Modeling Mental Qualities

Andrew Y. Lee

Rice University, Department of Philosophy

andrewyuanlee@gmail.com

http://andrewyuanlee.com/

ABSTRACT:

This paper develops a formal framework for modeling mental qualities that is more powerful than existing models. I begin by explaining why standard models of mental qualities of mental qualities cannot capture *precision structure*, or the phenomenal contrast between seeing an object as crimson (e.g., in foveal vision) versus seeing an object as merely red (e.g., in peripheral vision). Then I explain how my new framework provides a unified way of modeling similarity and precision structure, how different formal constraints within the framework correspond to different classes of theories about the space and structure of imprecise qualities, how the framework differentiates two distinct dimensions of phenomenal similarity, how empirical methods can be used to construct models for particular domains of mental qualities, and how the framework can be used to model the intransitivity of perceptual discrimination. An upshot is that the structure of the mental qualities of conscious experiences is fundamentally different from the structure of the physical qualities of external objects.

INTRODUCTION

How do we model mental qualities, such as those involved in seeing blue, feeling pain, or smelling cinnamon? The standard approach is to target a particular domain, such as color experience, and to develop a *quality-space model* for the mental qualities of that domain. These models represent mental qualities via points in multidimensional spaces, where points that lie closer in the space represent mental qualities that are more similar to each other. For example, in the canonical three-dimensional model of color qualities, any particular quality of color experience can be specified via its values along the hue, saturation and brightness dimensions, and color qualities that are more similar correspond to points closer in the space. The result is a systematic representation of the structure of color qualities.

This framework for modeling mental qualities is highly promising, since any domain of mental qualities can be structured via similarity relations. While the model for color qualities is the most developed, in recent years quality-space models have also been proposed for temporal experience, pain, auditory experience, and olfactory experience.[1] And even for domains that are difficult to empirically investigate, such as emotional or cognitive experience, it is plausible that a quality-space model could be developed for those domains if we only knew the relevant facts. If we wish to map the structure of experience, it seems that a significant part of the project will consist in constructing quality-space models for different kinds of experiences.[2]

This paper argues that the standard framework for modeling mental qualities is structurally inadequate, and develops a new framework that is more powerful and flexible. The core issue is that the standard framework for modeling mental qualities cannot capture *precision structure*. For example, when you see an object in foveal vision you might see it as crimson, but when you see it in peripheral vision you might see it merely as red (rather than any particular shade of red). In such a case, your

[1] See Klincewicz [2011] on temporal experience, Kostic [2014] on pain, Renero [2014] on auditory experience, and Young, Keller, & Rosenthal [2014] on olfactory experience.

[2] For discussion of the generality of the quality-space model framework, see Clark [1993, 2000] and Rosenthal [2010, 2015].

peripheral visual experience is less precise than your foveal visual experience. The basic problem for standard models of mental qualities is that imprecise qualities do not have particular values along the dimensions of standard models, and so do not correspond to any points in those models. Though this may at first appear to be a merely technical challenge, I will explain why solving the problem has both theoretical ramifications for our understanding of the structure of experience and methodological ramifications for the empirical investigation of experience.

The core idea behind my new framework is to model mental qualities using regions, rather than points, in multidimensional spaces. As a result, the framework I develop provides a unified way of representing the similarity, magnitude, and precision structure of mental qualities. Along the way, I will also explain how different formal constraints on models within the framework correspond to different classes of hypotheses about the structure and space of mental qualities, how the framework differentiates two dimensions of phenomenal similarity (which I call qualitative similarity and precision similarity), how empirical methods can be used to construct models for particular domains of mental qualities, and how the framework provides a way of modeling the intransitivity of perceptual discrimination. An upshot is that the structure of the mental qualities of conscious experiences is fundamentally different from the structure of the physical qualities of external objects.

In **§1**, I explain why standard models of mental qualities cannot capture precision structure. In **§2**, I develop my new framework for modeling mental qualities, which I call the *regional framework*. In **§3**, I discuss applications and extensions of the regional framework.

§ 1 | THE STANDARD FRAMEWORK

Let me begin by providing a brief overview of quality-space models. After that, I will explain what imprecise qualities are and why they pose a problem for standard quality-space models.

QUALITY-SPACE MODELS

A *quality-space model* is a model of a set of qualities.[3] The standard approach to modeling mental qualities is to represent qualities via points in multidimensional spaces. Under this approach, a quality-space model aims to represent every quality of a given domain so that there is one-to-one correspondence between qualities of the domain and points in the model and so that qualities that are more similar are represented by points that are less distant in the model.[4] As an example, a quality-space model for colors represents particular colors via points in a three-dimensional space such that points that are closer in the space represent colors that are more similar to each other.

Quality-space models can be developed for either the mental qualities of conscious experiences (such as phenomenal red) or the physical qualities of external objects (such as red).[5] The focus of this paper is specifically on mental qualities. As I will explain later, precision is not a property of physical qualities, so the framework I develop is inapplicable to physical qualities. As a terminological matter, whenever

---

[3] There is often ambiguity between talking about the formal representation of a domain of qualities and the domain of qualities itself: for example, consider Clark [2000]'s characterization of a quality-space as an "'ordering of the qualities presented by a sensory modality in which relative similarities among those qualities are represented by their relative distances." To disambiguate, I will always use the term 'quality-space' to mean the domain of qualities and the term 'quality-space model' to mean the formal representation of those qualities.

