

We want to understand the relations these abilities bear to one another, in current cognition and evolutionarily:

- (1) our ability to comprehend *expressions* (word, phrases, sentences)
- (2) our ability to interpret *utterances*
- (3) our ability to interpret *communicative behaviour*
- (4) our ability to interpret *intentional behaviour*
- (5) our ability to *figure things out* (all-purpose reasoning; inference to the best explanation)

1 MINDREADING AND MODULARITY

The idea: *utterance* interpretation—whether we are dealing with linguistic or non-linguistic utterances—is a form of *mindreading*, construed as the capacity to infer the mental states of others (e.g. beliefs, desires, intentions) from their behaviour. Just an instance? Or a special case involving proprietary capacities? It is widely assumed that answering that question requires understanding the extent to which utterance interpretation is *modular*.

Fodor (1983). Our cognitive architecture contains dedicated *input systems* (e.g. the perceptual systems). On Fodor’s account they are *fast, mandatory, domain-specific, informationally encapsulated*. And they manifestly include a *linguistic comprehension* system. But does this yields representations of *sentence meanings*? Of *what is said*? Of *what is meant*?

Cosmides and Tooby (1994), and Pinker (1997) have provided an alternative, evolutionary perspective: Modules are “designed” by natural selection to solve some particular adaptive problem faced by hunter-gatherers in the Pleistocene—finding food, choosing a mate, avoiding danger. Sperber and Wilson find this more plausible. An utterance interpretation module is an *adaptation*.

2 RECURSION

Attitude Reports and Speech Act Reports

- (I) George believes (thinks, assumes, doubts, hopes) that *Hesperus is Phosphorus*
 George wonders whether *Hesperus is Phosphorus*
 George_i intends (wants/expects) *PRO_i to visit Phosphorus*
- (II) George said (stated, asserted, implied, implicated, argued, insisted,...) that *Hesperus was Phosphorus*
 George asked (inquired, discussed, disputed) whether *Hesperus was Phosphorus*
 George_i promised *PRO_i to visit Phosphorus*.
 1. S → S CONN S (CONN = connective, e.g. ‘and’, ‘or’, ‘if’, ‘when’, ‘whilst’, ‘because’)
 2. S → NP VP
 3. VP → IV (IV= intransitive verb, e.g. ‘snore’, ‘laugh’)
 4. VP → TV NP (TV= transitive verb, e.g. ‘love’, ‘respect’)
 5. VP → SV S (SV= sentential verb, e.g. ‘believe’, ‘say’, ‘implicate’)

Immediate recursion: In Rule 1, the LHS is reintroduced on RHS

Deferred recursion: In Rule 5, LHS is reintroduced by *another* rule, Rule 2, whose LHS is reintroduced on Rule 5’s RHS

- (1) [_S George snored] and [_S Mary laughed]
- (2) [_S [_{NP} George] [_{VP} snored]]
- (3) [_S [_{NP} George] [_{VP} [_{IV} snored]]]
- (4) [_S [_{NP} George] [_{VP} [_{TV} likes][_{NP} Mary]]]
- (5) [_S [_{NP} George] [_{VP} [_{SV} believes (that) [_S Mary snores]]]

In principle we can have as much recursion as we like:

- (6) Jane doubts that [George insists that [Mary thinks that [Paul complained that [Jane implicated that [Mary snores]]]]]]].

3 METAREPRESENTATION

“Metarepresentations are representations of representations, but not all representations of representations are metarepresentations in the relevant sense. The human metarepresentational capacity we are interested in here is, first and foremost, a capacity to represent the *content* of representations.” (Sperber, 2000, 117)

Sperber’s Hypotheses: “[1] Metarepresentational ability developed in the ancestral species for reasons having to do with *competition, exploitation, and co-operation* and not with *communication per se*....[2] A metarepresentational ability is plausible as an adaptation quite independently of communication....[3] a well-developed metarepresentational ability makes certain forms of communication possible quite independently from any code or language.” (117-123)

Sperber's Sequence of Scenarios

First scenario	Peter has a first-order metarepresentational belief : that Mary believes [1] that these berries are edible	
Second scenario	Mary has a first-order metarepresentational intention : that Peter should believe [1] that these berries are edible.	(An <i>informative</i> intention)
Third scenario	Peter has a second-order metarepresentational belief : that Mary intends [1] that he should believe [2] that these berries are edible.	
Fourth scenario	Mary has a third-order metarepresentational intention : that Peter should believe [1] that Mary intends [2] that he should believe [3] that these berries are edible.	(A <i>communicative</i> intention)

• So Mary's **first-order** communicative intention is a **second-order** informative intention, hence a **third-order** ordinary intention (which is why DS calls it a "third-order metarepresentation").

• So Mary has **two** informative intentions here: "a **first-order** informative intention that Peter should believe that the berries are edible, and a **second-order** informative intention that Peter should be aware of her **first-order** informative intention." (125)

• "What reasons might she have to have the second-order informative intention?...Peter's awareness of Mary's **first-order** informative intention, provided he trusts her, may give him an extra reason to believe her. In other words, *the fulfillment of the second-order informative intention may contribute to the fulfillment of the first-order informative intention.*" (125) This is essentially Sperber's explanation for why Grice and others have been drawn to the notorious "third clause". Yes, a full-blown *meaning* intention is often (typically?) present. But this does not mean the existence of an intention that rich needs to be built into the definition of the principal phenomenon, which S&W take to be perfectly well definable in terms of (their counterparts) of the first two clauses. Nifty.

