non-perishable materials such as steatite) are exceedingly short. Despite claims that there must have been long texts—indeed a whole lost literature—on perishable materials, no evidence of *non-perishable* paraphernalia required to produce such texts—writing utensils, ink pots—has been forthcoming. Nor, despite 700 years of use, do the Indus symbols seem to show the kind of evolutionary shape changes found in writing systems that are in wide scribal use (Kelly et al., 2021). Such a situation, stable over 700 years, flew in the face of historical developments found in all known literate civilizations (Farmer et al., 2002).

Hence his question to me: nobody knew what the symbols meant, but would it be possible to tell just by the distribution of the symbols in texts, whether they represented language or not? To put it in modern terms, if you didn't know either written English or mathematical symbology, could you tell that the former represented a language and the latter was instead a formal non-linguistic system? At

the time, I didn't think there were any reliable tests, and indeed I still think that, a topic I take up later in this book.

But that initial email led to a collaboration that resulted in a paper (Farmer et al., 2004) that managed to convince people or offend them in roughly equal measures. This led to further controversy about the Indus symbols in particular and more generally whether statistical methods are informative about the kind of things graphical symbols denote, a topic which, again, I take up later in the book.

But it also got me to thinking about non-linguistic symbol systems and what their status is. Written language is prominent in our minds because it is so ubiquitous in the modern world, and we tend to think much less about the various graphical notation systems that communicate information, but do so without reference to language. Yet these systems are also ubiquitous, also communicate various different kinds of information, and, most importantly for one of the themes we will develop in this book, long predate writing in the history of human civilization.

They are also largely misunderstood, apparently. Many of the critiques of our 2004 paper involved the presumption that non-linguistic systems are simple and structureless. As we shall see in this book, nothing could be further from the truth. Yet it is this misunderstanding that has led people, as we shall see, to assume that evidence of structure in a symbol system is ipso facto evidence of *linguistic structure*.

One purpose of this book is to try to put such misconceptions to rest, and one way to achieve that goal is to examine in some detail a variety of non-linguistic symbol systems with different symbol set sizes, different ways of combining symbols, and different functions and, in the process, develop a taxonomy of such systems. It will become clear as a result of this that structure in a symbol system simply relates to the complexity of the domain that the system is used to represent.

Another purpose of the book is to understand better the relationship between non-linguistic symbol systems and a type of symbol system that has a special place in the history of humankind: true writing. To this end I delve in some depth into how writing works, how it relates to the speech that it encodes, and how speech and writing differ in terms of what they can and cannot easily express. But there is another aspect of the relationship between non-linguistic systems and writing that needs to be understood: *evolution*. It is widely accepted that writing evolved out of formerly non-linguistic systems. In Mesopotamia, where the evidence is clearest, the non-linguistic system was an accounting system. Also widely accepted is that the key point in that evolution was the realization that symbols could be used not just for what they *mean* but for how the words for the concepts the symbols originally encoded *sounded*. This allowed for the transfer of symbols to represent other words that sounded similar. In this book, I explore what this must have meant in neurological terms, and I offer a hypothesis as to the institutional context in which the symbol-sound correspondences would naturally have been *trained*. I offer a computational simulation in support of this hypothesis.

If I am successful at these various quests, it will have been in large measure because of the many scholars I have interacted with over the years, who have helped me understand the limits of my thoughts, and helped me revise them. Most people who have come to the study of writing systems and graphical symbol systems

more generally have come to it from a background in the humanities—often with a specialization in one or another writing system, ancient or modern. My background is different. Trained as a formal linguist, I took up computational linguistics when I moved to AT&T in the mid-1980s, and thence moved into an area that was a research topic of interest at Bell Labs at the time: text-to-speech synthesis. I became interested in the problem of language processing with a view to having a system—a computer— read text. As a result I became interested in the relation between written language and the speech that one can generate from it, which led to my becoming interested in writing systems and to the 2000 book referenced above. That first email from Steve Farmer was the impetus for me to try to understand more about graphical symbol systems that were not tied to language, and the relation between these non-linguistic systems and written language.