[4] Quality-space models also aim to capture magnitude relations between qualities so that higher magnitudes are represented by higher values along the dimensions of the model (e.g., if color quality A is brighter than color quality B, then A has a higher value along the dimension representing brightness than B). To simplify the discussion, I will focus only on similarity relations (rather than magnitude relations) for the rest of the paper, but the points I make will largely applicable to magnitudes as well.

[5] See Byrne [2011] for a more general discussion of the distinction between different kinds of qualities. Some philosophers, such as Rosenthal [2010], have argued that mental qualities can be instantiated even in the absence of consciousness. This paper is neutral on this issue and focuses only on mental qualities of conscious experiences.

I use the term 'quality' without qualification I will mean mental qualities. And for brevity, I will use terms such as 'hue' rather than 'phenomenal hue' to designate the dimensions of mental qualities. Though I distinguish mental qualities from physical qualities, my discussion is neutral on the metaphysical nature of mental qualities. In particular, none of the discussion in this paper is meant to suggest dualism about the mental and the physical or to imply that mental qualities cannot be understood as representations of physical qualities.

The target of this paper is the standard *framework* for modeling mental qualities, rather than any particular *model* within that framework. A model is a formal representation of a particular domain of qualities, whereas a framework is a general schema for developing models. In order to construct a model for any particular domain of experience (e.g., olfactory qualities), we must empirically investigate the relevant domain. But in order to develop an adequate framework in the first place, we must specify what kind of formal structure is required for modeling any domain of mental qualities, regardless of how any particular model is structured. I will soon argue that in order to capture the precision structure of mental qualities, we must make fundamental challenges to the whole quality-space model framework (rather than just revisions to particular models). Towards the end of the paper, I will briefly discuss how empirical methods can be used to construct particular models within the new framework I develop.

There is a diverse body of literature in both cognitive science and philosophy pertaining to the modeling of mental qualities. The relevant cognitive science literature has predominately focused on issues about the psychophysical relations between physical stimuli and mental qualities and the challenges in measuring mental qualities.[6] The relevant philosophical literature has predominately focused on questions

[6] For a classic text in psychophysics, see Fechner [1860]. For more recent overviews of psychophysics, see Murray [1993] and Gescheider [1997]. For an overview of the application of measurement theory to psychological models, see Luce & Krumhansl [1988]. For discussion of models of color qualities in particular, see Logvinenko [2015]. For some recent formal models of consciousness (targeting different explananda than my model), see Tononi [2007] and

about the relationship between mental qualities and physical qualities and the individuation of mental qualities.[7] However, in both disciplines, research that directly addresses the modeling of mental qualities tends to focus only on similarity structure and magnitude structure, leaving out precision.

This lacuna may be partly due to the common assumption that models of mental qualities are isomorphic to models of physical qualities. The standard methodology for constructing a model of mental qualities is to first use data concerning perceptual discrimination (or similarity or ordering) judgments to construct a model of a physical quality-space, and to then extrapolate from that model to a model of mental qualities. The justification is that mental qualities can be individuated by their perceptual roles: in particular, it seems that subjects make perceptual discriminations between physical stimuli only on the basis of being in mental states with different mental qualities. Since physical qualities (as opposed to physical stimuli)[8] are also individuated by subjects' perceptual discriminatory capacities, this suggests that the structure of mental quality-spaces is isomorphic to the structure of physical quality-spaces.[9] I will soon argue that this isomorphism thesis is false: models of mental qualities require more structure than models of physical qualities.

Prentner [2019]. For an approach to geometrically modeling concepts that shares some (though not all) formal features with my framework, see Gärdenfors [2014].

[7] For some classic and contemporary philosophical texts addressing these questions, see Goodman [1954] and Clark [2000], Rosenthal [2000, 2015]. For discussion of the role of phenomenology in psychophysical theorizing, see Horst [2005].

[8] Following convention, I distinguish between physical stimuli (e.g., wavelengths of light) and physical qualities (e.g., colors, which are multiply realized by different physical stimuli). Only the latter are held to be isomorphic to the models of mental qualities. See Clark [2000] for more discussion of the relationship between physical stimuli and physical qualities.

[9] See Sellars [1963], Clark [2000], and Rosenthal [2016] for discussion of this extrapolation step. As examples, Sellars [1963] talks of an "isomorphism of acts of sense and material things," and Rosenthal [2016] talks of extrapolating "from the quality space of perceptual discriminations to an isomorphic quality space of the mental qualities that enable those discriminations."

For the rest of the paper, I will call the approach to modeling mental qualities outlined above the *standard framework*, and I will call any particular model within that framework a *standard model*. More specifically, we can think of standard models as formally specifiable via a set of points (representing particular qualities), a set of dimensions (representing the respects in which qualities can vary from each other and be similar to each other), and a similarity metric (representing degrees of similarity between qualities). In what follows, I will explain why the standard framework is structurally inadequate.

