

2

Why Maps are Not Propositional

Elisabeth Camp

1 Introduction

Contemporary theorists of representation often assume a dichotomy between two basic modes of representation: ‘imagistic’, pictorial, or depictive representations, and ‘discursive’, logical, or propositional ones (Fodor 2007).¹ Rational thought and inference are assumed to require implementation with representations that are internally structured and systematically interconnected; but only one sort of structure is assumed to suffice. In one version of the view, found in philosophers like Rey (1995), Bermudez (2003), and Devitt (2006), genuinely logical thought is assumed to require a specifically linguistic medium. A more ecumenical stance allows for the possibility of non-linguistic logical thought, but only because any differences in representational structure are assumed to be ultimately notational (Sober 1976, Szabo 2012).

Given this background picture, establishing that apparently non-linguistic representational formats are logically respectable entails establishing that they are propositional. And establishing this in turn requires specifying what ‘propositional’ (and ‘pictorial’) amounts to. But this can seem like a frustratingly fruitless and abstract terminological matter (Johnson 2015). Given that predication is both relatively well-understood and the canonical mechanism for constructing propositional representations, an alternative strategy is to establish that maps are specifically predicative. This strategy has been pursued by Pratt (1993), Casati and Varzi (1999), and Blumson (2012), among others. Against this, extending work by Sloman (1978), Camp (2007), and Rescorla (2009a, 2009b), I argue that cartographic structure is neither predicative nor propositional. Maps work in an importantly different way from sentences.

The existence of a multiplicity of representational formats has potentially significant implications for the philosophy of mind, which I will not explore here. One

¹ Thanks to audiences at the Conference on Non-Propositional and Imagistic Content at the University of Antwerp, the Southern Society for Philosophy and Psychology, the CUNY Graduate Center, the NYU/Columbia Graduate Conference, the UCLA Department of Philosophy, and the Workshop on Pictures and Propositions at the University of London. Specific thanks to Sara Aronowitz, Russell Epstein, Peter Godfrey-Smith, Gabe Greenberg, Jeff King, John Kulvicki, Daniel Miller, Michael O’Leary, Michael Rescorla, and Dmitri Tymoczko for helpful comments and discussion.

notable implication is that, by attending to the distinct profiles of expressive and computational strength and weakness generated by different representational formats, we can begin to achieve a clearer understanding of key differences among human and non-human cognitive abilities. Instead of an apparent chasm between mere imagistic perception and full-fledged linguistic representation, we can identify families of distinct, sometimes overlapping representational systems that exploit various representational principles in different combinations (Camp 2009, 2015; Camp and Shupe 2017).

My discussion here aims only to establish the need for a formal, non-propositional semantics for an important but restricted class of maps. Actually providing such a semantics is a topic for another day, as is extending the conclusions here, which will focus on external, physical maps, to questions about cognitive architecture. Further, in order to focus on establishing the core claim that maps are not propositional, I will simply help myself to the assumption that maps do have some sort of formal structure, and in particular that many familiar maps are compositional. This should not be highly controversial for very simple maps like seating charts (Camp 2007), or even for more sophisticated systems that employ a fixed base of markers and coordinates—which are the sorts of maps that have been at issue in the debates between advocates and skeptics of predicative and propositional analyses. These cartographic systems permit the construction of indefinitely many maps, such that the representational import of any one map is a rule-governed function of the values of its constituent markers and the locations at which those markers are placed. Thus, such systems are clearly systematic and productive—thereby satisfying the classic motivations for compositionality.

Ultimately, I believe we can construct a formal, compositional semantics for a much larger class of maps, including those that lack a fixed base of markers, like familiar highway maps and atlases. Such map systems employ lines and areas of potentially continuous variability to denote the shapes of objects like roads, forests, and lakes. They often also employ multiple further dimensions of potentially continuous variation: for instance, a map might use variations in hue, saturation, and texture to represent variations in the total quantity, duration, and intensity of rainfall (MacEachren 2004). At the same time, however, they still permit a robust formal analysis, insofar as marks can be typed in wholly formal (ultimately physical) terms and assigned contents via a straightforward interpretation function on those types, and the entire map's content is fully determined by the way those marks are combined. And like their simpler siblings, they can employ markers whose connection to their denotation is arbitrary rather than imagistic (Camp 2007, 2015).

I believe these maps are still fundamentally compositional, despite the absence of a finite lexicon. If we reflect on the reasons for caring about compositionality in a theory of meaning, we should accept a more abstract notion of 'part' on which the representational import of an entire representational unit is an exhaustive, bottom-up function of its parts and mode of composition, so long as these 'parts' are

formally individuated.² I bracket this more controversial claim about compositionality here, however. Readers can either focus just on maps with a finite base and mode of combination, or else grant that a formal semantics is possible for a wider range of maps, whether or not it deserves to be classified as compositional.

Our question, then, is whether maps are propositional representations. In section 2, I argue that if the question is to be a substantive one, we should interpret propositionality in terms of functional structure; and that propositional structure is characterized by being digital, universal, asymmetrical, and recursive. In section 3, I present the two most developed semantics for maps, by Pratt (1993) and Casati and Varzi (1999), both of which argue that maps have a specifically predicative structure. In section 4, I examine the elements and constructions that are posited by a specifically predicative analysis and by a propositional analysis more generally: names and predicates, and complex propositions and truth-conditions. In each case, I argue, the application to maps is awkward at best. Standard maps exhibit expressive limitations and advantages that are left unexplained on a propositional analysis. By contrast, a non-hierarchical, holistic structure better explains the distinctive contours of cartographic representation.

2 Propositionality and Non-Propositionality

If the question about whether maps are propositional representations is to be a fruitful one, the term ‘propositional’ needs to be interpreted in a way that integrates with other standard uses of the term, that is specific and robust enough to make substantive claims, and that marks a useful distinction among representational formats. In particular, it should not follow trivially, just from the operative notion of propositionality, either that all representations are propositional, or that only sentences are. But ‘propositional’ is used in many different, sometimes cross-cutting ways in philosophy and related fields. In this section, I survey the leading philosophical understandings of propositionality, and articulate a notion of propositionality that is grounded in, but more general than, the functional structure of language. My aim in this section is just to establish a useful, non-question-begging framework for conducting the debate about maps’ propositionality—a framework I think is implicitly shared by all current parties to the debate.

2.1 *Propositions as Sets of Possible Worlds*

Many philosophers hold that a representational state or entity has propositional content just in case it can be the object of belief, can be shared by multiple cognitive

² Cf. Kulvicki (2015a). Even with this expansive notion of compositionality, the claim that natural languages are compositional is best seen as a regulative ideal (Szabo 2012); the operative notion of ‘language’ in most discussions of cognitive architecture is that of formal predicate calculus rather than e.g. English (Camp 2015).

agents, and can be evaluated for truth and falsity. A particularly elegant way to capture these requirements treats propositions as partitions in a space of possibilities (Lewis 1968, Stalnaker 1976). This sense of propositionality clearly does not trivialize the question whether a given intentional state or entity is propositional: for instance, merely considering mental states, it counts both objectual and expressive attitudes as non-propositional.³ Moreover, it links propositional content closely to information, in a way that has clear theoretical utility, as evidenced by the ubiquity of possible-worlds accounts in, for instance, philosophical and linguistic accounts of conditionals, epistemic modals, and many other constructions (e.g. Stalnaker 2014).

Although the possible-worlds model has a lot going for it, more generally interpreting the question whether maps are propositional in such coarse-grained terms runs a significant risk of trivializing questions about representational format, by classifying all truth-assessable representations as propositional. Like pictures, maps have informational contents: they represent situations, and make cuts in possibility space. As a result, they are plausibly assessed as true or false—although that assessment is typically relativized to certain purposes, and often also to just a subset of a given map's overall representational content.⁴ Further, again like pictures, the informational content of any given map can plausibly be captured in an (extremely long) sentence, at least in the weak sense that the map or picture itself can be recovered from that sentence by pixelating the picture to an arbitrarily fine degree of resolution and specifying the values of colors (perhaps using demonstratives) for each pixel (Anderson 1978, 253; Crane 2009, 460).

It is natural to respond that such a weak sense of informational equivalence ignores crucial implicit assumptions about how users extract information about the world from those representations. This response is correct, but of limited use in establishing the utility of the coarse-grained analysis of propositions in this dialectical context. Within a possible-worlds framework, many of the cognitive, semantic, and pragmatic phenomena for which propositions are traditionally invoked must instead be explained by appeal to the different *ways* or forms in which a common informational state can be presented. And since maps, sentences, and pictures do take different forms, and produce distinct cognitive, semantic, and pragmatic profiles, the questions that other theorists attempt to explore by asking whether maps are propositional will simply end up being recapitulated within the possible-worlds framework in terms of those forms and functions.