Fifth scenario	Peter has a fourth-order metarepresentational belief : that Mary intends [1] that he should believe [2] that she intends [3] that he should believe [4] that these berries are edible.	
----------------	---	--

"Peter might come to have this belief when he notes that Mary is ostensibly making sure that he is paying attention to her behavior by, say, establishing eye contact with him, picking the berries in somewhat formal manner, and so forth. Mary's first-order informative intention is now an "overt" (Strawson, 1964) or "mutually manifest" (Sperber & Wilson, 1995). We have reached a level where **communication proper** occurs, though **no code or language** is involved." (125)

"Once the level of metarepresentational sophistication of our fifth scenario is reached, a dramatic change indeed occurs. Whereas before, in order to fulfil her first-level informative intention, Mary had to engage in behavior - picking the berries - that was best explained by attributing to her the belief that the berries were edible and the desire to collect the berries, she can now resort to **symbolic** behavior, the best explanation of which is simply that she is trying to fulfil an informative intention, her desire to inform (or misinform) Peter that the berries are edible. Instead of actually picking the berries, she might, for instance, **mime** the action of eating the berries.

Typically, **symbolic** behavior such as *miming* has no plausible explanation other than that it is intended to affect the beliefs or desires of an audience. This generally is the one effect of such a behavior that could clearly have been intended by the agent. Thus, for Mary to **mime** eating the berries does not feed her, does not help her in any easily imaginable way, except through the effect it has on Peter. Her **miming** behavior triggers in Peter's mind, through the perception of a **resemblance** between the **miming** behavior and the behavior mimed, the idea of eating the berries. Moreover, if it is at all relevant to Peter to know whether or not the berries are edible, Mary's behavior suggests that they are. Provided that Peter sees this **miming** behavior as intended by Mary to inform him of an informative intention of hers, he is justified in assuming that the very idea triggered in his mind was one Mary wanted him to entertain, elaborate and accept.

The same result achieved by miming could, provided Peter and Mary shared an appropriate code, be achieved by means of a coded signal or by means of some combination of iconic and coded behavior." (125)

4 RELEVANCE AND THE COGNITIVE BASIS OF INTERPRETATION

Goal: A Theory of Communication that

- a. accommodates the many different *varieties* and *purposes* of communication;
- b. acknowledges that communication takes place *at a risk*;
- c. explains how communication can often *succeed*; and
- d. is compatible with what is known about *human cognition*.

Speaker's meaning amounts to is a complex *mental state* involving:

An informative intention: An intention to inform the audience of something.

A communicative intention: An intention that the audience should recognise the informative intention.

Fundamental problems for human cognition

1. Our **senses** monitor too much information for us to be able to attend to it all.
2. Our **memories** contain too much information for us to be able to retrieve and use it all.
3. Our **inferential** systems can draw too many conclusions for us to attend to them all.

Claims made by relevance theory

Human cognition is *relevance-oriented*: we pay attention to information that seems relevant to us. Our cognitive systems are so constructed that:

1. Our **perceptual systems** tend to pick out the potentially most relevant information.
2. Our **memories** tend to supply contextual assumptions most likely to contribute to relevance.
3. Our **inferential systems** tend to maximise relevance by deriving greatest cognitive effects.

Cognitive Principle of Relevance

Human cognition tends to be geared to the maximisation of relevance.

(This governs *all* types of information-transmission.)

Accidental. When someone speaks, we pay attention to any information we can pick up that seems relevant to us, whether derived from the content of their utterance, their facial expressions and gestures, their accompanying behaviour, their pauses, hesitations, and so on; and process this information in a context that is likely to maximise its relevance. This sheds some light on *accidental* information transmission.

Covert. Someone who wants to manipulate us can covertly exploit these basic facts about human cognition, by planting clues that our cognitive systems are likely to pick up. This sheds some light on *covert* information transmission. And it is against this cognitive background that overt intentional communication takes place.

Overt. What is unique here is that, approaching an utterance or other piece of communicative behaviour addressed to us, we can have not just hopes but reasonable expectations of relevance.

Communicative Principle of Relevance

Utterances (or other ostensive stimuli) create presumptions/expectations of relevance.

The Communicative Principle of Relevance is *not* a Gricean maxim. It can't be violated, and doesn't have to be learned: it follows from basic assumptions about human cognition. Speakers can't help creating expectations of relevance. (Of course, the expectations may be disappointed.)

Co-operation is not obviously necessary for communication

Does communication really have to have an 'accepted purpose or direction' to which speakers are expected to contribute? Grice (1989: 368-72) himself considers hostile or adversarial cases, casual chat, etc. as possible counterexamples.

Speakers are not really expected to 'be as informative as is required'

Grice suggests that speakers should be 'as informative as is required', even if they don't have the information to give, or if it would go against their interests to give it. Relevance theory claims that speakers are not expected to give required information if they are *unable* or *unwilling* to do so.