Phenomenal Structuralism*

David J. Chalmers

[It is] possible to characterize phenomenal properties in structural terms, and there are views on which all phenomenal truths and even all truths are [structural]. We might call such views *phenomenal structuralism*.

Carnap's view for much of the *Aufbau* is a sort of phenomenal structuralism. He characterizes experiences wholly in terms of the relation of phenomenal similarity among them. [...] Goodman's basic worry is that a single phenomenal similarity relation is not rich enough to recover the full character of phenomenology. Here, I think that the phenomenal structuralist has various options. Even if Carnap's account of phenomenal structure is deficient, other characterizations may be available.

At this point, a natural move for the phenomenal structuralist is to move from a single phenomenal similarity relation to multiple such relations, each corresponding to different respects or dimensions of phenomenal similarity. For example, distinct color experiences can be similar in red-green respects (having the same amount of redness), yellow-blue respects, or brightness respects. Visual experiences can also be similar in various spatial respects. Still, residual problems in the spirit of Goodman's may remain. For example, mere similarity information among a set of experiences may not suffice to recover structure when the number of total experiences is limited. If we are told that there are two experiences that are different in all relevant respects, this seems to leave their character underdetermined. A move to graded similarity relations may help to some extent, but problems will remain. The picture also gets complicated when we move to more complex experiences, such as that of a full visual field.[1]

*Excerpted from *Constructing the World* (Oxford University Press, 2012), pp. 412-22.

[1]To handle a visual field in this framework (where visual fields have the structure characterized below), one might appeal to an entire manifold of graded similarity relations, three for each location in the visual field, corresponding to similarity in the redness, blueness, and brightness respects at that location. One could then have three higher-order similarity relations that hold along these relations: an on-off relation that holds iff the lower-order relations involve sameness in the same color respect, and two graded relations that measure similarity of the two points corresponding

A better move for the phenomenal structuralist is to move from respects of phenomenal similarity to parametric information about locations along phenomenal dimensions. All this is most straightforwardly done by using quasi-mathematical specifications of phenomenal states analogous to the mathematical specifications of physical states discussed above.

For concreteness, let us adopt the fiction that phenomenal states are entirely visual and that (as on Carnap's own model) visual phenomenal states involve only the distribution of phenomenal colors in a two-dimensional visual field. On this model, locations in the visual field can be represented by (x, y) co-ordinates, where x is intuitively a left-right co-ordinate and y is an up-down co-ordinate. Phenomenal colors can be represented as (a, b, c) co-ordinates, corresponding to locations on a red-green axis, a yellow-blue axis, and a brightness axis. We can assume that each of these five co-ordinates is constrained to lie between -1 and 1 inclusively. Then if A is $[-1, 1]$, the set of real numbers between -1 and 1 inclusively, locations in the visual field can be represented as members of A^2 and phenomenal colors can be represented as members of A^3. The total phenomenal state of a subject can then be represented as a function from A^2 to A^3, in effect assigning a phenomenal color to each location in the visual field.[2]

What sort of additional constraints are needed to characterize the dimensions? The strongest constraints specify explicitly that the three dimensions of A^3 represent degrees of phenomenal redness, phenomenal blueness, and brightness respectively, and that the two dimensions of A^2 represent location on a left-right axis and an up-down axis in the visual field. This treatment makes

to the two lower-order relations in left-right respects and up-down respects respectively. From here, given a sufficiently rich set of total experiences one could recover much of the geometric structure of the parametric model below. The similarity-based model is arguably somewhat closer to the spirit of Carnap's method of "quasi-analysis", but the parametric model is much more straightforward, and is also somewhat more powerful (it allows one to more easily capture the distinct status of unique hues, for example).

[2]Of course visual experiences have a much more complex structure than this, but the phenomenal structuralist can reasonably hold that this structure can be characterized in more complex mathematical terms. The same goes for other perceptual experiences. Some putative experiences, such as the experience of thinking, or the experience of perceptually recognizing a given person, seem harder to characterize in mathematical terms, but the phenomenal structuralist could adopt a "thin" view of experience on which it is exhausted by the experience of low-level features such as color, shape, and location. The temporal aspects of consciousness raise further issues: here the phenomenal structuralists might appeal to phenomenal temporal qualities analogous to phenomenal spatial qualities (alternatively, a phenomenal/temporal structuralist could appeal to the temporal properties of experience). Phenomenal structuralism can also be adapted to a representationalist framework, characterizing experiences in terms of phenomenal awareness of certain properties (color and spatiotemporal properties, for example), where these properties are characterized in the structural terms above. Such a framework may invoke a primitive concept of awareness as well as concepts for relations among various primary and secondary qualities, so we can think of it as phenomenal/quality structuralism.

it reasonably plausible that the full character of the total phenomenal state will be [specified], but its credentials as a variety of structuralism are dubious. The notions of phenomenal redness and phenomenal blueness are intuitively far from structural notions, and they seem to pose the sort of problems for objectivity and communicability that structuralist views are supposed to avoid. It is arguable that something similar applies to primitive concepts of (phenomenal) left and right, and up and down.