³ Thanks to Alex Grzankowski for pressing this point.

⁴ Some theorists, such as Crane (2009), argue that pictures are non-propositional on the ground that truth is an absolute property and accuracy a matter of degree, and that pictures can only be accurate, not true. I do think the difference between truth as an on-off matter and accuracy as a graded, respect-relative one points to a correlative difference between propositions as discrete, digital representations, and more holistic, analog modes of representation. However, I also think that we do often assess maps for truth as well as accuracy, and that such assessment is more natural for maps than for pictures. Further, intuitive assessments of truth are themselves often interest- and context-relative, and it is often at least as natural to assess complex collections of sentences, like books, for accuracy as for truth. Thanks to Peter Godfrey-Smith for discussion.

In order for the question of representational format to be a substantive one, we need to interpret it as concerning *how* maps and other representational systems work, where this involves more than identifying either what information they represent or the physical marks that constitute any one representational vehicle. Instead, we need to determine how such marks function within a larger system, including how users access and manipulate that representation, what other representations the system generates, and how altering any one aspect of any given representation affects its overall representational content. The identification and assessment of representational form is thus inevitably directed at a complex package of actual and potential representations and interpretive processes (Anderson 1978, 263). At the same time, given that the same functional pattern of representations and operations can be instantiated in physically distinct ways, the identification of these representations and processes is also not appropriately directed at the representational system's ultimate underlying implementational mechanism.

In his classic (1978) discussion of the mental imagery debate, John Anderson proposes a minimal addition to the purely informational, possible-worlds model of propositions which shifts the analysis to the level of a representational system's functional profile. He takes propositions to possess three features (1978, 250). First, propositions have truth-values. Second, propositions are abstract, in the sense of not being essentially tied to a particular mode of expression; more specifically, Anderson interprets this condition in terms of invariance across at least some range of paraphrases (that is,) substitution of semantically equivalent marks. Third, propositions have explicit formation rules, which determine well-formedness and impose a "structural aspect" that underwrites psychological laws and/or logical rules of inference.

By understanding propositions as truth-evaluable, abstract, and structured, Anderson shifts the analysis from a purely semantic one, concerning only what is represented, to a partially syntactic one, about form—what principles govern the system's formation, compilation, and manipulation of representational states or entities? At the same time, by understanding this form abstractly, in terms of permutations over multiple classes of distinct expression-types, he frees propositions from any essential connection to particular vehicles.

I think this locates the question at the right level. Anderson's own aim in proposing this model is merely to distinguish propositional from imagistic representations. And for this purpose, the combination of abstractness and structure is arguably sufficient, since it is far from obvious that images do have formation rules that determine well-formedness or underwrite systematic manipulation. However, Anderson's criterion simply assumes that all structured representations have fundamentally the same structure—an assumption that, as we saw at the outset, is ubiquitous in discussions of representational type. In the current context, though, it answers the question whether maps are propositional by stipulation. Maps, graphs, diagrams, and sentences all meet all of Anderson's criteria. In particular, they are all abstract in his sense of not being essentially tied to any particular implementation: many different types of marks

can represent trees on a map, or magnitudes in a graph. Further, any given cartographic or diagrammatic system might employ multiple markers for what is in fact the same entity: a map of the night sky might have functionally distinct markers for the Morning and Evening Stars. At the same time, these representational systems also differ from each other and from sentences in their modes of operation, patterns of breakdown, and expressive power (Camp 2007). What we want to ask is whether these differences are mere *ad hoc* contingencies or reflect a more fundamental difference. And to even ask this question, we need to distinguish, at least in principle, among types of structured representations. In particular, we need to identify a type of combinatorial structure that is distinctive to propositions, without being stipulatively identified with sentences, so that we can ask whether other structured representations exhibit that same structure. Thus, I now turn from coarse-grained, possible-worlds accounts of propositions to those that posit internal structure.

2.2 Structured Propositions

The two classic models of structured propositions, Fregean and Russellian, differ in their constituents and criteria of individuation. The choice between Russellian and Fregean propositions turns on theoretical matters outside our concern, which is just about propositional *structure*; and in this respect the two are quite similar. Both count as abstract by Anderson's criterion, in allowing multiple distinct representations—most obviously, translations into different languages—to express the same proposition.

On Frege's model, functional application is a general combinatorial mechanism. In the basic case, an object serves as input to a special kind of function, which Frege calls a 'concept', that in turn delivers Truth or Falsity as its output; a Fregean proposition, or Thought, is a sense (mode of presentation) of Truth or Falsity which combines the senses of the represented object and concept. The different functional roles played by objects and concepts—the former serving as input to the latter—establishes a corresponding asymmetry between types of senses, and in turn between the signs that express them. As Frege (1892, 54) famously puts it, "For not all parts of a thought can be complete; at least one must be 'unsaturated', or predicative; otherwise they would not hold together." More complex propositions are formed by recursive application of the same functional machinery.

Where Fregean propositions are abstract thoughts, individuated by cognitive significance, Russellian propositions are ordered n -tuples of objects and properties. Thus, where Frege introduces the vague metaphor of 'saturation' as a gesture at how the senses of concepts form whole thoughts when 'completed' by the senses of objects, Russellians must confront the question of what unifies an ordered sequence into a truth-evaluable whole more directly. The most common answer appeals to the metaphysical structure of instantiation: atomic propositions are complexes of object(s) instantiating properties, and complex propositions are property-object complexes instantiating higher-order relational properties. An alternative answer,

picking up on an equivalence between predication and saturation suggested by Frege above, appeals directly to predication (King 2007).

Given that both Frege and Russell were interested primarily in constructing the predicate calculus as an ideal formal language for scientific inquiry and secondarily in analyzing natural language, it is no accident that both analyze propositional structure in a way that establishes a close connection to predication. However, for our purposes, in the absence of an independent grip on what predicative structure itself is, we risk distorting the question of whether maps or other formats are propositional by fixating on features of logical or linguistic structure that arise from contingent implementational factors that are specific to language. Rescorla helpfully characterizes predication as “a compositional mechanism whereby [denoting] terms fill the argument-places of a predicate that carries their denotations into a truth-value” (2009a, 177). Like Rescorla, I will argue that maps do not exhibit predicative structure, albeit on somewhat different grounds. However, I am primarily concerned with this question as an instance of the more general issue of propositional structure, as a way of achieving a more systematic understanding of significant dimensions of variation among representational structures. So for our purposes, we still need a characterization of the propositional genus of which predication is a canonical species.

Moreover, there are lively debates within linguistics about the combinatorial structure of language, and the status of predication in particular. Some take predication to be a special case of a fully universal principle of functional application; while many think that additional mechanisms, such as predicate modification, are also required (Heim and Kratzer 1998). And many others posit a different, more general mechanism, such as Merge—an operation of set formation targeting two objects (Chomsky 1995, Collins 2011). Finally, we should not foreclose by stipulation the possibility that maps lack predicative structure at the atomic level, but that entire maps can still be analyzed using a propositional logic that combines non-predicative atomic units into larger wholes.

We thus have several reasons to inquire into propositional structure without tying it essentially to either language or predication. Stepping back, we can identify four key features shared by all five of the candidates for propositional structure mentioned so far (that is, by conceptual saturation, property instantiation, predication, functional application, and Merge). First, the candidate combinatorial operations are highly *digital*: they take a small number (typically, a singleton or pair) of discrete elements as inputs. Second, they are *universal*, or at least highly general: they can combine a wide range of elements. Thus, instantiation is a highly general metaphysical relation binding many different kinds of properties and objects into situations or facts; likewise, functional application is a highly general operational relation taking many kinds of objects and functions and delivering truth-values. Third, the candidate combinatorial operations are *asymmetrical*: either just one element must be of a type

that enables it to serve as input for the other; or the operation itself creates an asymmetry among elements via the order of application.⁵ Fourth, the operation is *recursive*, so that its outputs can serve as inputs to the same operation: thus, quantifiers are functions that take functions as inputs and deliver truth-values as outputs. The recursive application of an asymmetrical relation produces *hierarchical* structure, containing nested iterations of the same type of representation. No one of these features is distinctive to language; in particular, other representational formats, like genealogical trees, employ combinatorial operations that exhibit some of these features to at least some significant degree (Camp 2009, 2015). And some will deny that even language exhibits all of these features. But together, they offer a cluster of features that underwrite a substantive notion of ‘propositional’ that is both familiar and relatively well-specified.



