In Between Worlds: The Art and Science of Interstellar Message Composition, Cambridge: The MIT Press, 2006, Douglas A. Vakoch, Ed.

# Tongue-Tied: The disadvantages of using natural language for interstellar message composition

**Kim Binsted**
**Information and Computer Sciences Department**
**University of Hawaii**
**binsted@hawaii.edu**

Our goal is to construct a message that another intelligence, completely unfamiliar with humankind, would be able to understand. It is assumed that the recipient is so far away that interaction in any reasonable time frame would be unfeasible, so the message must be entirely self-contained. Furthermore, since we know nothing about our target audience, we would like to keep our assumptions about their scientific knowledge, physiology, cognitive processes, and so on, to a minimum. If they have to be very human-like in order to understand the message, we have failed at our task. All we should demand of our audience is that they have the technical wherewithal to build equipment capable of receiving our message.

The problem is, we simply do not know how unusual humans are in the greater scheme of things. If the universe were full of intelligences differing from humans primarily in the shape of their foreheads, as in Star Trek, then the task of message composition would be much easier. Unfortunately, there is little reason to expect this to be the case. At the other extreme, if humans are but one data point in a huge range of possible physiologies, modes of communication, cultures, and so on, then it would be much more difficult – perhaps impossible – to compose a comprehensible message.

Even if the recipients were human, we could not necessarily expect them to be able to decode an arbitrary piece of English text. Historically, human archaeologists have had great difficulty decoding natural language texts written by other humans. In every case analogous to the interstellar decoding problem, to decipher the text they have relied on intermediary encodings (such as the Rosetta Stone, or modern versions of ancient languages) and a familiarity with human nature (e.g., that writings on a tomb are likely to be about the entombed person's life) (Finney & Bentley, 1998).

So, the message we send should not be arbitrary. Rather, it should be designed to be understood by an intelligent recipient with minimal knowledge of the sender. It has been argued that a large corpus of the written form of a natural language (English, French, or Chinese, for example) would serve this purpose. The main advantage to this approach is that there are large bodies of existing text that could serve as the bulk of the message. The question is, would it be understood? Heidmann (1993, p. 234), who proposed sending an encyclopedia, claimed that "the alphabetical coding [of English or French] can be deciphered using just a few pages, as well as the grammatical structures." We are not so optimistic. To the contrary: we believe that a natural-language text message, no matter how large and carefully composed it is, is extremely unlikely to be decodable by a non-human intelligence.

This conclusion might be counter-intuitive. After all, small children pick up natural languages with astonishing ease. Any human baby can learn any human language, barring brain damage or severe social deprivation. If a baby can do it, why can't an extraterrestrial intelligence? There are two reasons: specialization and interaction.

The evidence suggests that the ability to learn a human language is a specialized, evolved feature of the human brain, as opposed to being the result of the application of general-purpose intelligence.[1] Children succeed at the extremely difficult task of language acquisition, without formal lessons or consistent parental feedback, because they do not start as blank slates. Instead, they are born with the tools and structures necessary to acquire a human language. Thus, their task is more one of setting their internal parameters to match the particular language to which they are exposed, than one of learning that language 'from scratch,' as our recipient would need to do.

Secondly, children do not – they cannot - learn language merely by looking at written texts or by listening to recorded speech. Instead, they learn through direct linguistic interaction with other humans.[2] This interaction allows them to acquire vocabulary very quickly and efficiently, and to test and refine grammatical rules as they learn them. The recipient of our message, however, will be too far away to interact with us in real time, so will be in the position of a child plunked in front of a TV and expected to learn English.

It is possible that the range of human natural languages is coextensive with the range of natural languages for intelligent communicators everywhere. That is, Betelguesese and English might be no more different than, say, Mandarin Chinese and English. In this case, despite the problems described above, we might expect that a human could learn an extraterrestrial language as a second language, and vice versa. However, statistical analysis shows that, despite apparent surface differences, human languages are extremely similar to each other structurally (Elliot, 2002). That is, English and Mandarin Chinese are not, statistically speaking, all that different. Given the wide range of conditions under which a communicating species might evolve, and the influence of our particular conditions on the form of human languages, it would be surprising to find human-type communication to be typical of intelligent communication in general.

Humans have evolved to communicate primarily by inhaling a gas into their lungs, and exhaling it via specialized apparatus in their throats (vocal cords) and mouth (tongue, palate, teeth, and lips). This process produces sounds which, in Earth's atmosphere, are audible up to several hundred meters, depending on local conditions. These sounds are intercepted by specialized organs (ear, eardrum, tympanic nerve, etc.) on either side of the human head, which translate them into neural activation patterns that encode some variation on the intended content of the communication.