The weakest constraints here will simply specify that the dimensions are *phenomenal* dimensions, and will say nothing beyond that. This sort of specification requires an unanalyzed notion of phenomenology, or of a phenomenal dimension, just as nomic and spatiotemporal specifications require unanalyzed notions of lawhood and spacetime, but it plausibly counts as (weakly) structural to roughly the extent that these do. It likewise counts as structural to the extent that Carnap's own specification in terms of phenomenal similarity does. The most obvious problem for the weak model is one of underdetermination. It is natural to suggest that merely specifying that a phenomenal state is characterized by a certain function from A^2 to A^3 leaves open whether it involves phenomenal color (for example) at all. For example, perhaps there can be a phenomenal state isomorphic to this one involving the distribution of auditory qualities in a two-dimensional field. Still, a certain sort of structuralist about phenomenal properties will deny that this is possible, holding that the difference between visual and auditory phenomenal properties is ultimately a matter of their structure.[3]

An intermediate model specifies that the dimensions of A^3 are color dimensions and that the dimensions of A^2 are spatial dimensions, but does not specify their nature beyond this. As one might expect, the intermediate option is subject to watered-down versions of the objections to the weak and strong models. First: the notions of phenomenal color and phenomenal space are not structural notions, although perhaps they are less objectionable to a structuralist than phenomenal redness or phenomenal blueness. Second: specifying a phenomenal color in terms of its locations along three dimensions leaves open whether it is phenomenal redness or phenomenal blueness. For example, unique phenomenal redness and unique phenomenal blueness can both be represented as $(1, 0, 1)$ in the different co-ordinate systems: the [red, blue, brightness] and the [blue,

[3]Leitgeb (2011) makes a different phenomenal structuralist proposal on which there is a single basic property of "qualitative overlap" applied to sets of experiential tropes, for example when they all involve the same shade of red in the same area of the visual field. In effect overlapping sets will correspond to Carnap's phenomenal quality circles while avoiding Goodman's problems of companionship and imperfect community. The respects of overlap are specified only as phenomenal respects, so this specification imposes a version of the weak constraints in the text.

red, brightness] systems respectively. So the mere claim that a given phenomenal state can be represented as $(1, 0, 1)$ along phenomenal color axes does not enable one to determine whether it is phenomenal redness or phenomenal blueness.

The second underdetermination problem is a version of the problem of the inverted spectrum, which plagues all structuralist accounts of phenomenology. A structuralist may reply by appealing to further structural constraints that distinguish the various dimensions. But it is arguable that related problems will always arise. One way to see this is to note that Frank Jackson's Mary in her black-and-white room (chapter 3) could in principle be told any set of mathematical and structural facts about the phenomenal state someone is in when they see roses, but she still would not be in a position to know what it is like to see roses. So she will not be in a position to know that a certain mathematically specified state is a certain sort of phenomenal redness (specified under a pure phenomenal concept). Even if she is told that certain parameters represent locations along phenomenal dimensions, or along phenomenal color and space dimensions, this will not help. If she is told that these represent locations along phenomenal redness and blueness dimensions, where she somehow has mastered the pure phenomenal concepts of phenomenal blueness and redness, then she will be in a position to know what the state is like. But assuming that this sort of characterization in terms of specific phenomenal dimensions does not count as structural, then the considerations here suggest that certain phenomenal truths are simply not scrutable from truths about phenomenal structure, and that any form of phenomenal structuralism is false.

Still, a phenomenal structuralist is likely to be deflationary about phenomenal knowledge. Carnap himself would probably not have been too worried about inverted spectrum hypotheses or about Mary's new phenomenal knowledge: these are precisely the sorts of putative hypotheses and knowledge that he wants to reject as meaningless. I think that the most consistent line for the phenomenal structuralist is to adopt only the weak constraints on which it is simply specified that the relevant parameters are phenomenal dimensions, and to deny the claim that this phenomenal structure leaves some phenomenal truths underdetermined.[4] This model shares much of the spirit of Carnap's model in the *Aufbau*. In effect, Carnap's single relation of phenomenal similarity has been expanded into many such relations (corresponding to unspecified respects of similarity), and these have then been reconstrued as parametric phenomenal dimensions.

[4]If phenomenal dimensions are too cheap, then a version of Newman's problem will arise. For example, if any function from total phenomenal states to [-1, 1] counts as a phenomenal dimension, then given any set of phenomenal states of cardinality no greater than the continuum, there will be phenomenal dimensions under which these states can be mapped to A^3. Something similar applies to mappings from phenomenal states to functions from A^2 to A^3. These

[...] My own view is that no robust version of structuralism [about reality] is correct, because of problems associated with consciousness. If a structural specification is something that Mary can grasp from inside her black and white room, [...] then it is likely that many phenomenal truths (for example, that what it is like to see roses is such-and-such) will be inscrutable from this specification. And if a specification cannot be grasped by Mary inside her black and white room (for example, because it uses the notion of phenomenal redness), then it will be structural in at best a highly attenuated sense.