I will argue that maps exhibit a very different sort of structure: that they are non-hierarchical, holistic modes of representation. In section 3, I briefly present the leading formal analyses of maps, which treat them as having specifically predicative structure. In section 4, I survey various features we would expect maps to display if they did employ predicative or more generally propositional structure, and find the application lacking.

3 Predicative Map Semantics

In this section, I present two leading proposals for a cartographic semantics, both claiming that maps have propositional, specifically predicative structure. On the account offered by Roberto Casati and Achille Varzi (1999), a map is a collection of colored regions. First, map-regions are defined by their mereotopological relations, and are assigned to regions in the world by a mereotopologically well-behaved interpretation function f , so that map-regions are parts of or contiguous with one another if and only if the world-regions to which they are assigned are related in the same way (Casati and Varzi 1999, 194). Second, a color patch covering such a region is “an unsaturated predicate, which gets saturated when it is juxtaposed to a map region” (1999, 192). Third, an “atomic map-stage” is a mereological fusion of all regions of a single color, F , on an entire map. Atomic map-stages are assigned a two-part truth-condition: all of the world-regions denoted by F -colored map-regions must have the F -associated property, and the complement of those regions must lack that property. Casati and Varzi hold that the semantics is “compositional” and “recursive” insofar as entire maps are constructed by layering atomic map-stages, and the truth-condition for the entire map “results from the conjoined truth” of all of its

⁵ For example, the set formed by Merging the singleton set of Bob with the set of love and Mary is {Bob {love, Mary}}, which differs from the set formed by Merging the singleton set of Mary with the set of love and Bob: {Mary {love, Bob}}.

map-stages (1999, 195). The upshot is that maps “are propositionally structured albeit in a peculiar way” (1999, 191).

In the interest of tractability, Casati and Varzi restrict their semantics to a narrow class of formalized cases. They do not consider maps containing arbitrary labels; nor do they address richer geometrical relations, such as distance, direction, and orientation. By contrast, Pratt’s (1993) account hews more closely to a linguistic model, by taking as its basic case the locating of word-like symbols. Following Schlichtmann (1985), Pratt treats an individual “map-symbol” as having two aspects: a “substantive component” X (e.g. ‘’) and a “locational component” (x, y) . Each “component” has a distinct interpretation function: the symbolic function I assigns symbols to a domain of individuals in D (e.g. assigning ‘’ to the post office); while the spatial function μ assigns locational components to places in a space S . An individual map-symbol is true just in case $I(X)$ is located at $\mu(x, y)$ (Pratt 1993, 80). Symbols covering extended regions are treated as true (roughly)⁶ just in case the element assigned to that symbol-type is present at all locations within the relevant region (1993, 86). Finally, an entire map is true if and only if, first, all of its individual map-symbols are true, and second, the overall extensions of the substantive components are *minimal*, so that there are no elements in D of the type assigned to X by I that are not located on the map (1993, 82).

The semantic analyses proposed by Casati and Varzi and Pratt are very different, in ways we will consider shortly. But they both arguably get the truth-conditions largely right—at least, for the classes of maps they consider. They thereby demonstrate that many maps are indeed formally structured in such a way that it is possible to assign truth-conditions to entire maps by identifying stable, formally individuated parts, and assigning a stable semantic interpretation to those parts and to their mode of combination. By contrast, it is much less obvious that an analogously formal semantics is constructable for pictures. Further, if it worked, a propositional model would have important explanatory advantages. Although there will obviously be some implementational differences, a common propositional code would streamline informational integration across representations couched in apparently different formats (Pylyshyn 2003). It might also enable us to draw on more sophisticated aspects of predicate logics in extending the analysis to address more complex maps.

Despite these potential advantages, I will argue that neither a predicative nor a more generally propositional analysis captures the way maps work. There is little positive evidence that maps employ a structure that is digital, universal, asymmetrical, recursive, and hierarchical. On the one hand, the expressive power of distinctively propositional structure is representationally inert if it is present at all; and on the other, the constraints associated with propositional structure should impose restrictions on how maps are constructed and manipulated that don’t appear to be borne out.

⁶ Pratt argues that issues involving figure-ground ambiguity undermine a straightforward interpretation along these lines; but his various disambiguated analyses all share the common structure stated in the text, and the worries about ambiguity are irrelevant for current dialectical purposes. I return to this issue below, in fn. 9.

4 Problems with Propositionality

4.1 Names

Casati and Varzi explicitly state that “map regions are to be considered symbols in themselves, akin to individual constants such as linguistic proper names” (1999, 191). While Pratt does not say that map-regions are names, he too analyzes map-locations as object-level inputs to predicates; and like Casati and Varzi, he treats map coordinates as mere ‘tags’, where the assignment to map-locations of world-location-values is determined directly by the interpretation function, without any mediation by condition-satisfaction. So the first question we need to ask is whether map-locations do behave like names.

The most literal, straightforward understanding of the predicative model would seem to entail that it is essential to a map’s well-formedness that it be assigned worldly coordinates. That is, if the simplest map *markers* (arbitrary symbols on Pratt’s analysis; colored regions on Casati and Varzi’s) are predicates, and if a map’s coordinates or regions are names to which those predicates are applied, then a map which lacks an assignment from map-locations to world-locations should just be a set of predicates denoting a set of properties, rather than a well-formed, ‘saturated’, truth-condition-determining whole.

At least two factors make it less than straightforward to assess this prediction. First, where linguistic names are distinct identifiable marks (e.g. ‘Bill’), maps typically lack marks that function just to denote locations; the closest common analogue might be grid quadrants (e.g. ‘C4’), but these are often marked only at the map’s periphery. So it’s not obvious, in purely syntactic terms, what a collection of mere unsaturated map predicates would look like, and how it would differ in form a map with assigned world-locations. Second, the interaction between syntactic structure, semantic content, and pragmatic supplementation is considerably less systematic in general in maps than in language—where it is already far from straightforward. In particular, just as utterances of isolated linguistic predicates (e.g. ‘gone’) can be elliptical for complete sentences and used to make complete assertions (e.g. ‘Gone!’) and other speech acts (e.g. ‘Gone?’), the proponent of a predicative analysis can hold that unarticulated coordinates are implicitly present in maps.

Even with these caveats lodged, though, we still have reason to conclude that having assigned world-locations is not essential to a map’s being either syntactically well-formed or semantically interpreted, because maps without assigned locations can still fix substantive satisfaction-conditions. For instance, I might construct a map of the estate I hope to construct when I win the lottery, deciding where to place the pond relative to the flower and vegetable gardens, the circular driveway, and the grand house’s corner tower. Not knowing the extent of my future riches, I might refrain from determining a scale for the whole estate; nor have I determined where it will be built. Yet my map still rules in some ways the world could be and rules

out others. When I do win, I could present the map as a command to my realtor/architect team: make the world be *this* way, leaving the site selection and scale up to them. Likewise, if I were CEO of a drugstore chain obsessed with maximizing efficiency and the uniformity of patron experience, I might construct a map indicating where each type of item (toothbrushes, cosmetics, soda) is to be placed in every store, and present it as an order to my workers: arrange your stores *that* way. Once they have carried out my instructions, I might distribute copies for patrons to use in navigating any individual store.

The coherence of these cases undermines the *prima facie* prediction of the predicative model that maps with uninterpreted coordinates should be at least semantically, and possibly syntactically incomplete. It also suggests that map coordinates do not always function as names, because they are not necessarily rigid: my architectural blueprint could be, and my drugstore map eventually is, satisfied by many different locations. Finally, these cases point to another way in which our actual use of maps fits poorly with assigning an essential role to map-locations as names. Often, our only practical interest in, and epistemic access to, the world-locations that are denoted by map-locations goes through the objects and properties located at them. (Indeed, this is reflected in Casati and Varzi's semantics, insofar as map-regions are constituted *as* syntactic elements, with an object-level denoting function, *by* being colored.) More generally, many maps—not just informal ones drawn on napkins, but published tourist maps—provide no explicit coordinate assignment, relying instead on notable landmarks to implicitly fix the interpretation function from map-locations to worldly ones.