In this journey I have benefited from discussions with and feedback from many people including, at various times over the years, Michael Witzel, William Boltz, Christopher Woods, Edward Shaughnessy, and Suyoun Yoon; as well as audiences at talks I have given over the years on related topics at *Empirical Methods in Natural Language Processing*, the *Berkeley Linguistics Society*, Kings College London, York University, the Max Planck Institute for Evolutionary Anthropology in Leipzig, Saarland University, Johns Hopkins University, Carnegie Mellon University, and the Signs of Writing Conferences in Chicago (2014) and Beijing (2015).

I have benefited from extremely detailed comments on previous versions of this work by Jacob Dahl, Kyle Gorman, Alexander Gutkin, Steve Farmer, Zoltan Somogyi, and Brian Roark. I also thank Sven Osterkamp and one anonymous reviewer for feedback and suggestions.

My interactions with Rajesh Rao and Rob Lee and their colleagues (see Chap. 8), while often heated, have proved very useful in that they forced me to think deeper about the issues surrounding putative statistical tests for whether a symbol system is or is not true writing. As is often the case, the people who disagree with you the most can often be the ones who most force you to understand the issues better.

Finally, I would like to thank Alexandru Ciolan, my editor at Springer, for his support during the production process.

Shibuya City, Japan   Richard Sproat

Contents

Chapter 1
Introduction

1.1 What's in a Symbol?

Most people in the world will be familiar with the octagonal red stop sign. While the sign is typically accompanied by a written text message—e.g. *STOP*—that written message is not really needed for the sign to convey its meaning. To see this, one only needs to consider the stop sign in Fig. 1.1: even if you do not read Arabic, it is likely that if you see this sign at an intersection, you would know what it means. The red octagon is a *conventional symbol* for the notion "stop".

This of course is a simple case: it involves a single symbol on its own conveying a simple piece of information. The stop sign is in turn one of a set of more or less conventional symbols used to convey information along roads. Typically these are also simple symbols, conveying relatively simple information. Sometimes such symbols may be combined: one may see a sign indicating a filling station, and next to it a sign with a knife and fork indicating a restaurant. In this case the "message" is that nearby one can fill one's vehicle, and get a meal. The messages that one typically finds involving road signs are invariably simple like this, and the ways of combining them—e.g. just lining them up—are also simple. More on this in what follows.

But not all symbol systems are simple. Let us consider an example of a system that involves far more complicated combinations, namely an example of medieval British heraldry. Let us say that I would like to have a coat of arms. Of course, to get one I would have to apply to the College of Arms,[1] which application would almost certainly be turned down. But let us ignore that for now. I always liked the color blue as a child, so I have decided my background will be blue, or *azure* in the terminology of the formal language *blazon*. (See below in Sect. 3.5 for more on blazon.) I want to keep the design simple, so apart from the azure *field*, I'll just add a single *ordinary charge*, namely a *bend*. According to the *first rule of heraldry*, since my field is a *color*, my bend must be of one of the two designated *metals* (it could also be a *fur*),

---

[1] www.college-of-arms.gov.uk.

Fig. 1.1 An Arabic stop
sign. Source: Wikipedia.
https://en.wikipedia.org/wiki/
Stop_sign#/media/File:
Saudi_Arabia_-_Road_Sign_
-_Stop_(Arabic).svg, Author:
Qrmoo3. License: CC BY-SA
4.0

Fig. 1.2 The shield of my
hypothetical coat of arms:
*Azure, on a bend argent three
brussels sprouts proper*

namely gold (*or*) or silver (*argent*). I am going to pick argent, because *or* invariably
renders as yellow, and yellow on blue is a bit garish, in my view. Finally, being
addicted to bad puns, I am going to design my arms to be *canting arms*, typically a
pun on the name of the bearer. In my case I will pick brussels sprouts since *sprout*
is similar to *Sproat*, and in fact some people in my experience never seem to be able
to make the distinction, and finally because, due to this resemblance, my nickname
in elementary school was *brussels*. I will pick three brussels sprouts, since three is
a common number for a charge, I will place them on the bend, and since brussels
sprouts would look just like green circles otherwise, I am going to pick a realistic
looking rendition for them, in blazon terminology *proper*. And there we have it:
Azure, on a bend argent three brussels sprouts proper. See Fig. 1.2.

