

Does Thought Require Language? Part II: The Language of Thought

Last week we examined one aspect of the question, “Does Thought Require Language?”, focusing on whether nonlinguistic animals can think. This week, we’ll look at a very different version of the idea that thought requires language: specifically, the idea that *thought itself takes place in a language*.

1. The Representational Theory of Mind

Propositional Attitudes = mental attitudes towards propositions (e.g., *belief, fear, imagining, hope, desire*).

Big Question: What does it take to have a propositional attitude?

- For example, suppose Annabelle believes that Hong Kong is north of Singapore. What does that involve? What does it take for her to have this belief?
- Or suppose Cassandra hopes that she’s win the lottery. What does this involve? What does it take for her to have this hope?

Representational Theory of Mind (RTM):

S has some propositional attitude towards a proposition p if and only if:

- i) S stands in some suitable relation towards a mental representation R ,
- ii) R means that p .

Examples:

- Annabelle’s belief that Hong Kong is north of Singapore involves two things: (i) it involves her having a mental representation R whose content is the proposition, \langle Hong Kong is north of Singapore \rangle . (ii) R needs to be in her “belief box”: it needs to play the belief-role in her psychology.
- For Cassandra to hope that she’ll win the lottery likewise involves two things: (i) it involves her having a mental representation R^* whose content is the proposition, \langle Cassandra will win the lottery \rangle . (ii) R^* needs to be in her “hope box”: it needs to play the hope-role in her psychology.

RTM breaks down the Big Question into two sub-questions:

Q1: What does it take for S to stand in a suitable relation to a mental representation?

Q2: What does take for a mental representation to mean that p – i.e., to express a certain proposition?

In answer to Q1, many contemporary philosophers of mind appeal to *functionalism*. The core idea behind functionalism is that what makes a particular mental state qualify as a belief rather than, say, a hope is the state’s *functional role*.

For example, consider the following simple functional role theory of belief:

Input condition: Typically, if it perceptually appears to you that p , you'll come to believe p

Output condition: Typically, if you believe p , then you'll be disposed to act as if p is true.

Combined with RTM, the idea would be that what makes a mental representation R (with content p) count as a belief, rather than, say, a hope – is that R satisfies the input and output conditions associated with belief. Hope will have its functional role – its own input and output conditions – that distinguishes it from belief.

2. The Language of Thought Hypothesis

What about Q2? One partial answer is provided by the Language of Thought Hypothesis (LOT):

LOT (Rough Version): Mental representations are linguistic – they represent the world in the same way that language does.

To flesh this out, we need to say what it means for a representation to be linguistic: how does language represent the world?

Some Elements of Linguistic Representation

- *Complexity:* Linguistic representations are often complex, in the sense that they are made out of more basic elements.
 - For example, the sentence:

(1) The dog chased the lizard.

is made out of the words, “the” “dog”, “chased”, “lizard”.

- *Transportability:* The basic elements are transportable, in the sense that they can be rearranged.
 - For example, we could rearrange the elements in (1) to get the sentence:

(2) The lizard chased the dog.

- *Compositionality:* The meanings of complex representations are determined by the meanings of the basic elements, together with their mode of combination (i.e., the manner in which they are arranged).
 - For example, the meaning of (1) is determined by the meanings of “the”, “dog”, “chased”, “lizard” (similarly for (2)).

The Language of Thought Hypothesis, in more detail

According to LOT, mental representations do not necessarily occur in any natural, spoken language (e.g., English, Chinese, Thai, etc.). Rather, they occur in “Mentalese.” But Mentalese is still linguistic, in the sense that it has all of the key properties of natural languages. Specifically:

LOT (More Detailed Version): Mental representations exhibit complexity, transportability, and compositionality. In these respects, they are just like linguistic representations.

Example: Suppose Seng Leong believes: *The dog chased the lizard.* According to LOT, this mental representation is basically a sentence in his head. Specifically:

- It is built out of simpler parts – namely, the mental representations for *the dog*, *chased*, *the lizard*.
- These parts are transportable, meaning that they could be rearranged to create the mental representation, *The lizard chased the dog.*
- The complex mental representation is compositionally determined by the meanings of its parts, together with their mode of combination.

3. Arguments for LOT

Why believes LOT? Fodor, one of the main proponents of LOT, offers two main arguments.

Argument 1: The Argument from Productivity

Why think that linguistic representation has the various properties we’ve ascribed to it – specifically, complexity, transportability, and compositionality?

One main reason is that language is *productive*, in the following sense:

Productivity of Language: A competent speaker of a language has the ability to produce and understand a potentially infinite variety of new sentences.

For example, you’ve probably never encountered the following sentence before:

- (3) The sleepy rhinoceros ate a jar of peanut butter, and then jumped over the Marina Bay Sands.

But you have no trouble understanding it.

This is something we taken for granted, but it’s worth taking a step back and noticing how remarkable it is. What makes productivity possible? That is, what enables us to understand a potentially infinite variety of new sentences?

Explanation: Complexity, transportability, and compositionality! (3) is a complex representation, whose meaning is determined by the meanings of its parts (its individual words), together with their mode of combination. Speakers

learn the meanings of the words, together with their mode of combination, and this enables them to learn the meaning of the whole.

Now, Fodor thinks the exact same argument extends to *mental representation*. In other words, thought is productive:

Productivity of Thought: A normal thinker has the ability to think a potentially infinite variety of new *thoughts*.