At a more fundamental level, it is notable that the metaphysical distinction between objects and properties, which is closely mirrored in language by the syntactic distinction between 'saturated' subjects and 'unsaturated' predicates, is not nearly as important in maps. Insofar as the distinction between individuals and properties shows up, it seems to be an ad hoc constraint (a 'meaning postulate') on whether the marker can be multiply tokened on a single map. This might seem like an expressive limitation of maps: that it fails to mirror a deep metaphysical feature of the world (Fodor and Pylyshyn 1988). But it might also, instead, be construed as an expressive advantage, given the long tradition of skepticism about the notion of subsisting essences that the object/property distinction seems to entail. (Some theorists in this skeptical tradition maintain that individuals and species just are homeostatic clusters of properties (Boyd 1999); while others hold that no real reference to individuals is possible without quantificational structure (Quine 1995, Davidson 1999)—or at least, that there is a level of cognitive representation at which no genuine distinction between individuals and kinds can yet be drawn (Millikan 1998).)

All of this suggests that rather than map-locations functioning as names, their closest propositional analogues might instead be quantifiers. On this analysis, a particular map coordinate or region on a map that employs a merely implicit, landmark-relative interpretation function would have roughly the significance of 'the place/region in

the world, whatever it might be, where the denotation(s) of this marker is located'; the map as a whole would be true just in case all of the object/properties denoted by the map's markers were related to each other in a spatial configuration that was appropriately isomorphic to the spatial relations among the markers.⁷ My architectural blueprint would then impose the requirement that there exist some set of appropriately spatially related world-locations which instantiate the object/properties denoted by the markers at each of the corresponding map-locations. And my store locator maps would be true of a particular store if it arranged its stock in the specified pattern; while my master map would be true only if and only if all of my stores were laid out that way.

Thus, in the face of the challenges canvassed earlier in the section for treating map locations as names, the predicative theorist can hold on to the claim that map markers (whether colored regions or symbols) do function as predicates by invoking an analysis of map locations as quantifiers rather than names. In section 4.3, I canvass the possibility of more direct evidence for quantificational structure in maps, and find it largely lacking. First, though, in section 4.2, I examine the core assumption of the predicative view: that map markers predicate properties of locations.

4.2 *Predicates*

In this section, I argue that there are at least two deep disanalogies between predicates as they function in language and the way they should function in maps if a predicative analysis were apt.

4.2.1 EXPRESSIVE LIMITATIONS ON PREDICATION

The first difference follows directly from the account developed in section 2.2 of propositional structure. A distinctive feature of predication, as of propositional operators more generally, is that a wide range of predicates can be predicated of a wide range of objects. In its strongest form, any expression of predicative type can be combined with any expression of subject type. Many theorists have wanted to impose categorical restrictions on such combination, on the ground that sentences like 'Julius Caesar was a prime number' are too absurd to express genuine thoughts. I reject such semantic restrictions (Camp 2004). But actual natural languages impose many syntactic restrictions within expression types, some but not all of which appear to have at least some semantic basis (Johnson 2004). Even if we include both semantic and syntactic restrictions on predication, however, it remains true that language enables us to say of a very wide range of things that they are a very wide range of

⁷ In some cases, it might be essential to the denotation that it be located in a specific location: perhaps the Grand Canyon could not be located elsewhere and still be the Grand Canyon, while London Bridge could be relocated to Alaska. Cases of the former type would then be essentially anchored to a particular location, but in virtue of the semantic interpretation assigned to the marker, rather than in virtue of the interpretation function assigned to the map-location at which that marker is placed.

ways. And a key reason for this generality is that predication itself contributes very little to the overall semantic import of the sentence, beyond simply *applying* the denoted property to the denoted object. The point holds with even greater force for other candidate propositional operations: saturation, functional application, and Merge (Camp 2015).

By contrast, the ‘propositions’ that result from placing map-markers at locations are much more expressively restricted: they always say ‘*X is here*’. This restriction does not arise from any inherent constraint on the meanings of the markers themselves, which can be just as arbitrary as in language (indeed, map-markers often are words). Rather, they arise because the *placement* of a marker at a map-location has a fixed significance: that the denoted object/property is located at the denoted world-location (where, as we saw in section 4.1, those world-locations may be specified in relational rather than absolute terms). The fact that it is impossible to ‘predicate’ object/properties of anything *but* locations in maps thus strongly suggests that the combinatorial operation itself is something other than predication.

To see the contrast more clearly, consider a branching tree structure, interpreted as instantiating either a genealogical or a propositional structure (Camp 2009). In both cases, the combinatorial principle is digital, recursive, and asymmetrical, and thus hierarchical. But in a genealogical tree, a common node means something very specific: common ancestry. As a result, such trees can only represent entities *as* being related by ancestry and descent. In principle, we could interpret genealogical trees as having propositional structure, and claim that their users are only ever interested in representing a single topic, or only ever interested in one kind of hierarchical structure (Cheney and Seyfarth 2007); but the imposition of this expressive limitation would be *ad hoc*. It makes more sense to treat the combinatorial principle itself as semantically loaded, in a way that constrains the sorts of inputs for which it can deliver true representations. Similarly for maps, we could posit an *ad hoc* constraint blocking maps from exploiting the full generality of propositional structure; or we could posit a distinct combinatorial principle—of using spatial structure to represent spatial structure—from which this expressive limitation follows directly.

4.2.2 PREDICATES AND ABSENCE

The second point also concerns a respect in which the predicative view is committed to treating predication as having an unexplained representational import in maps as compared to language: the ‘Absence Intuition’. In section 4.1, I considered collections of map-markers that lack assigned world-locations, and concluded that they still comprise a well-formed map, which can sometimes also be assessed for truth, depending on the pragmatic ‘force’ with which it is presented. What about the converse case, of map-locations without markers? If map-markers are predicates, then *prima facie* one would similarly expect unmarked map-locations to lack

truth-conditions: they should be ‘bare’ location-denoters. Likewise, one would expect deleting a single marker from a map to produce a syntactically incomplete expression—or at least, to delete just one proposition. However, both Pratt and Casati and Varzi construct their semantics quite differently, with the aim of respecting the widespread intuition that the *absence* of a marker-type at a map-location represents the absence of the denoted object/property from the denoted location.

As noted in section 3, Casati and Varzi implement this intuition with a ‘complementation’ clause on the truth-conditions of atomic maps: a map stage for color F is true just in case all of the F -colored regions on the map possess the associated property, *and* the complement of those regions does not possess that property (1999, 194). The second clause entails both that maps never do contain bare, ‘unsaturated’ locations, and also that deleting a marker introduces the absence of the denoted property rather than simply erasing the associated positive proposition.

We might already worry that the first clause makes the atomic propositions in maps inappropriately ‘large’: insofar as we can in fact intuitively assess individual sub-regions within an atomic map-stage for truth, an entire atomic map stage does not constitute the most basic unit of cartographic information and assessment. Their second clause amplifies this worry at a more fundamental level, because it makes the basic unit of assessment dependent, not just on how things are at all of the world-regions where the denoted object/property is represented as being located, but also on how they are at world-regions where the denoted object/property is *not* represented as being located. But, as Michael Rescorla (2009a) argues, predication just *is* a function from the denotation of the predicate’s input to a truth-value. Moreover, as I argued in section 2.2, propositional structure more generally employs a recursive, asymmetrical operation which delivers a unique output given the values of its inputs, where the number of inputs is very small, typically a singleton. Casati and Varzi’s semantics either makes the value of the whole proposition depend on something other than the values of its parts, or else individuates atomic propositions in a way that doesn’t fit with the way that they themselves define the individual parts.

Pratt’s treatment of Absence avoids both of these problems. His basic unit of assessment is a marker at a single point. A region is assessed as true iff only all points within that region instantiate the object/property denoted by the symbol. More importantly, he treats Absence as arising from a general, format-neutral feature of default reasoning: an inference from the failure to explicitly say that something is abnormal to the conclusion that it must be normal. This general inference can be captured within a predicate logic by stipulating that an entire set of formulae counts as true in a model just in case the individual formulae are true and the predicate *abnormal*(x) has no unmentioned members. Analogously, Pratt holds that an entire map is true iff each of its individual symbol-pairs is true, and the overall extensions of the symbols’ substantive components are all minimal (or *circumspect*), so that there are no unmentioned elements of those extensions within the entire represented domain (1993, 82). Thus, unlike Casati and Varzi’s analysis, this model does assign

truth-conditions that depend only on the values of a given symbol-location pair; Absence is enforced as a condition on the truth of an entire map.

