Mental Iconicity Seminar

Prof. Gabriel Greenberg

W4: Map-like Thought

1.27.25

Question I: is the Fodorian argument for a language of thought compatible with a cartography of thoughts?

Camp reconstructs the argument for LOT

The classic reason for thinking that thought must be language-like is that only this assumption can explain or justify the systematicity of thought. (146)

> There are systematic relations among the contents that a thinker can represent and reason about.
>
> Systematic relations in content must be reflected by correlative structure in a thinker's representational and reasoning abilities.
>
> Structured representational abilities require a system of representational vehicles which are composed of recurring discrete parts combined according to systematic rules.
>
> Any system of representational vehicles composed of recurring discrete parts combined according to systematic rules is a language.
>
> Therefore: there must be a language of thought.

Spoiler: Camp denies Premise 4.

Camp replies: maps as systematic, productive, compositional

> Maps as counterexamples to Premise 4: composed of recurring discrete parts combined according to systematic rules is a language.
>
> "Maps are clearly constructed out of recurrent formal elements that make a common semantic contribution each time they occur: for instance, on many maps any solid line of a certain width signifies a street, any blue line or blob signifies a river or lake, and any cross signifies a church. Further, the representational import of the entire map is a systematic function of the way in which those elements are combined: if two lines intersect, with a blob in one quadrant and a cross in the other (Figure 1a), then this represents two intersecting streets with a church across from a pond. By contrast, if the two lines are drawn in parallel, with the cross above the blob (Figure 1b), then these same elements represent a different but related situation, in which a church is north of a pond and between two parallel roads." (154)

Weak vs. Strong LOT

Better tests for LOT?

Maps vs. Language

Toy Isomorphism Semantics for Maps:

Assume an assignment [[X]] from markers to object or object types.

Map M is accurate at geography G iff For every pair of markers m1, m2 in M: distance(m1,m2) in M = distance ([[m1]], [[m2]]) in G.

Maps:

Arbitrary first-order rules: "The semantic principle which maps those constituents to objects and properties in the world can be quite indirect and arbitrary. Road maps often represent churches with a cross, four-lane highways with a red line, state capitals with a star, and cities by their names."

Non-arbitrary first-order rules: "In particular, the physical features of the icons themselves often reflect salient physical features of the objects or properties being represented. Thus, a straight line represents a straight street and a crooked line a crooked street; a blue blob represents a pond of that very shape and exploits similarity of color to indicate that it's water; and a green blob represents a park of that very shape and exploits similarity in color to indicate that it's filled with vegetation."

Isomorphic second-order rules: "Maps represent by exploiting isomorphisms between the physical properties of vehicle and content. But maps abstract away from much of the detail that encumbers pictorial systems. … maps only exploit an isomorphism of spatial structure: on most maps, distance in the vehicle corresponds, up to a scaling factor, to distance in the world."

Language:

Arbitrary first-order rules: "The semantic principles mapping the vehicle's constituents to represented contents are clearly highly arbitrary and conventional:"

Non-isomorphic second-order rules: "The syntactic principles combining those constituents are less arbitrary, but they too clearly abandon any appeal to physical isomorphism. Instead, some sort of functional relation among syntactic constituents maps onto some sort of logical or metaphysical relation among the semantic values of those constituents."

Atomism: "Because sentential systems represent by combining discrete, conventional symbols in an abstract structure, they are highly digital: they deliver chunks of information about discrete states of affairs."

Moral: the deep and consistent contrast here is really at the level of *how atomic parts are combined* (i.e. 2nd order structure).

The Conjunction Test (aka semantic holism)

> "Maps automatically conjoin information about the spatial locations of all the objects and properties they represent."
>
> I.e. if you represent R(A,B) and R(B,C) in a map, then you also represent R(A,C).
>
> Suppose I have the following sentences specifying the locations of Bob, Ted, and Alice:
>
> Bob is at the grocery store at 10th and South.
>
> Alice is at the cafe? at 11th and Pine.
>
> Ted is at the park at 9th and Spruce.
>
> 9th Street is east of 10th Street.
>
> 10th Street is east of 11th Street.
>
> Lombard Street is north of South Street.
>
> Pine Street is south of Spruce Street.
>
> Lombard Street is between South and Pine Streets. (161)
>
> **Ask:** how many blocks is Bob from Alice?
>
> Contrast the process of constructing this list with that of constructing the following map:
>
> **Ask:** how many blocks is Bob from Alice?
>
> **Camp's explanation.** "No single, syntactically isolated portion of the map represents just where Bob is, without also representing Bob's location relative to Ted and Alice and everything else that is represented on the map. Each icon contributes to the overall spatial configuration, and the location of each object and property is given in terms of this overall configuration. As a result, any alteration in the location of the 'Bob' icon automatically alters the represented relations between Bob and everyone else."

Question II: what format are we committed to when we posit maps in thought?

Literal maps, functional maps, and pseudo-maps

> Let's attempt to distinguish various map concepts:
>
> **Strict vs. loose maps** (from Rescorla)
>
> "A cognitive map in the **loose sense** is a mental representation that represents geometric aspects of the environment."

“A cognitive map in the **strict sense** is a mental representation that has the same basic representational properties and mechanisms as an ordinary concrete map.”

Spatial map: like a public map, uses spatial relations to represent spatial relations.

Type I functional map: like a spatial map, but uses some other physical relations (e.g. wiring connection, local coactivation) to represent spatial relations.