- For example, a normal thinker can think the thought, *The sleepy rhinoceros ate a jar of peanut butter, and then jumped over the Marina Bay Sands*, even if they have never thought this before.

What explains this? According to Fodor, we should give the same answer here that we gave for language.

Explanation: Complexity, transportability, and compositionality! Thoughts (mental representations) are often complex; when they are, their meanings are determined by the meanings of the individual elements, together with their mode of combination. A thinker masters the meanings of the individual elements, together with their mode of combination, and this enables them to grasp the meaning of the whole.

Question: Do you find this a compelling argument?

Argument 2: The argument from Systematicity

A related, but distinct, argument comes from the phenomenon of systematicity:

Systematicity in Language: If some sentence of the form “*aRb*” is meaningful in a language, then so is “*bRa*.”

- We’ve already seen an example of this: sentence (1) “The dog chased the lizard” is meaningful. And so is (2) “The lizard chased the dog”.

Here too, this is something we take for granted. But what explains it? Once again the answer seems to be:

Explanation: Complexity, transportability, and compositionality! (1) is made up out of simpler elements. Since these elements are transportable, they can be rearranged to get (2). By compositionality, (2) is guaranteed to be meaningful as well.

Here too, Fodor argues that systematicity applies not just to language, but also to thought.

Systematicity in Thought: If someone is capable of thinking something of the form, *aRb*, they are also capable of thinking something of the form, *bRa*.

- For example, if someone is capable of thinking, *the dog chased the lizard*, they are also capable of thinking, *the lizard chased the dog*.

And here too, Fodor thinks we should explain the systematicity in thought in the same way we explain the systematicity in language:

Explanation: The complex mental representation, *The dog chased the lizard*, is made up of elements – e.g., the mental representation of the *dog*, the mental representation of *chasing*, and the mental representation of *the lizard*. And these elements are transportable, so they can be rearranged to create a different representation such as *the lizard chased the dog*. Moreover, if mental representation is compositional, this rearrangement is guaranteed to be meaningful as well.

Argument 3: The Explanatory Argument

Beliefs have causal effects. For example, believing that it's raining outside will have certain typical effects on the agent's behavior. Likewise, believing that there's beer in the fridge will have certain typical effects.

Now, what about the belief that it's raining outside *and* there's beer in the fridge? This too will have characteristic effects, which will be similar to both the effects of the belief that it's raining, and to the effects of the belief that there's beer in the fridge.

What explains this similarity? Fodor takes the following to be a plausible methodological principle:

Methodological Principle: from similar effects, infer similar causes.

LOT seems to do this: it says that the belief that it's raining outside and there's beer in the fridge is composed of both the belief that it's raining outside and the belief that there's beer in the fridge.

4. Objections to LOT

Incompleteness Objection

LOT explains the meanings of complex mental representations in terms of the meanings of basic elements (the atoms). So, for example, the meaning of *Jane enjoys philosophy* is explained in terms of the meanings of the Mentalese atoms, *Jane*, *enjoys*, and *philosophy*. But LOT doesn't explain the meanings of the atoms. Why is it that the Mentalese atom *Jane* refers to Jane, rather than someone else? Why is it that the Mentalese atom *enjoys* refers to enjoyment, rather than hatred (or water, or dogs, etc.)? Without some answer to this, we won't have fully answer Q2: we won't have given a full explanation of what it takes for a mental representation to be about something in the world.

Standard Reply:

The standard reply to this objection is to concede that LOT is incomplete: it provides only a partial answer to Q2. To give a complete answer, it would need to be supplemented with some story about how the atoms succeed in representing what they do.

This issue will become important later in the class, when we discuss theories of mental content. While these theories do not necessarily presuppose LOT, proponents of LOT can – and often do – appeal to these theories to try to provide the missing account of the meanings of the atoms.

An Alternatives to LOT? The Map Hypothesis

Some authors (such as Braddon-Mitchell and Jackson) suggest that an alternative to LOT is to think that mental representation has a *map-like* structure.

Map Hypothesis: Mental representations represent the world in a way that is more similar to maps than to sentences.

Note that the map hypothesis doesn't say that we literally have maps in our heads. The idea is that thoughts have the abstract structure of maps rather than sentences.

So, what are the differences between maps and sentences?

Some candidates (Braddon-Mitchell and Jackson):

- Maps are *informationally rich* – If a map tells you what the largest island is, it only does this by also telling you the sizes of all of the other islands. In other words, maps can't tell you a little bit of information without telling you a lot of information. In this regard, they seem to differ from sentences – we can have sentence that says, “Borneo is larger than Singapore” without providing any other information.
- There is no natural way of dividing a map at its truth-assessable joints – that is, no natural *minimal* unit of truth-assessable information. By contrast, sentences have a natural minimal unit of truth-assessable information – namely, sentences!

Proponents of the map hypothesis argue as follows:

- 1) Maps represent the world in importantly different ways than sentences.
- 2) The map hypothesis is a genuine alternative to LOT. (from 1)
- 3) At the same time, maps, like sentences, are *productive* and *systematic*.
- 4) So, the map hypothesis provides an equally good alternative explanation for the productivity and systematicity of thought as LOT does. (from 2, 3)

Question: Is this a compelling objection to the Language of Thought Hypothesis?

Things to think about:

- Do you agree that there are important differences between the way maps represent the world and the ways sentences do?
- Is it true that maps are productive and systematic?
- Are maps and sentences expressively equivalent? That is, if a sentence can represent something, can a map represent it equally well, and *vice versa*?