Although Pratt's analysis avoids the letter of Rescorla's criticism and rehabilitates the intuition that particular placed markers can be assessed for truth, it still introduces a fundamentally holistic element to the determination of truth-conditions. Rescorla claims that this holistic element has no place in a predicative analysis, and that the Absence Intuition "reflects [a] fundamental representational disparit[y] between attaching a marker to a coordinate and attaching a predicate to a singular term" (2009a, 197). I am not convinced that the difference between maps and languages is quite so stark in this particular respect. Although the Absence Intuition is very strong when it comes to published maps designed for general navigation, many other maps are more informal, more restricted in their representational purposes, and/or weaker in their epistemic authority; and for them the intuition is much weaker (Camp 2007, Blumson 2012). For instance, if I draw a map showing you the route from the hotel to the party, then my placing a gas-station-denoting marker next to a McDonald's-denoting-marker just before a crucial left turn would not normally be taken to represent the absence of other gas stations along the route. The assumption in such cases is only that the map will contain markers for all known relevant objects and properties, where what counts as relevant may be quite temporally and functionally restricted.⁸ Further, as Pratt makes clear, parallel forms of default reasoning involving circumspection are pervasive in linguistic communication, perhaps most notably in the calculation of scalar implicatures. Linguists generally take such inferences to be merely pragmatic, although some such inferences arguably produce genuinely semantic effects (Chierchia 2004). Thus, we should at least entertain the possibility that a parallel pragmatic analysis is appropriate for at least many maps.

Like Casati and Varzi and Rescorla, John Kulvicki (2015b) takes the Absence Intuition to be sufficiently pervasive and robust to warrant semantic treatment. Further, like Casati and Varzi and Pratt, he claims that maps predicate properties of locations. However, he argues that they "organize" those predicates in a fundamentally different way than sentences. This is partly because they exploit spatial structure (as we'll discuss in section 4.3); but because they introduce families of incompatible predicates "en banc" (2015, 158). So, for instance, if blue denotes water and green denotes land, then the incompatibility between being simultaneously blue and green mirrors a corresponding incompatibility between being simultaneously sea and land. (Incompatibilities among markers may be defined purely conventionally.) In the simplest case, the incompatibility is just between the presence or absence of a single isolated marker; but many "incompatibility classes" contain indefinitely many marker-types, of arbitrarily fine-grained difference. The number of contrasting marker-types in the class defines a "degree of freedom" for how things could

⁸ Similarly, Kulvicki (2015, 152) cites official governmental navigational maps, which mark only those structures visible from the sea that are relevant for navigation.

be at that location: for instance, an incompatibility class containing blue, green, and beige, representing land, water, and beach respectively, has three degrees of freedom. Markers from distinct incompatibility classes are mutually compatible: for instance, if the same map uses distinct textures to represent variation in elevation—say, smooth, rugged, and mountainous—then a single map-location could be marked as both green and notched.

In effect, Kulvicki saves the claim that maps predicate features of specific locations from the threat of holism posed by the Absence Intuition by proposing a systematic meta-semantic constraint on relations among predicates. The presence of a marker at a location simply predicates the denoted property of that location. But the absence of any marker from within a given incompatibility class counts as a predicate in its own right—it denotes the “zero value” for that class, the property of having *none* of the features in the class.

The proposal that maps employ families of markers linked by relations of incompatibility is highly attractive, for a number of reasons. It explains Absence by citing a deep and ubiquitous feature of maps. It also naturally generalizes from absolute incompatibilities—for example, that green-land entails *not* blue-water—to explain the fact that maps frequently use families of marker-types to express positive relational information: for instance, the use of greater color saturation to represent higher density of rainfall.⁹ Nevertheless, because Kulvicki’s analysis retains the imputation of a fundamentally predicative structure, it does nothing to explain *why* maps exhibit this systematic interpretive constraint, and in particular why the contrast between maps and language is as robust and pervasive as Kulvicki, following Casati and Varzi and Rescorla, takes it to be. In this respect, Pratt’s pragmatic analysis of Absence in terms of format-neutral principles of default reasoning offers a deeper explanation, one that fits more smoothly with the predicative model.

Kulvicki’s discussion of how map systems “organize” their constituent expressions into families points to a distinct dimension, beyond spatial organization, on which maps are holistic, in contrast to the highly digital structure that is characteristic of propositional representations. As I argued in section 2.2, predication and other candidate propositional combinatorial principles generate individual propositions by recursive application of an operation which takes a small input and delivers a small (typically single) output. Any non-pragmatic treatment of Absence fits uncomfortably with the propositional model, simply because it introduces a significant degree of holism into the analysis.¹⁰ I will argue in section 4.4 that even putting

⁹ Further, it plausibly helps to address the worries that Pratt (1993, 84 ff.) raises about figure-ground ambiguities in connection with Absence, which Pratt takes to ultimately undermine the possibility of a fully formal semantics that could explain map use.

¹⁰ Again, the contrast with actual natural languages (as opposed to formal models of predicate logic) should not be overstated. In particular, it is a deep, productive insight in linguistics and the philosophy of language that sentences occur not in isolation but within extended discourses, with systematic semantic effects (e.g. Roberts 1996/2012, Kehler 2002, and Asher and Lascarides 2003).

Absence aside, maps' holistic organization is not well modeled by a propositional analysis. Thus, although I believe we should adopt something like Kulvicki's proposal about families of marker-types, I think we should do so without grafting it onto a predicative or propositional foundation.

4.3 *Complex Propositions*

So far, I've been assessing the propositional model at the basic level: of predicates that take object-level inputs and form atomic propositions. When we turn to the ways in which maps compile information about particular locations into larger wholes, the disparities with propositionally structured formats become more stark. On the one hand, propositions are constructed by a digital, asymmetrical, recursive mechanism which creates hierarchical structure of indefinitely nested depth; and this structure is not obviously manifested in maps. On the other hand, maps encode information in a holistic, relational way; and this structure is not directly captured by a propositional mechanism.

I mentioned in section 4.1 that the predicative theorist could plausibly accommodate cases of 'unanchored' maps that do not represent specific locations, such as blueprints, by appealing to quantification. However, such cases aside, there is notably little role for quantification in maps. In particular, because maps work by placing markers for object/properties at locations, maps cannot represent purely quantificational information. For instance, on a map, one can only represent that someone with a red shirt is carrying a gun by representing a particular individual or type of person as being at some more or less specific location (where again, that location may be relational) (Camp 2007, 165). Likewise, one cannot explicitly represent that everyone wearing a red shirt is carrying a gun: at best, one can co-locate all the markers for red-shirt-wearers and gun-carriers, and rely on Absence to communicate that no one else is; but this still requires co-locating the red-shirt-wearers and gun-carriers *somewhere*, and thus includes location-specific rather than purely quantificational information.

This limitation suggests that maps are primarily first-order modes of representation: they just place objects and properties at locations. One might be tempted to hold that this is all maps can do. In support of this, one might point to the implausibility of modifying a predicative map-marker which denotes gun-carrying by an adjectival map-marker which denotes red-shirt-wearing, or of applying a predicative marker which denotes red-shirt-wearing to an individual-marker which denotes Steve. Instead, it seems that all one can do is to co-locate the markers for Steve, red-shirt-wearing, and gun-carrying. This suggests that map markers all have the same basic syntactic status, and that the semantic import of co-locating markers is always just the conjunction of their denoted object/properties.

Although I think something like this is on the right track, I don't think the restriction to a flat, first-order mode of representation can be quite so absolute. For one thing, there are arguably cases of genuinely nested maps—say, a computer map

of a city, where icons for individual buildings expand to become maps of the buildings' interiors. More importantly, maps often deploy multiple "incompatibility classes," some of which have second- or higher-order significance (MacEachren 2004). For instance, a cartographer might deploy hue to denote the most common disease in the denoted region; saturation to represent the frequency of that disease there; and textural density to represent the reliability or evidentiary status of the representation of disease as having that frequency in that area. The information thereby conveyed is hierarchically structured: not just that disease x plus frequency y plus reliability z are all present at this location, but that there is reliability z of disease x occurring with frequency y here.