The point of this bit of fiction is to illustrate some of the features of symbols. In
the unlikely event that I applied to the College of Arms to be granted this shield,
and that my application were granted, then first of all these arms would represent me
and would in effect *mean* me: I could have them printed on household items or on

letterhead and it would identify those items as being related to me. More technically, the *denotation* of the arms would be me—though denotation cannot generally be thought of in such simple terms as a symbol being related directly to an object in the real world. (After all, the denotation of a symbol may not even exist in the real world, as is the case for the denotations of words like *dragon* or *unicorn*.)

The symbol itself is really a message of sorts composed of simpler symbols, some of which are abstract—e.g. the bend and the colors and metals (a color and metal being basic parts of the symbol system of heraldry); and some iconic, like the sprouts, being depictions that clearly evoke an object. As with many complex symbol systems, the rules of heraldry restrict how things may be combined, hence the limitation on how I can color the bend, noted previously.

So much for the denotation and a crude characterization of the basic symbols making up this complex symbol. There are in addition the various *connotations*, which are evoked by my choice of symbols. Most obviously there is the pun on my name, along with a bit of personal history embedded in the design. There is also my preferences in colors. The pun will likely be obvious to the observer, the other choices perhaps not so obvious, but all of them are part of the history or, to couch it in linguistic terminology, the *etymology* of the symbol.

All symbols carry with them a set of denotations and a set of connotations. The denotations are simply what they are used for. The 'A' at the beginning of this paragraph is used to represent, in its most basic linguistic use, a particular vowel—or in English, anyway, a set of vowels depending on the context. In other uses it has other denotations: for example it can denote the highest letter grade in an academic setting. Or a particular set of musical notes. Or a particular train in the New York City subway system. But depending on the particular use case, the other denotations effectively become connotations, insofar as they are typically irrelevant in the given context: If I am reading the letter 'A' in the word 'insofar', the academic grade sense of that symbol is irrelevant. And there are connotations that may arise because of the history of a symbol. In the case of 'A', the letter descended over nearly 4000 years from what was originally, in the ancestral Proto-Sinaitic scripts, a picture of an ox's head—e.g. Phoenician *'lp* 'ox', originally representing a consonant, a *glottal stop*, and only later reinterpreted by the Greeks as a vowel when they borrowed their alphabet from Phoenician (Gnanadesikan, 2009). If you happen to know this history, that connotation is lurking in there somewhere.

It is fair to say that a lot of the perceived 'magic' of symbols relates to their connotation as much as to their denotation. Symbols have often been believed to have special powers, and hidden meanings. Kabbalists such as Franciscus Mercurius van Helmont (1614–1699) found hidden meanings in the letters of the Hebrew alphabet far beyond the mundane use of those symbols to represent the sounds of the Hebrew language. See Sect. 3.3. Symbols are a *technology*, where most of the time the purpose of the technology is to communicate some sort of information. As James Burke showed in his 1970s era book and TV series *Connections* (Burke, 1978), any given piece of technology has a history that can often involve many unexpected twists and turns: the (now thoroughly antiquated) computer Hollerith card had its origins in cards used to control the operation of the Jacquard loom. In similar fashion, symbols

often have a complex history and thus a lot of excess baggage—to which is often added imagined baggage of the kind that the Kabbalists were fond of. But ultimately symbols are just a technology.

We explore some of the properties of that technology in this book, and one of those properties is that in many symbol systems, there are clear rules on how one can combine basic symbols into more complex messages, like the color restriction ("first rule of heraldry") noted for my imagined coat of arms described previously. We turn now to a brief introduction to the topic of *syntax*.

1.2 Syntax

I have five different signs in the top row of Fig. 1.3, represented as five differently colored hexagons. What, if anything, can we say about the fact that these five different signs are ordered in the particular way they are?

If you think about this for a moment you will quickly realize that the answer must depend upon what the signs are, i.e. what they represent. Suppose for example that the black hexagon represents the English word *the*, the white hexagon the word *dog*, the red hexagon *is*, the blue hexagon *very* and the yellow hexagon *hairy*. Then the five signs make up the sentence *the dog is very hairy*. In this case one can say a lot about the order, namely that in this particular order the signs make up a meaningful and grammatically correct English sentence; and that in most other orders they do not. In other words, the words are constructed into a sentence according to the rules of the *syntax* of English.

