Naturalising representational content
Advanced Topics in MLEC — Week 1

Mark Sprevak

University of Edinburgh

`mark.sprevak@ed.ac.uk`

Reading for this week

- Paper that Shea wrote before his book

- *Philosophy Compass* specialises in useful intros

Naturalising Representational Content

Nicholas Shea*
King's College London

Abstract

This paper sets out a view about the explanatory role of representational content and advocates one approach to naturalising content – to giving a naturalistic account of what makes an entity a representation and in virtue of what it has the content it does. It argues for pluralism about the metaphysics of content and suggests that a good strategy is to ask the content question with respect to a variety of predictively successful information processing models in experimental psychology and cognitive neuroscience; and hence that data from psychology and cognitive neuroscience should play a greater role in theorising about the nature of content. Finally, the contours of the view are illustrated by drawing out and defending a surprising consequence: that individuation of vehicles of content is partly externalist.

The contents list is as follows:

1 The Project of Naturalising Representational Content
2 The Explanatory Role of Content
3 Existing Theories
4 Pluralism
5 Externalist Syntax
6 Conclusion

1. The Project of Naturalising Representational Content

Some things in the world have semantic properties. Spoken and written sentences are paradigm cases. They are perfectly ordinary particulars in the causal order: ink marks on the page and vibrations in the air. But they also have more exotic properties: they can be true or false, or, in the case of imperatives, they can be satisfied, or go unsatisfied. That is, they are associated with a condition, and it makes an important explanatory difference whether that condition actually obtains, comes to obtain, or fails to obtain. Thoughts too have semantic properties in this way, and on all naturalistic views of the mind the semantic properties of thoughts are also associated with ordinary physical particulars: either whole people, or things going on within people, almost certainly involving processes in the brain.

In the 19th Century Franz Brentano identified the closely-related idea of intentionality and argued that it is a peculiar feature of thoughts (Brentano 1874/1995). Thoughts can be about objects and properties that are not present to the thinker, that are distant in time and space, that are hypothetical or may only be actualised far in the future, or that are entirely imaginary. If thinking is a physical process realised by or within people, how then can thoughts reach out and be about such things? Indeed, even when the object is right in front of me, how does my perception and thought manage to be about it, when the object is out there and the thought is in here (in some metaphorical sense)? A perfectly ordinary feature of everyday life – the fact that sentences and thoughts have

Main question of the paper

How should we naturalise representational content?

- Focus on naturalising mental representation

- Big unsolved problem in philosophy

- Previews the approach that Shea's book later adopts

Section 1

Some useful background

Watch my videos on the problem of representation!!

video 1   video 2   video 3   video 4

What is representational content?

Representations have a vehicle and a content

- Vehicle : concrete physical particular that *'is'* the representation

- Content : what the representation is *'about'*

Content is sometimes also called:

- semantic content, intentional content, representational content, mental content, encoded content, meaning, distal content, . . .

    NB. Some authors use these terms in different ways

What does it mean to *naturalise* content?

How do you get representational content from a world of non-semantic, non-mental, non-normative facts?

Explain how representations with content arise out of

purely physical ingredients

Why is the naturalising project *important*?

Representations are pretty puzzling in themselves . . .

. . . but the issue gets a special edge from the mind–body problem

Mind–body problem

How does our mental life arise from purely physical ingredients?

Chunk of our mental life involves mental representations

How does this aspect of the mind arise from purely physical ingredients?

Conventional representations & mental representations

Representations come in 2 kinds:

1. Conventional representations (e.g. words, diagrams, road signs)

2. Mental representations (e.g. beliefs, mental maps, percepts)

End goal is to naturalise them all

In this course, we will just focus on mental representations

The 'mental representation is fundamental' view

Conventional representations $\rightarrow$ Mental representations $\rightarrow$ Physical facts

Lewis (1969) & Grice (1957)'s 'intention-first' model of content:

- Conventional representations gain their content because of our mental intentions

- Fundamental challenge is to explain how these mental representations arise

NB. Skyrms (2010) tries to naturalise public language content directly, without appealing to mental representation. Hutto & Myin (2013) try to explain mental representation in terms of public conventional representations.

Subpersonal mental representations

Don't focus on personal-level mental representations (beliefs, desires, thoughts, percepts)

Focus on subpersonal mental representations (e.g. representations inside the early visual and motor systems, or in unconscious decision making)

- Not introspectively accessible

- Often not conscious

- Studied largely from 3rd-person point of view

- Vehicles would be single neurons, populations of neurons, computational states inside the brain

Realism about X

If you are a 'realist' about X then – roughly speaking – you think that Xs are really 'out there'.