At the same time, I don't think the presence of such layered information warrants an ascription of hierarchical syntactic structure to maps. When maps do support a layered semantics, I claim, this is because the markers themselves have been assigned values that mandate this interpretation, not because their syntax is hierarchical. Thus, an interpretation function that assigns evidentiary status to a certain family of textures entails, merely in virtue of what evidentiary status *is*, that the presence of that texture on the map can only coherently be interpreted as modifying the significance of some other marker with more basic significance, rather than the other way around. By contrast, in propositional systems, the fact that the combinatorial operation itself is asymmetrical entails that same set of expressions—say, 'Bob', 'loves', and 'Mary'—can be combined in a different order to produce semantically distinct results. It may be for lack of imagination, but I have been unable to construct analogous syntactically generated truth-conditional differences in maps. Nor do maps appear to permit the construction of indefinitely complex hierarchical structures by repeated iteration of the same operation, as in "Everyone who is carrying a gun is standing next to someone who is wearing a red shirt, and owes money to someone who was wearing a blue shirt yesterday." Indefinitely deep hierarchical recursion is highly characteristic—some would say definitive (Hauser, Chomsky, and Fitch 2002)—of propositional structure. By contrast, cartographic structure appears to be largely, if not absolutely, flat.

4.4 *Conjunction and Holism*

The semantics proposed by both Pratt and Casati and Varzi comport with this basically flat structure, insofar as they analyze entire maps as mere conjunctions of atomic representations, and conjunction itself imposes no hierarchical relationship between the conjuncts. An exclusive focus on conjunction as the mechanism for compiling entire maps from atomic parts also makes sense, given that it is not obvious that maps usually represent other sorts of connections between, or modifications of, basic states of affairs, such as disjunction, conditionalization, temporalization, and modalization.

Pratt's and Casati and Varzi's circumspection about such connectives might thus be seen as a tacit endorsement of the common assumption that maps lack the

capacity to express such sorts of conditions and modifications—although it might also just derive from their focus on simple cases in an initial proof of principle for a formal cartographic semantics. However, extending the discussion of layering from section 4.3, it is important to note that much as we imagined a map in which texture indicates degree of credence in the reliability of information about the density and quantity of rainfall, so too it is possible for maps to employ markers that achieve the effect of disjoining, conditionalizing, or modalizing one or more states of affairs. So, for instance, disjunction might be represented by coloring disjuncts in stripes or with flashing lights; or a solid color might be used to indicate the antecedent of a conditional and stripes or flashes to represent the consequent (Camp 2007, 164). Again, though, this effect appears to be achieved semantically rather than syntactically—in virtue of constraints on coherence originating from the values assigned to the marks, rather than in virtue of the way the marks themselves are combined.

Putting aside the question of whether and how maps can represent more complex relations between basic, independently variable bits of information, I will focus in this section on the question of whether conjunction offers a plausible model for capturing the relational structure among located map markers, as Pratt and Casati and Varzi both assume. I will argue that it does not: the conjunctive model fails to explain the direct, explicit way in which maps represent multiple spatial relations simultaneously as an integrated whole.

Recall that Pratt treats each located marker as a symbol-pair (e.g. $\langle \square, (x, y) \rangle$), with a spatial interpretation function μ assigning the second element to a world-location. However, his semantics is largely silent about μ , and in particular about how map coordinates are themselves represented and interpreted. As far as the syntax, semantics, and meta-semantics go, a map simply consists of a list of ordered symbol-pairs. As a result, any constraints there might be on relations among map-locations, and on the implications of those constraints for relations among the corresponding world-relations—for instance, that $(x+1, y+1)$ is a proximate diagonal neighbor of (x, y) , and that therefore $\mu(x+1, y+1)$ must also be a proximate diagonal neighbor of $\mu(x, y)$ —must be encoded meta-semantically as conditions on μ , much as constraints on the interpretation of symbols, such as the fact that markers denoting individuals can only be uniquely tokened on a given map, are.

I am happy to concede that, given suitable constraints on μ , Pratt's semantics can derive the appropriate truth-conditions for any given map within its representable class. And if so, then a suitably constrained version of μ plus the substantive interpretation function and the atomic symbol-pairs will also appropriately support the extraction of any given piece of information about spatial relations from a map. The question we need to ask is *how* the propositional model, and conjunction in particular, encodes and recovers that relational spatial information. On Pratt's account, spatial relations among particular denoted objects/properties—for instance, that the post office is inside the strip mall, that it is closer to the park than to the grocery store, or that it is on the way to the grocery store from the park—will be

generated by inference from the individual atomic symbol-pairs plus a set of general propositions—which are themselves either explicitly encoded in, or else derived from the constraints on, μ —about relations among the locations that are denoted by the locational component of the symbol-pair. That is, information about spatial relations among object/properties is inferentially derived, in part by appealing to meta-semantic constraints. This is exactly what we should predict from a propositional analysis, on which predicates just ascribe properties to individual locations. But intuitively, and as I will argue later, this is not how maps work.

Unlike Pratt, Casati and Varzi do explicitly invoke maps' topological structure—the fact that map-locations themselves stand in relations of containment and contiguity—and explicitly constrain their interpretation function to preserve that structure. (While they appeal only to topological structure, their interpretation function can easily be supplemented to incorporate more robust geometrical features like direction and distance.) Moreover, because they treat all applications of a given color to map-regions (plus the non-coloring of complementary regions) as a single predicate, there is an important sense in which for them, unlike for Pratt, spatial relations among distinct instances and absences of any given object/property F are represented directly, rather than being derived. However, insofar as Casati and Varzi analyze entire maps as conjunctions of monochromatic atomic map-stages, they are still committed to treating spatial relations among *distinct* object/properties as derived.

The distinction between direct and derived information has some intuitive appeal. But what does it mean to say that a predicative analysis of conjunction in maps is committed to treating the spatial relations among object/properties that are denoted by distinct symbols or colors as derived rather than direct? Much as predication just *is* an operation that delivers a truth-value based on the value of its input, so too conjunction just is an operation that takes multiple propositions¹¹ and delivers truth if and only if they are all true. This means, first, that the conjuncts are treated as distinct representational units, and second, that the truth-value of the whole depends only on truth-values of the parts: conjunction itself sets no further constraints on how those conjuncts are related or represented.

Given this, as far as Pratt's analysis goes all that a map represents directly is that all of the constituent atomic propositions expressed by each of the map's symbol-pairs are true; spatial relations among denoted object/properties are merely guaranteed by a background condition on what it takes for those propositions to be true. Any other way of encoding the individual symbol-pairs which preserves the meta-semantic constraints on μ is representationally equivalent. But intuitively, a map makes the spatial relations among denoted object/properties directly available in a way that a list of symbol-pairs plus μ does not (Camp 2007, 161).

¹¹ Or other type-similar expressions; our focus here is just on conjunction as an operation that conjoins atomic propositions, as per Pratt's and Casati and Varzi's analyses.

Likewise, as far as Casati and Varzi's analysis of conjunction goes, a series of atomic map-stages that are placed side-by-side on separate sheets of paper is representationally equivalent to a series of map-stages printed on top of one another. But intuitively, only the latter genuinely compiles the information into a unified whole; and only such a compilation directly represents the spatial relations among distinct object/property types. As it's sometimes put, what would be an active inference from premises to conclusion in a predicate calculus comes along as a "free ride" in maps and diagrammatic systems (Shimojima 1996; Larkin and Simon 1987; and Shin 1994, 2003).

We might still worry that while appeals to ease of use, or to derived versus direct access, are highly intuitive, they are also suspiciously psychological, even phenomenological. Interpreting representations obviously requires significant processing; but much of it is not introspectable, and what is effortful for one user may be automatic for another. In order to make more principled sense of the distinction between direct representation and inferred information, it helps to shift focus from the generation of static truth-conditions to the dynamics of updating representations in light of new information.

Maps are integrated representational wholes, in which a single token (which may itself be complex, in combining marks from multiple "incompatibility classes") represents how things are at a given location, and in which spatial relations among denoted object/properties are represented by the spatial relations among the constituent tokens. Because spatial relations in the map represent spatial relations in the world, moving or otherwise altering any given marker on a single map automatically updates the represented spatial relations among that denoted object/property and all of the other represented object/properties. It is physically impossible to alter the map in a way that fails to update those relations. A map that failed to fully update the represented spatial relations among its denoted object/properties would be ill-formed in a very strong sense: it could not be drawn. By contrast, altering the object/property predicated of a location within one sentence still leaves further operations to be performed in order to propagate that new information to other propositions in the set. It is entirely possible, indeed all too common, to fail to fully update the represented spatial relations; and the resulting failure is simply one of semantic inconsistency, not of syntactic incoherence.¹²

Because Pratt treats each symbol-pair as a separate proposition—just as a predicative theorist should—his analysis appears to wrongly predict the possibility of partial updates. Casati and Varzi's analysis captures the phenomenon of automatic updating for represented spatial relations among instances of a single denoted object/property (and its absence), because they analyze all instances of a given color (and its absence) as a single predicate, and because the regions to which that predicate applies are defined by actually standing in the relevant topological (geometric, etc.) relations to one another.