To aid communication over space and time, humans often transcribe these sounds into a binary image, i.e., writing. This binary image is composed of elements (characters) that sometimes represent sounds and sometimes meanings. In fact, some Chinese characters are composed of two halves, with each half conveying a different sort of information: the left half provides a clue to the meaning of the character, while the right half gives a hint about how to pronounce it. But the connection between a character and its sound or its meaning is sometimes fluid. The letter "a" can represent the [ah] or [ay] sounds, and the same character that represents a human in Japanese can be pronounced as either [hito] or [nin], amongst other possibilities. Similarly, the character "+" usually refers to the mathematical operation of addition, although its pronunciation varies widely over languages. Despite all this variation, the written forms of all human languages are linear, with characters strung one after the other, reflecting the linear nature of speech, which unfolds over time.

This complex system works for us because it has evolved as we have evolved, so that we, as a species, have the organs, brain functions, instincts, and cultural apparatus to handle it. There is no reason to expect our recipients to have followed a similar evolutionary path. For example, if their primary mode of communication is visual or olfactory, their language might best be represented in two or three dimensions, rather than linearly. If this is the case, then it might be impossible for us to adapt our linguistic apparatus to handle their language, and vice versa.

One might object, however, that our recipients are not children, but rather technologically-skilled adults, presumably with powerful analytical tools at their disposal. They do not, after all, have to become fluent in English – they merely need to be able to decode it, no matter how unnatural that decoding process might be to them.

However, there are two problems with natural language that even the best analysis is unlikely to overcome: grounding and ambiguity. The grounding problem refers to the difficulty of establishing a connection between a particular word and the concept to which it refers. Ambiguity is when the interpretation of the text cannot be determined from the text alone.

The grounding of words for concrete objects can arguably[3] be accomplished by providing a dictionary of images. For example, the entry for the word "dog" might be a set of photographs of dogs. Similarly, concepts from mathematics and, perhaps, physics can be grounded via reference to formulae and universal constants. However, this still leaves a large number of words ungrounded, including some we would probably want to use in our message ("peaceful," for example, or "cooperate"). Clever analysis by the recipient might reveal a great deal about the structure of our message, but without some way of tying meaning to words, the message would be content-free.

The problem of ambiguity is common to all natural languages, and can cause miscommunication even between skilled human language users. Humans are usually able to disambiguate sentences using their background knowledge about the world and about the probable intentions of the speaker. There are many different kinds of ambiguity, but here we will look briefly at syntactic ambiguity and referential ambiguity, both of which are common in English as well as other natural languages.

Syntactic ambiguity occurs whenever a text can be parsed grammatically in two or more ways. For example, the sentence "John put the ball on the table in his pocket" has two parses: either John took the ball from the table, and then put it in his pocket, or John took a ball, and then put it on a table in his pocket. Most humans will see only the first interpretation, because we know that tables do not typically fit into pockets. Unfortunately, our message recipient is unlikely to have access to such knowledge, and so will be unable to resolve the ambiguity.

Another very common kind of ambiguity is referential ambiguity, as seen, for example, in inter-sentential pronouns. The sentences "The ball was on the table. John put it in his pocket." are ambiguous, since "it" could refer to either the ball or the table. Again, it is impossible to resolve this ambiguity without some real world knowledge about the relative sizes of balls, tables, and pockets.[4]

The above examples are simple, and perhaps do not adequately convey the scale of the problem. Consider the following passage, taken more-or-less at random from a children's novel. Some (but not all!) of the background information that would be required to understand this passage is in the square brackets.

> 'Perhaps it [the dragon referred to in the previous paragraph] will go [in the sense of changing locations] away [from here, in the near future, without taking any other significant actions],' said Lucy.
>
> 'It'll [the situation will] be worse [than the situation is now] if it [the dragon] does [go away],' said Edmund, 'because then [if the dragon goes away] we [the individuals in the situation, not including the dragon] shan't [will not] know [in the sense of come to understand] where it is [the new location of the dragon, which is important information because the dragon is a threat]. If [hypothetical situation, analogous to the current situation]

> there's [there is] a wasp in the room [any room, not this particular room] I [Edmund] like to be able to see it [the wasp] [so that I, Edmund, can determine the location of the wasp, which is important information because the wasp is a threat].' (Lewis, 1952, p. 85)[5]

All natural languages are rife with ambiguities like those described above. Of course, we could remove some, if not all, of these problems by avoiding or removing the features that cause them. To reduce the ambiguities in English, we would have to eliminate variations based on sound (contractions, "an," etc.), make spelling uniform, remove most pronouns, and rewrite all those sentences which have ambiguous syntax, amongst many other changes. However, if we were to do this, the result would no longer be a natural language. It would be as artificial and as formal as mathematics, predicate logic, or any of the formal encoding schemes (e.g., Freudenthal, 1960) that have been proposed for interstellar message composition. Moreover, the main advantage of using natural language would be lost, since the existing texts that were to form the bulk of the message (e.g., encyclopedias) would have to be completely rewritten in the new, formalized language.