Move now to the second row: here we have a row of five road signs telling us that at this place one can buy fuel, find a telephone, get a meal, get a room for the night, or
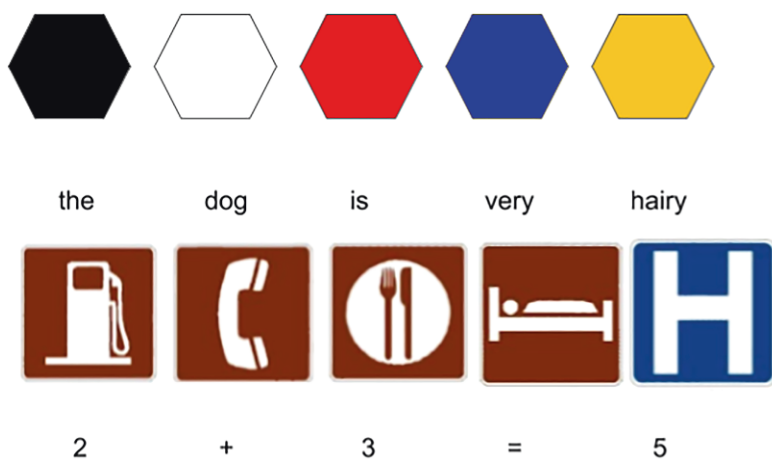

Fig. 1.3 Messages with five signs

Fig. 1.4 A typical
arrangement of informative
signs for fuel, food and
lodging

find a hospital. Like the English sentence these signs also communicate information. But in this case the order does not matter: one could reorder these signs in any way, and the "message" would be the same. There is no syntax here. Or is there? While it is true that it would actually make no difference if the signs were in a different order, there are certainly conventions on how these signs tend to get placed. Fuel signs tend to be placed first, perhaps because buying fuel is one of the most common things one does on highways, followed by getting a meal or needing a telephone. People tend to look for lodging only at certain times of the day, and with any luck, most people will never need a hospital while they are traveling. The arrangement in Fig. 1.4 is fairly typical. So convention tends to place these signs in certain orders. This is not really syntax: it would not be *wrong* to place them in another order. But imagine for a moment that one did not know what the signs meant, and that one merely observes that these signs tend to recur again and again in the same order. Just statistically one might well conclude that there is a syntax here, even when in fact there is not.

But those issues aside for the moment, *linguistic* messages such as *the dog is very hairy* have a definite syntax, whereas non-linguistic signs such as our road signs do not. So is syntax a sign of language? If I can definitely ascertain that there is a syntax to a set of messages, can I conclude that this means we are dealing with something that relates to language?

The final example should disabuse the reader of that notion. The arithmetic expression $2 + 3 = 5$ also has a definite syntax. I can change the 2 and 3 around without violating the syntax or the meaning. I can swap the 5 and the 3 and have something that is syntactically well formed—but of course will now be false. But, under the familiar arithmetic notation system being assumed here, I cannot for example swap the 3 and the $+$ or the 5 and the $=$ and still have a well-formed equation. The rules of the syntax of mathematical expressions are of course very different from those of natural language at least in part because the former is a formal and artificial system that was developed by convention, whereas English is a natural language that evolved naturally over time. But mathematical expressions have syntax nonetheless. Syntax or structure is not determinative of something being natural language.

Perhaps though we need to clear up one possible misunderstanding about expressions like $2 + 3 = 5$. Obviously I can *read* such an expression in English (*two plus three equals five*), so one might be inclined to think that this is after all just language. But while $2 + 3 = 5$ can be read as a linguistic message, one can in fact do this in practically any language (or at least in languages for cultures that have non-trivial counting and some notion of arithmetic) so that one could easily render this expression into French or Japanese as one could into English. The fact that one

can *read* a non-linguistic message using language is not the same as saying that the message represents language or linguistic information.[2]

All of this may seem trivial and obvious but it is surprising how often this point is misunderstood. In this book we will make a clear distinction between *linguistic* symbol systems—i.e. true writing systems; and *non-linguistic* symbol systems, which is any other graphical symbol system that conveys some sort of information but is not tied to language. Indeed we will use the term *writing* throughout as *synonymous with a linguistic symbol system*. We acknowledge that there are those who use the term *writing* to denote *any* system of conventionalized graphical marks, but that is not the way we will use it in this book. It is important to be clear about this point since much confusion has been sown by vague uses of the term *writing*. We return to this issue more in Chaps. 4 and 8.