Still, to say this much is compatible with holding with Carnap that structural expressions play a special role in science and in communication. One might hold that nonstructural concepts, such as that of phenomenal redness, pose special problems for science and communication, in that grasp of these expressions depends on one's prior history, and in that one cannot be certain that others are using their corresponding expressions to express the same concept. Claims like these need to be formulated carefully, as I think it is plausible that science can say a good deal about phenomenal redness, as when phenomenal colors are decomposed along three basic dimensions. Still, it is plausible there is a core part of the science of color experience that can be understood by Mary inside her black-and-white room. The same goes for human scientists studying nonhuman sensory modalities. This core part can be cast largely in terms of phenomenal structure. Grasp of specific phenomenal dimensions such as phenomenal redness certainly enriches our grasp of the science, but it is not obviously necessary in order for the science to proceed. For that purpose, structural notions are enough.

All this connects interestingly to the thesis of structural realism in the philosophy of science, which says roughly that scientific theories concern the structure of reality.[5]

It is standard to distinguish ontological structural realism, according to which reality (as described by science) is wholly structural, from epistemological structural realism, according to which we can know (through science) only structural aspects of reality. As an epistemological thesis, structural scrutability is more akin to epistemological structural realism, but it is somewhat stronger. We might see structural scrutability as a form of *conceptual* structural realism, holding that the only true hypotheses that we can entertain about reality are structural (here I assume that a hypothesis can be entertained iff it can be expressed, and that any truth scrutable from a structural truth is itself a structural truth). Conceptual structural realism is intermediate in strength between

structural specifications will be satisfied by any set of phenomenal states of small enough cardinality. To avoid this problem, there must be constraints on phenomenal dimensions, perhaps requiring them to correspond to natural or fundamental aspects of phenomenology.

the ontological and epistemological varieties. It is plausibly entailed by ontological structural realism but not vice versa: the two will come apart if there is a nonstructural character to reality that we cannot entertain hypotheses about. And it plausibly entails epistemological structural realism but not vice versa: the two will come apart if there are true nonstructural hypotheses that we can entertain without knowing them to be true.

[...] The consciousness-based objections to structuralism do not apply to all versions of structural realism. It is common to cast structural realism as a claim about the *non-observational* aspects of reality. These views in effect allow that reality more broadly can be characterized in structural and observational terms, in effect allowing certain observational expressions in the base. On this reading, structural realism is somewhat weaker than [a more global structuralism], which does not make special allowance for observational expressions. Corresponding, even if phenomenal truths count against [global structuralism] (as I have suggested), they need not count against structural realism so understood, as long as one counts phenomenal expressions as observational expressions. Something similar applies to secondary qualities: even if these pose an obstacle to structural scrutability, they do not pose an obstacle to this variety of structural realism.

[...] Even if one accepts nonstructural aspects of reality such as consciousness, quiddities, and primitive spacetime, it is plausible that there will at least be a structural core to our scientific theories. As in the case of the science of consciousness above, this structural core will capture the structural aspects of these phenomena and abstract away from nonstructural aspects. Such a theory will not be a complete theory of reality, but it will do much of the work that a complete theory can do.

[Note added 2022: I might have added that there are versions of each of ontological structural realism, epistemological structural realism, and conceptual structural realism in the domain of consciousness. The option of epistemological or conceptual structural realism without ontological structural realism is less well-motivated in the domain of consciousness than in the case of the physical world, since the nonstructural phenomenal qualities such as phenomenal redness that cause problems for ontological structuralism are plausibly knowable and conceptualizable, unlike nonstructural quiddities in physical reality.

However, we can distinguish a fourth sort of structural realism, *methodological structural*

[5]Structural realism was introduced under that name by Grover Maxwell in "Structural Realism and the Meaning of Theoretical Terms" (1970). For the distinction between epistemological and ontological structuralism, James Ladyman's "What is Structural Realism?" (1998). For discussions of Newman's problem as a problem for structural realism, see Ainsworth 2011, Ketland 2004, and Melia and Saatsi 2006.

realism, that seems especially apt in the domain of consciousness. This is the Carnap-inspired idea that methodologically, science should focus on structural properties (even if there are nonstructural properties, even if they are knowable, and so on). In the case of consciousness, methodological structural realism says that the science of consciousness should focus on the structural properties of consciousness. I am sympathetic with a somewhat attenuated version of this view as in the final paragraph above: nonstructural qualitative properties are not off-limits to science (one could theorize about the neural correlates of phenomenal redness, for example), but there is a central core of the science of consciousness that focuses on its structural properties.]