¹² The syntactic holism of cartographic spatial structure thus contrasts with the pragmatic, or at most semantic, holism manifested by the Absence Intuition.

However, what about spatial relations among distinct atomic map-stages? If a series of side-by-side atomic map-stages linked by conjunction is indeed possible, as appears to be permitted by their account, then we should predict possible failures of full update: after all, faced with the information that *F* is located otherwise than previously thought, a printer might merely replace the first, *F*-colored sheet in the series, but inadvertently leave the rest of the map-stages, with their old holes where *F* would be colored in, unchanged.

Casati and Varzi can rightly object that their analysis does indeed classify such a series of merely partially updated map-stages as syntactically ill-formed, since map-regions are defined by their topological relations, and the series of partially updated map-stages fails to define a well-formed set of regions (1999, 193). Indeed, they should object that even a fully consistent set of side-by-side maps is ill-formed, because topological relations among the represented world-regions are not preserved by the separate map-stages: for instance, a land-region that borders a given water-region is not represented by a green map-region that is actually contiguous with a blue map-region, if the green map-region and the blue map-region are on different sheets of paper.

However, the relevant question is not just whether their overall semantics delivers the right verdicts about well-formedness—although that is important. Our primary question is whether Casati and Varzi's analysis of that semantics in terms of the construction of a propositional structure via conjunction is warranted. The contrast between automatic and possibly partial updating provides a way both to diagnose the presence of multiple distinct but conjoined representations and also to see how the difference between direct and derived information might matter.

Derivation or inference is a mechanism for extracting new information by connecting previously distinct representations. Conjunction is a particular way of connecting distinct representations into a whole: sequentially and merely in virtue of their truth-values. As such, it is compatible with distinct side-by-side map-stages, and produces the possibility of merely partial update. By contrast, the sort of automatic updating among spatial relations that is exhibited by moving a marker on a map is only possible because all of the markers already occur within a single map, so that a single marker token simultaneously contributes to representing many distinct spatial relations. Given that maps exhibit automatic updating, the combinatorial operation by which they unite isolable bits of information into a larger whole must be more robust and specific than conjunction. Plausibly, it is a form of compilation that actually places the atomic map-parts into the represented spatial relations with each other.

It is a virtue of Casati and Varzi's account that their semantics appeals directly to actual spatial relations among map-regions. But to the extent that they use these relations to argue either that no actual compilation is necessary (because the map-regions themselves already stand in the appropriate relations and they are merely being colored in stages), or else that a specific form of compilation is required (of successively printing a single sheet of paper), they thereby undermine their claim that

maps are constructed by the conjunction of a set of atomic predicative propositions. On the other hand, to the extent that they, like Pratt, do not appeal to these more robust features of compilation, then they make predictions about how maps function that are not borne out.

Knowing that a represented domain *has* spatial (topological, geometrical, etc.) structure already itself licenses a rich set of assumptions about relations among elements in that structure—most notably, about containment, contiguity, direction, and distance; and these in turn license a rich set of inferences about the effects of altering structural relations among those elements. These assumptions and inferences can be encoded and implemented in any format; and they can be used to generate appropriate illustrations of the resulting state of affairs. However, maps do more than just illustrate: they represent those structural relations among elements by instantiating them. As Rescorla (2009a, 197–8) puts it, “A map’s geometric structure is not just another element to be listed alongside its markers and coordinates . . . Rather, the markers and coordinates *stand in* geometric relations, and those relations bear representational import.” A direct consequence is that permutations of those relations within the vehicle automatically instantiate the relevant permutation to how the world is represented as being.

If maps don’t have predicative or propositional structure, then how should a formal cartographic semantics go? We can now say that rather than employing a combinatorial operation that is digital, universal, and asymmetrical, we need one that is holistic, spatial, and symmetrical, and that permits automatic updating by isomorphic permutation. Because a predicative analysis first identifies denoted locations, then assigns denoted objects and properties to those locations, and finally conjoins those located object/properties, it inherently assigns a derivative status to spatial relations among denoted object/properties. By contrast, as Rescorla says, we should begin with the fact that maps’ markers themselves stand in spatial relations.¹³

Casati and Varzi almost achieve this, by taking seriously the fact that the regions to which markers are applied are themselves spatially related. But this compromises their claim to a propositional analysis. More importantly, their predicative analysis commits them to treating map-regions as being individuated and interpreted prior to and independently of the markers that are applied to them. If we instead treat a map as a spatial distribution of markers, we can hold that an entire map is true just in case there exists a region in the world at which the objects and properties denoted by the markers are distributed in a way that preserves the relevant spatial (topological, geometrical, etc.) relations of that distribution. Implementing this suggestion, however, is a task for another day.

¹³ Familiar maps on paper employ spatial structure to represent spatial structure by replicating relevant aspects of that structure. Other maps—say, in a computer, or in a brain—implement geometrical, topological, and other, broadly ‘spatial’ structures in other ways. I think they are still maps if they have the function of representing a domain that has such a structure, and if their constituent elements actually stand in relations such that appropriate operations on those relations automatically produce structurally isomorphic representational effects.

5 Conclusion

The semantics offered by Pratt and by Casati and Varzi are formal: they offer conditions for constructing and determining the truth-conditions of entire maps, stated in terms of relevant, formally individuated features.¹⁴ As Blumson (2012) emphasizes, a formal semantics of this sort suffices to address the traditional requirements invoked by defenders of the Language of Thought Hypothesis for a ‘rational’ cognitive architecture: systematicity and productivity. By contrast, it is not obvious that a similarly formal semantics can be offered for pictures, let alone perception.¹⁵ In particular, the representational significance of images often appears to be highly local and context-dependent: changes to the marks which make little semantic difference at one location produce significant semantic difference, or destroy the representation altogether, at another.

However, the possibility of such a formal cartographic semantics does not suffice to demonstrate that maps have propositional structure, or the same sort of structure as language, as Blumson (2012, 414) also claims. Rescorla (2009a) argues that the Absence Intuition demonstrates that the semantics proposed by Casati and Varzi is not truly predicative. Although I am less sanguine about Absence, I have argued that there are other reasons to reject a predicative analysis, along with a propositional interpretation more generally.

The contrast between maps and propositional representations shows up especially clearly when we consider the dynamics of map construction, modification, and compilation. This suggests the need for not just a formal cartographic syntax and semantics, but for a logic of map dynamics. As Pratt (1993) and MacEachren (2004) emphasize, the interaction between map semantics and pragmatics is highly complex. But as linguists and philosophers of language have increasingly recognized, the same also holds for language, with many linguistic features having the conventional function of tracking, expressing, and modifying both truth-conditional and non-truth-conditional aspects of both the objective and conversational context. One especially fertile area for future research concerns the different ways in which conventional semantic meaning and occasion-specific pragmatic use interact in language and in maps.

In arguing that maps do not have propositional structure, I have relied both upon a restricted class of maps that are clearly formalizable, and upon an idealized analysis of language as a predicate calculus. Natural languages depart in deep and pervasive ways from the clean, universal systematicity of traditional formal logics. But this has not prevented semanticists from developing detailed, predictive formal analyses of

¹⁴ Pratt (1993) argues that figure-ground ambiguity rules out the possibility of a fully formal semantics. Although I have not addressed his concerns here, I believe they can be resolved by distinguishing more clearly between cartographic semantics and pragmatics than he thinks possible, and by taking more seriously the sense in which an overall cartographic system includes principles for both constructing and manipulating maps.

¹⁵ Though see Greenberg (2013), among others, for a start.

apparently recalcitrant natural-language phenomena using tools and methods that are at least inspired by, and arguably extensions of, a predicate calculus. I have argued that the basis for a cartographic logic should be quite different; but I would urge a similar spirit of constructive optimism in this pursuit as for language.

The restricted class of formal maps and the idealized case of predicate logic also make the contrast between maps and language appear especially stark. When we turn to real maps and languages, we find much more intriguing complexity. Most obviously, many ordinary maps employ words, as well as shapes that purport to replicate the shapes of roads and lakes. In the other direction, Phillippe Schlenker (Schlenker et al. 2013) has argued that sign languages systematically employ, modify, and relate iconic or pictorial elements within linguistic structure. We also regularly encounter other representational formats—diagrams, graphs, musical scores—that are more or less map-like or language-like in different ways.