Returning to the main theme, one of the most common misunderstandings about non-linguistic systems is the assumption that such systems are ipso facto structureless and without syntax. While many are indeed that, many are not, and one of the goals of this book will be to convince the reader that non-linguistic systems may have structure, often *rich* structure, and whether or not a system displays syntax relates to the kind of message it is designed to convey rather than whether or not it is tied to language.

But does the syntactic structure of some non-linguistic systems still depend on language in a different way? Is it because of the structure that natural languages have evolved over tens or hundreds of thousands of years, that it is possible for humans to construct other, non-linguistic, communications systems that themselves have structure? Probably, though it would be hard to demonstrate that this must be the case. The evolution of language itself and what it was precisely that evolved—a distinct mental "module" devoted to language as nativists have argued, or simply a complex use of cognitive functions adapted for other purposes—is a contentious issue. How writing itself evolved from prior non-linguistic systems is also unclear, though we will suggest some possible mechanisms in Chaps. 6 and 7. As to how more general graphical symbol systems evolved, and to what extent their evolution depended on the prior existence of a developed language faculty we can only speculate. We do know that non-linguistic icons are processed differently in the brain from written language, and in a way more akin to the interpretation of pictures (Huang et al., 2015), so at least that suggests that there is a more tenuous connection to language for non-linguistic systems. But those experiments considered only non-linguistic symbols in isolation, not non-linguistic symbols from a syntactically complex symbol system being used in complex constructions: such cases might involve processing more akin to what happens in the brain when spoken language is processed. We return to this theme in Chap. 5.

---

[2] We will however argue later that the reading of non-linguistic symbols aloud was probably critical in the evolution of the first writing systems developed.

1.3   What this Book Is About

This book is about graphical symbols, what sorts of things they denote and how, in some systems, the symbols can be combined into complex messages. It is not about *signs* more generally: that is the domain of *semiotics* which has been characterized by one prominent semiotician as being "concerned with everything that can be taken as a sign" (Eco, 1976, page 7), and includes such things as smoke signals, or even rashes as symptoms of underlying diseases. The domain of this work is man-made graphical symbols,[3] and in particular ones that are conventionalized in that their form and meaning are agreed upon by a large community of users. We will define this notion more formally in what follows. And since semiotics as a field would seem to relate to the topic of this book, I will briefly review that field (Chap. 2), and point out where the topic of this investigation differs and in many ways goes beyond what semioticians have typically dealt with. We will see that, at most, semiotics and the topic of this book have a non-null intersection, but that there are areas of semiotics we will not cover and in contrast we will be concerned with areas that have at best received lip service by semioticians.

One of the distinctions we will introduce briefly in Chap. 2 is the important distinction between *non-linguistic* symbol systems, and writing, which encodes linguistic information. Writing is perhaps the most familiar symbol system since it is one that most of us use every day. Less familiar are the many varieties of non-linguistic systems. Thus, in Chap. 3, I will present a taxonomy of *non-linguistic* graphical symbols, according to what they denote, and what the possibilities are, in the given system, of combining the symbols into more complex messages. In order to understand the workings of a couple of these in greater detail, I will also present in that chapter an in-depth comparison of two systems of heraldry, European heraldry, and Japanese *kamon*, which served much the same function in the two cultures, but differed in their syntactic combinatorics.

Writing is a special case of a conventionalized graphical symbol system that has complex syntax, formally defined as a symbol system that encodes information from

---

[3] Since I use the term "graphical" throughout this book, one might wonder where *tactile* systems—the most notable example is Braille—fit into this schema. At the risk of overextending the meaning of "graphical", I will assume that tactile systems are also instances of graphical systems. As a practical matter, systematic tactile systems have been rare in history. Traditional symbol systems do include cases that probably at least had a tactile component: *khipu* (Sect. 3.6.12) may have been such a system. But the widespread use of a conventional symbolic system based on symbols that one could sense by touch seems to have been relatively modern. Braille itself was first developed in the nineteenth century by Louis Braille and was inspired by an earlier system by Charles Barbier (Barbier, 1815; Henri, 1952). Braille of course has spawned a whole family of tactile systems based on arrangements of dots, not only to represent a large number of written languages, but also numerical and mathematical information, among others. But again, this is a quite modern phenomenon. But in any case, it seems reasonable to assume that tactile systems fall under the rubric of graphical systems more generally. If nothing else, as we will mention in Chap. 5, Footnote 4, the processing of Braille by blind users seems to make use of the same areas of the brain as the processing of standard written forms by sighted users.