The reader who feels we have exaggerated the difficulty of the problem should recall that, despite decades of effort, human programmers have not yet been able to program a computer to understand a natural language in general, for the very reasons described above. The only successes have been in cases where the conversation is very specific in content and context, such as booking a ticket via a travel agent (Rudnicky et al., 1999). Why should we expect the recipient of our message to do better than human computational linguists, who are themselves already fluent in the language in question?

In summary, natural language only works for humans for two reasons. First, we have brains that are evolutionarily specialized for learning and using human languages. Second, we share vast amounts of common knowledge about the kinds of things we usually communicate about, i.e., the facts of life on Earth. This allows us to communicate very efficiently and expressively using natural language. Our recipient, however, could have a very different brain, and will almost certainly share very little of humanity's common knowledge.

All this bad news begs the question: If not natural language, what encoding <u>should</u> we use? Other chapters in this volume suggest various approaches, each with their advantages and disadvantages. Given how little can be known about the recipient, we would argue for a scattershot approach: send messages in every encoding we can think of, including natural language. For the reasons discussed above, we would expect that the message encoded in natural language would be understood late, if ever. Nonetheless, in the context of the messages in other, more formal encodings, the natural language text itself (as opposed to its content) might provide our recipient with insight into how the human mind works.

## Notes

1. For an excellent overview and analysis of how children acquire language, and what this reveals about the evolution of language and the brain, see Pinker (1994).
2. Ervin-Tripp (1973) conducted an interesting study of hearing children of deaf parents, whose only exposure to spoken language was from radio or TV. None of the children was able to acquire language from these inputs alone.
3. The idea of grounding concrete nouns via reference to pictures rests on a number of dubious assumptions. One assumption is that the recipient will be able to interpret two-dimensional images as representations of three-

dimensional objects. Another is that the recipient will be able to discern the correct central category from a set of images. For example, given the word "dog" and a set of five dog images, would the recipients understand that "dog" refers to the concept dog, or would they instead think it refers to four-legged animal or living being? It could even refer to those five dogs in particular, rather than the category of dogs in general. However, this debate is beyond the scope of this chapter.

4. Left out of this discussion, for simplicity, are several other possible sources of confusion. For example, if our recipients did not evolve in a gravitational well, the concept of one object being "on" another might be extremely difficult to convey.
5. The author begs forgiveness of fans of C.S. Lewis for mangling his prose. Think of it as a backhanded compliment, since the purpose of the exercise was to demonstrate that even the most clear and elegant natural language text contains ambiguities and depends on common knowledge for interpretation.

## References

Elliot, J. (2002, October). *The filtration of inter-galactic objets trouves and the identification of the lingua ex machina hierarchy*. Paper presented at the 53rd International Astronautical Congress, Houston, Texas.

Ervin-Tripp, S. (1973). Some strategies for the first two years. In T. E. Moore, ed., *Cognitive development and the acquisition of language.* New York: Academic Press.

Finney, B. and Bentley, J. (1998). A tale of two analogues: learning at a distance from the ancient Greeks and Maya and the problem of deciphering extraterrestrial radio transmissions. *Acta Astronautica, 42*(10-12), 691-696.

Freudenthal, H. (1960). *LINCOS: Design of a language for cosmic intercourse.* Amsterdam: North-Holland Publishing Company.

Heidmann, J. (1993). A reply from Earth: just send them the encyclopaedia. *Acta Astronautica, 29*(3), 233-205.

Lewis, C.S. (1952). *The Voyage of the Dawn Treader*. Harmondsworth, England: Puffin Books.

Pinker, S. (1994). *The language instinct.* New York: Morrow.

Rudnicky, A., Thayer, E., Constantinides, P., Tchou, C., Shern, R., Lenzo, K., Xu W., Oh, A. (1999). *Creating natural dialogs in the Carnegie Mellon Communicator system*. Proceedings of Eurospeech, 1999, 4, 1531-1534.