**W6: Structural Representation**

Mental Iconicity • Gabriel Greenberg

2.10.26

Quotes from Shea 2014

**Orienting questions**

What is it for a given neural state to be used *as* a representation?

What is it for a given neural representation to be used *as* a picture, map, or iconic representation?

Part of our build up to retinotopy:

W6: when is a representation used iconically?

W7: not all perceptual representations are used iconically.

W8: overview of visual neuroscience.

W9: the debate: are retinotopic areas used iconically?

W10: my argument: retinotopic areas *are* used iconically.

**Some Key Assumptions of Computational Theory of Mind**

**Computationalism**

Cognitive processes consist of inferential transitions between representational states.

Representational states are realized physical states of the brain that function to carry information.

Cognitive processes are realized by computational subsystems of the brain that function to process information.

**Two central motifs of computational systems:**

Computation sandwich:

input representation → computation → output representation

Representation sandwich:

producer computation → representation → consumer computation

**Use-criteria for representation**

Properties of a neural state can be representational only if they are used by "downstream" computation process (i.e. a consumer computation).

== Properties of a neural state can be representational only if downstream computation is sensitive to (responsive to) these processes.

Contrast: spike rate vs. color

## Case studies

The honey bee “dance”

Place cells in the hippocampus

Coactivation sequences

Hypothetical: measure distance by measuring duration of coactivation sequences

## Shea on structural representations

Shea’s take on iconicity.

“**Structural Representation:** A collection of representations in which a relation between representational vehicles represents a relation between the entities they represent.”

Maps are primary example.

**Shea’s hypothesis:** structural representations in the brain make use of exploitable isomorphisms with systems of relations in the external world.

**General position.** “Many kinds of representational content depend in part on there being an exploitable relation between representations and the things they represent, a relation that is made use of when the system relies on those representations in generating behaviour.”

## Shea on exploitable isomorphism

“*Exploitable Isomorphism:* (i) A one-to-one correspondence mapping between a set of putative representational vehicles in an organism or other system, and a set of entities to be represented, (ii) in which a projectable relation on the vehicles to which operations of the system could be sensitive, (iii)corresponds to (iv) a relation, of significance to the system, on the entities to be represented.”

## Isomorphism, defined

Relational structure

1-1 Mapping

<A,R> is isomorphic to <B,R’> iff there is 1-1 function f s.t. for all x,y in A, for all x’,y’ in B: R(x,y) iff R’(f(x),f(y)).

Isomorphism is matching of relational structures.

## Shea on exploiting a relation

**Make a difference to downstream computation. “**A necessary condition for an isomorphism to be exploited is that relations between putative representations should make a difference to downstream computations or behaviour. An *exploited isomorphism* is one where such a relation is made use of in virtue of its correspondence to a relation between the entities represented.”

**Case of the hypothetical route analyzer.** “The relation between place cell firing acts as input to a computational process of choosing the shorter chain, whose utility lies in the fact that co-activation relations over the place cells correspond to spatial relations over locations. So in this case the correspondence between co-activation structure and spatial structure is being exploited. It is also a plausible case of structural representation.”

**Shea on exploitable correlation**

“Another kind of exploitable relation is correlation. The simplest kinds of exploitable correlations concern variable matters of fact.”

The example of the Fly Detector.

The example of the Bee Dance.

**Limitations of Shea’s analysis**

**2nd order vs. 1st order iconicity.**

**Isomorphism vs. natural relations.**

**Exploiting a structural relation vs. exploiting an isomorphism**

**Extension 1: structural representation → iconicity**

Let’s broaden the scope from structural representation to iconicity generally.

Structural representations are 2nd-order iconic representations.

The same analysis of representational “use” should also extend to 1st-order iconic representations.

Shea’s analysis seems to collapse the distinction between 1st-order iconic representation and 1st-order symbolic representation.

See e.g. his discussion of the bee-dance

**Extension 2: isomorphism → natural relations**

We should move beyond isomorphism-based analysis in two ways.

First, there are 2nd order representations that can not. be defined in terms of isomorphism alone– e.g. projection relations, scale transformations

Second, not all forms of iconicity can be defined in terms of isomorphism (or any other relation between relational structures.

We should replace isomorphism with a more general notion of a “natural relation”

**Extension 3: exploiting a relation to the environment**

Shea’s definition of relation exploitation is almost exclusively focused on the structural relations— i.e. the least interesting part of his analysis!

We need a theory of what it is for a computational system to *use a relation to the environment*.

This isn’t trivial, since, by definition, computational mechanisms aren’t themselves sensitive to relations to the environment

**Reliable algorithm**

An algorithm reliably performs an **epistemic function.**

The standard case: **preserving truth.**

**Examples:**

The MP (modus ponens) algorithm reliably preserves truth.

The Bayesian update algorithm reliably tracks rational probability.

For an **algorithm to be reliable**, there must be…

An **environmental regularity** that it is tracking.

A **systematic semantic relation** between inputs/outputs of the algorithm and inputs/outputs of the environmental regularity. (“Semantics”)

**(Towards a definition of) Exploiting an external relation**

Compare: what is it for a *person* to exploit an external relation?

E.g. A uses the isomorphism relation between the map and the geography to successfully make navigational decisions.

**Definition of agent exploiting a relation:** Where *r* is a representation and *e* is a situation in the environment: an agent A *exploits* relation R between *i* and *e* iff

(1) A assumes R(*i*, e);

(2) R(*i, e*) holds;

(3) (1) and (2) promote the realization of A’s goals.

**Definition of agent exploiting a relation:** A system S exploits an external relation R iff

R must hold in order for S to reliably perform its epistemic function.

Put another: S exploits R iff the *justification* for why we implement S assumes R.

Making an implicit assumption

Following Johnson on implicit bias…