natural language. The literature on writing systems is significant, with quite a few books and other works having been added in just the last few years. Chap. 4 will merely review the main issues, focusing on the question of how writing systems encode linguistic information, what linguistic information is encoded, and what it takes for something to be a full writing system. As has been pointed out many times elsewhere, we will see that full writing systems must encode *phonology*, that is they must have some way to represent the sounds of the language: they cannot just represent meanings. Unfortunately this empirically derived observation that all full writing systems must encode phonology has been misinterpreted as "speech centric" by those who prefer to emphasize the commonality between writing and other graphical systems. But as we will see, there is really nothing to argue about here.

How are symbols processed in the brain? In Chap. 5 we review some of the literature on that topic, and point out some of the differences between how writing is processed and how other symbols are processed. Crucial in the processing of writing as opposed to non-linguistic symbol systems is the involvement with writing of the language processing areas, in particular those related to phonology.

Writing systems are special, but they evidently evolved from non-linguistic systems. While there have been hundreds of writing systems developed throughout history, as far as we know the *pristine* invention of writing happened in only three, or at most five places independently of one another—see Sect. 6.1. How did it happen? Unfortunately the archaeological evidence on that point is almost non-existent in all but one culture (Mesopotamia), and sparse even in that case, so the best we can do is speculate. But we can at least do one thing: we can simulate the evolution of written language from non-linguistic symbol systems using computational models. The evolution of writing will be the topic of Chaps. 6 and 7.

Chapter 8 will be about a topic that seems to be a source of confusion: given a symbol system where we do not know the meaning of the symbols—say a symbol system from an ancient civilization—what can we say about that symbol system *before* we have established what the symbols mean? One of the most common assumptions when faced with an unknown system is to assume that must have been some form of written language, especially if overtly it "looks like" writing. (We will also discuss what people seem to think it means for something to "look like" writing, and some of the pitfalls in the assumptions.) But can we ascertain this short of a decipherment? The past couple of decades have seen claims that one can determine the status of a system on the basis of statistical properties of the distribution of symbols in extant "texts" of the system. The little demonstration in Sect. 1.2, Fig. 1.3, ought to give a clue as to how successful such approaches are likely to be since they are more or less the computational equivalent of trying to determine for our sequence of colored hexagons, what sort of system one is dealing with. In any event, we review some of this recent work and point out some of the difficulties it faces.

Finally in Chap. 9, we will look at one question that seems to be a recurrent theme in the popular press, namely whether *emoji* or some similar form of non-linguistic symbol system could replace written language as an effective and complete form of communication. In Neal Stephenson's *The Diamond Age*, a large segment of

the population's written communication was via *mediaglyphics*, animated glyphs that are supposed to be able to communicate the same sort of information that conventional writing systems do. Then there are jocular exercises like *Emoji Dick*,[4] which attempts to translate Herman Melville's *Moby Dick* into emoji, which has led in turn to others wondering just how much of written communication can be replaced by sequences of little icons (WNYC, 2014). We close with some thoughts on these ideas.

---

[4] https://www.kickstarter.com/projects/fred/emoji-dick.

Chapter 2
Semiotics

The topic of this chapter is the field of *semiotics*, which is broadly interested in the topic of *signs* and their meanings. As such it would seem to be very related to the theme of this book. However, as we shall see, there are large areas of semiotics that are not directly relevant to the current study and, contrariwise, the current study goes well beyond semiotics in some ways. Indeed, the main reason for including this chapter is to explain how the present work differs in focus from much of what most semioticians are interested in. It is likely that semioticians will not be happy with my explanations, but I hope that the explanations will at least seem defensible.