- Xs exist as concrete particulars independently of our views about them

- We should employ X-talk, X-concepts, and practices associated with X in science

Realism about mental representations

If you are a realist about mental representations then you think that:

- Mental representations exist as concrete particulars (realised inside the brain) and really have specific contents (independently of how we interpret them)

- Science should appeal to those mental representations to, amongst other things, explain aspects of cognition and behaviour

Let's try to be realists about mental representation

What do we have to do?

1. Say something about which concrete particulars the `vehicles` are

2. Say something about how they get their `content`

Section 2

Theories of representational content

4 ideas about how to naturalise content

1. Representation arises from covariation

2. Representation arises from inferential role

3. Representation arises from structural isomorphism

4. Representation arises from natural functions

Idea 1: Covariation as the secret ingredient?

- Hubel and Wiesel's cat experiment

- The 'cow detector' thought experiment

- Dretske's *Knowledge and the Flow of Information* (1981)

Basic intuition behind the covariation theory

A vehicle X represents content Y if X tends to occur when Y is present, but does not tend to occur otherwise

(in such a case, vehicle X 'carries information about' content Y)

Problems for covariation theories

1. Covariation is everywhere, representation is not

2. Only applies to early sensory representations

3. Covariation with proximal or distal stimuli?

4. How is misrepresentation possible?

And lots more . . . !

Idea 2: Inferential role as the key ingredient?

- Think about the meaning of logical terms like 'and', 'not', 'or'

- Think about the meaning of theoretical terms like 'mass', 'force'

- Block's 'Advertisement for a Semantics for Psychology' (1986)

Basic intuition behind the inferential-role theory

A vehicle X represents content Y if, during inference, vehicle X plays the role associated with that content

(also known as conceptual role, inferential role, or procedural semantics)

Problems for inferential-role theories

1. Symbol grounding problem (Harnad, 1990): How does one connect vehicles to the non-linguistic world?

2. Out of all roles a vehicle plays, which ones confer content?

3. Fodor & Lepore (2002)'s dilemma: inferential-role semantics entails either implausible holism or commitment to idiosyncratic definitions

And lots more . . . !

Idea 3: Structural isomorphism as the key ingredient?

- Think about how a London Underground map represents

- Navigation maps in the hippocampus of rats ('place cells')

- Cummins's *Meaning and mental representation* (1989)

Basic intuition behind the structural-isomorphism theory

A vehicle X represents content Y if a structure-preserving mapping (isomorphism) exists between states of the organism and their relations and states of the environment and their relations, and according to that mapping scheme, $X \mapsto Y$.

Problems for structural-isomorphism theories

1. Isomorphism is symmetrical, but representation is not

2. Newman's objection to Russell: isomorphism are trivial to find

3. Some representations are not isomorphic to their intended content (e.g. France as a hexagon)

And lots more ...!

Idea 4: Natural function as the key ingredient?

- An evolved 'cow detector'

- Consider alarm calls in vervet monkeys

- Millikan's *Language, Thought and Other Biological Categories* (1984)

Basic intuition behind the natural-function theory

A vehicle X represents content Y if a 'consumer' system inside the cognitive agent has the natural function of 'using X as a representation of Y'.

'Natural function' should be understood in evolutionary terms (treating X as a representation of Y aided past survival, reproduction)

Problems for natural-function theories

1. Swampman (Davidson, 1987)

2. Many mental representations don't have any clear evolutionary role
   (e.g. *Paris is the capital of France*)

3. Content often left underdetermined (*snake!* or *danger on ground!*)

4. Discrete consumer systems are hard to find in the brain

And lots more . . . !

Section 3

Shea's approach

Pluralism

- 'Pluralism' means that there is not one single answer to the question of how representational content gets determined

- Different ingredients (1–4) play greater or lesser roles for different representations

- Mental representation should be naturalised in different ways in different cases

What should we do?

1. Identify a specific behaviour in relation to an environment

2. Look for a scientific explanation of that behaviour that appeals to representation of that environment

3. Analyse whether that explanation really succeeds (and whether it really relies on representation)

4. Identify which kinds of physical properties are needed in order for the explanation to succeed in explaining the behaviour

   ▶ covariation, isomorphism, inferential role, evolutionary natural functioning, . . .

   Those are the properties that naturalise that representation