The point of contrasting maps and languages is not to erect sharp taxonomic boundaries or identify deep natural kinds (Johnson 2015). Rather, my aim has been first, to argue that not all representations with a formally characterizable functional structure are therefore simply notational variants on language. Different representational systems employ different degrees and dimensions of resemblance and arbitrariness in the interpretive principle(s) that assign values to representational constituents, and also different degrees and dimensions of abstraction and holism in the syntactic principle(s) that combine constituents into larger wholes (Camp 2015). My second aim has been to identify more clearly and specifically which dimensions of difference make a difference, and how. This will help us, among other things, to better understand how particular representational systems function out in the world, to build systems that more efficiently meet our representational needs, and to make better inferences from behavioral evidence about what cognitive architecture a given agent may be employing.

References

- Anderson, John (1978) Arguments Concerning Representations for Mental Imagery. *Psychological Review*, 85: 249–77.
- Asher, Nicholas, and Alex Lascarides (2003) *Logics of Conversation* (Cambridge: Cambridge University Press).
- Bermudez, Jose Luis (2003) *Thinking without Words* (Oxford: Oxford University Press).
- Blumson, Ben (2012) Mental Maps. *Philosophy and Phenomenological Research*, 85(2): 413–34.
- Boyd, Richard (1999) Homeostasis, Species, and Higher Taxa. In R. Wilson (ed.), *Species: New Interdisciplinary Essays* (Cambridge, MA: MIT Press), 141–85.
- Camp, Elisabeth (2004) The Generality Constraint, Nonsense, and Categorical Restrictions. *Philosophical Quarterly*, 54: 209–31.
- Camp, Elisabeth (2007) Thinking with Maps. *Philosophical Perspectives*, 21(1): *Philosophy of Mind*, ed. J. Hawthorne (Oxford: Wiley-Blackwell), 145–82.

- Camp, Elisabeth (2009) A Language of Baboon Thought? In R. Lurz (ed.), *The Philosophy of Animal Minds* (Cambridge: Cambridge University Press), 108–27.
- Camp, Elisabeth (2015) Logical Concepts and Associative Characterizations. In E. Margolis and S. Laurence (eds), *The Conceptual Mind: New Directions in the Study of Concepts* (Cambridge, MA: MIT Press), 591–621.
- Camp, Elisabeth, and Eli Shupe (2017) Instrumental Reasoning in Non-Human Animals. In J. Beck and K. Andrews (eds), *The Routledge Handbook of Philosophy and Animal Minds* (London: Routledge), 100–8.
- Casati, Achille, and Roberto Varzi (1999) *Parts and Places: The Structures of Spatial Representation* (Cambridge, MA: MIT Press).
- Cheney, Dorothy, and Robert Seyfarth (2007) *Baboon Metaphysics: The Evolution of a Social Mind* (Chicago: University of Chicago Press).
- Chierchia, Gennaro (2004) Scalar Implicatures, Polarity Phenomena and the Syntax/Pragmatics Interface. In Adriana Belletti (ed.), *Structures and Beyond* (Oxford: Oxford University Press), 39–103.
- Chomsky, Noam (1995) *The Minimalist Program* (Cambridge, MA: MIT Press).
- Collins, John (2011) *The Unity of Linguistic Meaning* (Oxford: Oxford University Press).
- Crane, Tim (2009) Is Perception a Propositional Attitude? *Philosophical Quarterly*, 59: 452–69.
- Davidson, Donald (1999) The Emergence of Thought. *Erkenntnis*, 51: 7–17.
- Devitt, Michael (2006) *Ignorance of Language* (Oxford: Clarendon Press).
- Fodor, Jerry (2007) The Revenge of the Given. In Brian P. McLaughlin and Jonathan D. Cohen (eds), *Contemporary Debates in Philosophy of Mind* (Oxford: Blackwell), 105–16.
- Fodor, Jerry, and Zenon Pylyshyn (1988) Connectionism and the Cognitive Architecture of Mind. *Cognition*, 28: 3–71.
- Frege, Gottlob (1892) On Concept and Object. In *Translations from the Philosophical Writings of Gottlob Frege*, ed. Peter Geach and Max Black (Oxford: Basil Blackwell, 1977), 42–55.
- Greenberg, Gabriel (2013) Beyond Resemblance. *Philosophical Review*, 122(2): 215–87.
- Hauser, Mark, Noam Chomsky, and Tecumseh Fitch (2002) The Language Faculty: What is it, Who has it, and How did it Evolve? *Science*, 298: 1569–79.
- Heim, Irene, and Angelica Kratzer (1998) *Semantics in Generative Grammar* (Oxford: Blackwell).
- Johnson, Kent (2004) On the Systematicity of Language and Thought. *Journal of Philosophy*, 101(3): 111–39.
- Johnson, Kent (2015) Maps, Languages, and Manguages: Rival Cognitive Architectures? *Philosophical Psychology*, 28(6): 815–36.
- Kehler, Andrew (2002) *Coherence, Reference, and the Theory of Grammar* (Palo Alto, CA: CSLI Publications).
- King, Jeffery C. (2007) *The Nature and Structure of Content* (Oxford: Oxford University Press).
- Kulvicki, John (2015a) Analog Representation and the Parts Principle. *Review of Philosophy and Psychology*, 6: 165–80.
- Kulvicki, John (2015b) Maps, Pictures, and Predication. *Ergo*, 2(7): 149–73.
- Larkin, Jill, and Herbert Simon (1987) Why a Diagram is (Sometimes) Worth 10,000 Words. *Cognitive Science*, 11(1): 65–99.
- Lewis, David (1968) Counterpart Theory and Quantified Modal Logic. *Journal of Philosophy*, 65: 113–26.

- MacEachren, Alan (2004) *How Maps Work: Representation, Visualization, and Design* (New York: Guilford Press).
- Millikan, Ruth Garrett (1998) A Common Structure for Concepts of Individuals, Stuffs, and Real Kinds: More Mama, More Milk, and More Mouse. *Behavioral and Brain Sciences*, 21: 55–100.
- Pratt, Ian (1993) Map Semantics. In *Spatial Information Theory: A Theoretical Basis for GIS Lecture Notes in Computer Science*, vol. 716 (Berlin: Springer-Verlag), 77–91.
- Pylyshyn, Zenon (2003) *Seeing and Visualizing: It's Not What you Think* (Cambridge, MA: MIT Press).
- Quine, Willard V. O. (1995) *From Stimulus to Science* (Cambridge, MA: Harvard University Press).
- Rescorla, Michael (2009a) Predication and Cartographic Representation. *Synthese*, 169: 175–200.
- Rescorla, Michael (2009b) Cognitive Maps and the Language of Thought. *British Journal of the Philosophy of Science*, 60: 377–407.
- Rey, Georges (1995) A Not 'Merely Empirical' Argument for the Language of Thought. *Philosophical Perspectives*, 9: 201–22.
- Roberts, Craige (1996/2012) Information Structure in Discourse: Toward an Integrated Formal Theory of Pragmatics. *Semantics and Pragmatics*, 5: 1–69.
- Schlenker, Philippe, Jonathan Lamberton, and Mirko Santoro (2013) Iconic Variables. *Linguistics and Philosophy*, 36(2): 91–149.
- Schlichtmann, Hansgeorg (1985) Characteristic Traits of the Semiotic System "Map Symbolism". *Cartographic Journal*, 22: 23–30.
- Shimajima, Atsushi (1996) Operational Constraints in Diagrammatic Reasoning. In J. Barwise and G. Allwein (eds), *Logical Reasoning with Diagrams* (Oxford: Oxford University Press), 27–48.
- Shin, Sun-Joo (1994) *The Logical Status of Diagrams* (Cambridge: Cambridge University Press).
- Shin, Sun-Joo (2003) *The Iconic Logic of Peirce's Graphs* (Cambridge, MA: MIT Press).
- Slovan, Aaron (1978) *The Computer Revolution in Philosophy: Philosophy, Science and Models of Mind* (Atlantic Highlands, NJ: Harvester Studies in Cognitive Science Humanities Press).
- Sober, Eliot (1976) Mental Representations. *Synthese*, 33: 101–48.
- Stalnaker, Robert (1976) Propositions. In Alfred MacKay and Daniel Merrill (eds), *Issues in the Philosophy of Language* (New Haven and London: Yale University Press), 79–91.
- Stalnaker, Robert (2014) *Context* (Oxford: Oxford University Press).
- Szabo, Zoltan (2012) The Case for Compositionality. *The Oxford Handbook of Compositionality* (Oxford: Oxford University Press), 64–80.