That said, the reader who is not particularly interested in semiotics—or who already knows the field but is not particularly interested in how the present work differs from it—may safely skip most of this chapter. However he or she will still want to look at Sect. 2.5, which introduces the semiotic notion of *articulation*, which we use elsewhere in the book.

2.1 Introduction

There are many ways of communicating information. The communication may be *intentional*, as in my writing this text with a view towards explaining a particular view of symbol systems to a hopefully interested audience. The communication may be unintentional, as with an overt collection of symptoms that can be thought of as "communicating" the presence of an underlying disease. Or they may be intentional or unintentional depending on circumstance: when people speak they typically accompany their speech with facial expressions and gestures, of which they may or may not be aware depending on the case.

In addition to intentionality, another dimension is *conventionality*. The written symbols used to write this book are conventional since as readers and writers of English, we as a community agree that this set of symbols assembled in a particular way can be used to encode messages that are to be read and understood in English.

But if two people develop a secret code that only they understood, that is not conventional.

Yet another dimension is *graphicality*. Again, the written symbols I am using to write this text are graphical insofar as they are (or at least show up as) marks on a surface. But speech is not graphical. Neither is music—though *musical notation* is.

Consider, then, the following issues that one might study as forms of communication:

1. Written language and, more specifically, how writing encodes language.
2. Graphical signs that conventionally represent information but are not tied to language: mathematical notation, chess notation, traffic signs, guild signs, heraldic/emblematic signs …
3. Spoken languages and sign languages.
4. Spoken *argots* (codes), such as Boonville jargon (Rawles, 1966).
5. One-offs: you and I agree that an apple on the stoop means I am home but not to be disturbed. Or a *one-time pad*, which in cryptography denotes a formula for encrypting and decrypting a message that only the sender and receiver know, and which is only used once.
6. Birdsong. Frog mating calls. Dogs marking locations with urine.
7. Slime mold chemical signals.
8. Genetic codes.
9. Clear mucus and a chronic cough are symptoms of pertussis.
10. Various more or less indirect ways of communicating information: *metaphor*, *metonymy* ("the ham and eggs wants more coffee"), *irony*.
11. Myth and what it communicates.
12. The connection between texts.
13. The main character Sen in Miyazaki Hayao's *Spirited Away*, who is entrapped into working in a spirit world *onsen*, has to be understood in the context of child sexual slavery (Info, 2020).

All of these topics would fall somewhere in the broad field of *semiotics*, relating as they do to *signs* of one kind or another and what they communicate.

In this work we will be interested in exactly two of these: 1 and 2, both of which are *intentional*, *conventional*, *graphical* forms of communication. We will be largely unconcerned with the other examples, though to understand writing and how it works one needs to understand how it relates to the third item, language.

Of course a semiotician might argue that by limiting ourselves in this way we are automatically handicapping ourselves, since we are not considering the full range of possible ways in which information may be communicated. True, perhaps, but the first two items on our list above form an important natural class. Writing in particular has been termed the "technology of civilization" (Powell, 2009), and there is a lot of merit to that view: without writing, detailed record keeping, beyond what one could record using accounting symbols or other non-linguistic symbol systems, would depend on human memory, and there is much reason to believe that in most purely oral cultures (cultures that lack writing), records are in general much more fluid than in cultures that possess writing (Goody & Watt, 1968; Goody, 1977). And writing

in turn depends on non-linguistic symbol systems, our second item above, since it is generally believed that writing *evolved* from a prior non-linguistic system. Thus *our object of study will be those semiotic systems that might, in principle, have evolved into writing.* In any case, as we shall see, the first two already form a rich set of cases, from which we can glean a lot about information and how it is communicated.

Since we have introduced the term *semiotics*, it behooves us to say a little more about that field, which we do in the next section. We will in particular focus on why for the present study I am not making an active attempt to align my story strongly with that much broader field. In the subsequent two sections, we take up two themes, *iconicity* and syntax that, respectively, have and have not been the main focus of semioticians. Finally we will end the chapter by discussing *articulation*, a semiotic notion that will prove useful in our subsequent foray into a taxonomy of non-linguistic symbol systems.