Type II functional map: like a Type I functional map, but uses symbolic data-structure and symbol-mediated processing relations to represent spatial relations (e.g. by memory address or coordinate system).

Pseudo-map: uses sentences to express the same truth-conditions as spatial map.

Route plan (non-map): uses sentences to express sequence of actions required to follow route.

Camp on (Type I) functional maps

Objection: “But how are we to interpret this claim in the context of thought? If the claim that thinkers employ cognitive maps is read as the claim that there are spatial structures in the brain isomorphic to spatial structures in the world, the objection goes, then this is radically implausible. By contrast, precisely because sentential systems employ such abstract semantic and combinatorial principles, the claimthat a thinker employs a language of thought is compatible with an extremely wide range of plausible neurological implementations. Thus, although by itself Weak-LOT leaves open the possibility of thought with a non-sentential form, one might think, only Strong-LOT offers an empirically plausible implementation ofWeak-LOT”

Reply 1: “In response, however, note first that physically instantiated maps are in fact ubiquitous in the brain. Scientists have known since the 1940’s that the mammalian cortex represents many aspects of the world, especially the layout of one's own body and sensory stimuli, in such a way that the spatial structure of neural firing reflects the physical or psychological structure of the represented content”

Reply 2: “More importantly, the claim that thought might be map-like rather than sentence-like is best interpreted functionally.” “Likewise, the cartographic theorist can take claims about cartographic structure quite seriously, if not fully literally, by maintaining that relations like ‘next to’, ‘above’, and ‘intersecting’, which hold between symbols in a map, correspond to some physical relations—not necessarily spatial—among brain states.”

Conclusion: Type I functional maps are (strictly) map-like.

Rescorla’s argument: algebraic maps are strict maps

Setup. A computational system *M* stores the locations of objects in a geography as coordinate indices of the form <*F*, *x*, *y*>. E.g. <*danger*, 4, 5>, <*self*, 22, 300>. Coordinate indices are stored as an unordered list in memory.

Distance computation. *M* has a distance-computing sub-component *D*. *D* takes as inputs a pair of coordinate indices and outputs a numeral. We can say that *D* computes the following function:

$D(<F, x1, y1>, <G, x2, y2>) = z$ iff`

$\sqrt{}[(den(x1)-den(x2))2 + (den(y1)-den(y))2\ ) = den(z)]$

Route-planning. Route-planning is directly sensitive to the outputs of *D*. For example, to select its next foraging goal, *M* computes the distance to each known food source (according to *D*), and then selects the nearest food source as its next foraging goal.

Spatial structure v. geometric structure. "An essential feature of ordinary cartographic representation is geometric structure… This raises the question: what is it for a mental representation to have geometric structure? An analogous question arises for logically structured representations: what is it for a mental representation to have logical structure? But the question may seem more challenging for geometric than for logical structure. No one thinks that we can open up the head to discover a miniature map laid out in physical space. What, then, could it possibly mean to attribute geometric structure to a cognitive map?" (388)

Metric axioms.

The function computed by *D* accords with the axioms of metric space.

"Any entities, including representations manipulated by a mind or computer, may compose a geometric structure. One need simply ensure that the representations bear appropriate computational relations to one another, relations satisfying the relevant axioms." (389)

"Whether or not we want to say that m is geometrically structured in exactly the same sense as a concrete map, there is a clear sense in which it is geometrically structured. Similar remarks apply to many other represen- tations deployed in AI, Geographic Information Science, and computational geometry." (389)

Rescorla's conclusion: Type II functional maps are (strictly) map-like.

Reply: Type II functional maps are not language-like, but they are still symbolic. They lack key representational features of strict maps: e.g. holism. It is not enough to respect the metric axioms of the world.

Map-like vs. iconic

Cross-cutting distinctions. There seems to be a natural kind *map-like representation* which cross-cuts the natural kind *iconic* representation.

Map-like representations. R is a **map-like representation** iff there is a map M such that (i) R and M have the same truth-conditions C; (ii) R explicitly represents the spatial properties of each landmark in C represented by M.

Map-like vs. iconic. Map-likeness doesn't put a constraint on *how* contentful parts are represented. So a map-like representation could be either iconic or symbolic.

But it is notable that the notion of "map-like" is derivative from the iconic map.

Visual vs. iconic. Likewise, there is a notion of visual representation which cross-cuts the iconic/symbolic distinction. For example, a self-driving car might need to store visual representations.

Alternative view. An alternative view is that map-like representations are iconic at a high-level of abstraction (in which case, there is high-level syntax).

Seven Grades of Mental Map

Setup. An agent that navigates according to an internal representation.

Case I. A spatial grid-map (e.g. a piece of paper) composed of cells.

Case II. A network of cells connected by wires in a lattice, with each wire connection corresponding to a proximate cell in a given direction.

NB: what enforces the metric axioms?

Case III. A network of cells, each containing four addresses (in a numerical code), and each address corresponding to a proximate cell in a given direction.

NB: does holism still hold?

Case IV: A long look up table, with the contents and xy-coordinate of each cell recorded there.

Case V: A pair of look up tables: one containing a list of landmarks, one containing a list of all the distance relations between them.

Case VI: A set of sentences that are controlled by a proprietary algorithm that enforces conformity with geometrical axioms.

Case VII: A set of sentences that expresses the same truth-conditions as a map.