2.2   The Field of Semiotics

The field of semiotics, broadly construed, dates back millennia at least to Plato and Aristotle, but the modern incarnation of what we recognize as the field starts with the work of Charles Sanders Peirce, an American philosopher of the nineteenth century (Peirce, 1868, 1934). More recent figures have notably included Umberto Eco (1976) and Thomas Sebeok (1977; 2001), who is credited with being one of the founders of the field of biosemiotics. Good reviews of the field can be found in Bouissac (1998) and Chandler (2002).

In its broadest, most "imperialistic" charter, the field is "concerned with everything that can be taken as a sign" (Eco, 1976, page 7), and is in principle the theoretical discipline that deals with sign and symbol systems. Is the present study not, therefore, simply an instance of semiotics?

There are several reasons, however, why we will not be overly adherent to the theoretical notions developed in semiotics. As we noted in the previous section, the notion of *sign* that we intend is much narrower than what semioticians usually consider. To further support the previous discussion on this point consider the entry for *sign* in Bouissac's *Encyclopedia of Semiotics*, which tells us that:

> In scholarly writing, the term *sign* might include, for example, words, sentences, marks on paper that represent words or sentences, computer programs … pictures, ideograms, graphs, chemical and physical formulas, fingerprints, ideas, concepts, mental images, sensations, money, postures and gestures, manners and customs, costumes, rules and values, the orienting dance of the honeybee, avian display, fishing lures, DNA, objects made of other signs … and also nonrepresentational objects (perhaps in music or mathematics) that have types of structure characteristic of other signs. (Bouissac, 1998, page 572)

Clearly the graphical symbols we are interested in represent a far narrower concept than what Bouissac lays out.

Second, semiotics provides no formal theory of the combination of signs in text— in other words the *syntax*, as laid out in the previous chapter. Look in any text

on semiotics, and one will be hard pressed to find any serious discussion of the combinatorics of signs even though every semiotician recognizes that this topic is important. Only in linguistics has the issue received serious scholarly attention. But this is also an important topic for us, since one of the interests here is in combinatoric sign systems that might be mistaken for (written) language; or that might *evolve* into written language under the right set of sociological, political and linguistic circumstances.

Syntax depends on formal mathematical models for its proper characterization and part of the problem with much recent work in semiotics is that, to put it bluntly, the field has been effectively hijacked by deconstructionists, who are not particularly interested in formal mathematical models. But even non-deconstructionist semioticians, such as Eco, are not typically well versed in mathematical models of information. Cf. the following (incorrect) definition by Eco: "[according to the mathematical theory of information] information is only the measure of the probability of an event within an equi-probable system" (Eco, 1976, page 42), which appeals to the Shannon notion of information (Shannon, 1951) which, however, does not depend upon equiprobability.

We will also not make much use of distinctions, due originally to Peirce, between *symbols*, *indices* and *icons*. Briefly, an *icon* is a (graphical or auditory) device that resembles its intended referent, an *index* is a device that somehow points out its intended referent (a canonical example is a pointing finger), and a *symbol* is something that has an arbitrary but conventional relation to its referent, like the word *cheese* in English, which in no way "resembles" cheese and only has meaning and form it does in Modern English by accident of history.

We will have more to say about icons specifically in the next section, but it should be pointed out that not all semioticians agree with the distinction. Eco (1976), page 178, notably does not, rejecting the distinction largely for reasons internal to his own theory of signs, though he also points to the fact that even iconicity is conventional, thus nullifying some of the basis for distinguishing symbols from icons.[1] For our purposes, the distinctions are largely irrelevant. As we shall see, many of the signs we consider would be considered iconic, and *all* early writing systems made use of iconic symbols. We do not wish to deny that one *may* subclassify signs in this way, just that we will not make much use of the distinctions.

In this work we will thus restrict ourselves to using the term *symbol*, possibly incurring the wrath of semioticians like Sebeok (2001), who notes (page 56) that

> 'symbol' is the most abused term of those under consideration here. In consequence, it has either tended to be grotesquely overburdened, or, on the contrary, reduced to more general kinds of behavioural phenomena, or even to absurd nullity.

But here we will give it a rather specific meaning: a man-made graphical form, or more generally a *set* of such forms (since any given symbol may have variant forms) that is itself a member of a set, termed a *symbol system*—in semiotic

---

[1] But see Sebeok (2001) for some arguments against Eco on this point.