PRECISION

Consider your color experience in foveal vision versus in peripheral vision. In foveal vision, you may see an object as a specific shade of red, such as crimson. But in peripheral vision, you may no longer see it as any specific shade of red, but instead just as red. It is not that you merely see the object as a different shade of red across the two cases. Instead, whereas your foveal color experience is characterized by a specific value along the hue dimension, your peripheral color experience does not seem to correspond any specific value along that dimension. This difference in phenomenal character is what I am calling *precision*.

Though I will focus on precision with respect to color experience, there are other examples that may also be used to illustrate the phenomenon. For example, consider the contrast between your color experience of an object that is far away versus nearby, your spatial visual experience with vision correction lenses versus without vision correction, or your tactile experience touching a texture with your fingertips versus your tactile experience touching a texture with your back. In each case, the former experience is more precise than the latter.[10]

Precision is different in kind from phenomenal properties such as hue, loudness, or painfulness. These familiar phenomenal properties correspond to particular

---

[10] I take the term 'precision' from Block [2015]. A number of philosophers have used the term 'determinacy' for what I call 'precision', but I later explain why the term 'precision' is better. For further discussion and examples of precision, see Block [2015].

dimensions of quality-space models. By contrast, precision cannot be captured in the same way (at least without making substantive theoretical assumptions). To see this, consider how a visual experience could plausibly be precise with respect to color and imprecise with respect to shape or how qualities even across different modalities can be similar with respect to precision. Instead, precision is a structural feature of experience, in the same way that similarity and magnitude are structural features of experience. This hints at why modeling precision requires modifying the entire quality-space model framework, rather than just patching up particular models.

Questions about the nature of precision depend on more fundamental issues in the philosophy of perception. For representationalists, it is natural to think that precision is a matter of the specificity of representational content. For naïve realists, it is natural to think that precision is a matter of being perceptually acquainted with more determinate phenomenal properties. For qualia theorists, it is natural to think that precision is a structural property of the phenomenal character of one's experience akin to the resolution of an image. Since the quality-space model framework is largely theory neutral, the issues discussed in this paper will be relevant to theorists across the board. At times, I will talk about mental qualities representing physical qualities, but my discussion can straightforwardly be reframed in terms of other relations (such as perceptual acquaintance or causal correspondence).[11]

There is a mix of literature across both cognitive science and philosophy pertaining to imprecise qualities. In cognitive science, this includes research on perceptual discrimination capacities (such as spatial resolution and tactile discrimination), on how attention affects perceptual discrimination, and on the neurophysiological properties underlying these differences.[12] However, these discussions tend to

---

[11] For a general overview of theories of perception, see Crane & French [2017]. For argument against representationalism about precision, see Block [2015]. For argument against naïve realism about precision, see Cutter [2019].

[12] See Strasburger et al [2011] for an overview of the psychology of peripheral vision, Intriligator & Cavanagh [2001] and Block [2012] for discussion of the relationship between visual resolution and visual attention, Anton-Erxleben & Carrasco [2013] for discussion of attention, spatial resolution and the role of receptor cell density for perceptual discrimination, Bruns et

focus on perceptual capacities and their functional roles, rather than the structure of the mental qualities associated with those perceptual capacities. In philosophy, there has been recent work examining imprecise qualities in connection with philosophical theories of perception, generic phenomenology, and the representational contents of experience.[13] However, these discussions have not directly addressed how precision structure relates to similarity structure and how to integrate imprecise qualities into the quality-space model framework.

Before explaining why precision poses a problem for the standard framework, I need to address a common confusion about precision. The confusion is that precision is a matter of *determinacy*, the relation between determinables and determinates. A determinate is a way for a determinable to be instantiated, and determinables may be thought of as disjunctions of determinates. To see why precision and determinacy are independent, consider the maximally determinate phenomenal property characterizing the particular peripheral color experience you are currently undergoing. That property is maximally determinate since there is only one way for that property to be instantiated, but it is also imprecise since it does not represent any specific shade of color. Conversely, consider the determinable phenomenal property that has as determinates each of a series of precise color phenomenal properties characterizing your foveal visual experiences when looking at a series of red color chips in optimal conditions. That property is determinable since there are multiple ways for that property to be instantiated but each of its determinates is precise since

---

al [2014] for discussion of tactile spatial resolution, Denison [2017] for discussion of precision in relation to perceptual uncertainty, and Gescheider [1997] for an overview of the role of perceptual discrimination thresholds in psychophysics.

[13] See Block [2015] and Cutter [2019] for discussions of precision in connection with the philosophy of perception, Fink [2015] and Fazekas & Overgaard [2018] for discussion of precision in relation to generic phenomenology and empirical methods for investigating consciousness, Stazicker [2011] and Morrison [2016] for discussion of the representational contents of imprecise experiences, and Hellie [2005] for an explanation of the intransitivity of perceptual indiscriminability that appeals to imprecise qualities.

they all represent specific shades of color. Since there are both maximally determinate imprecise phenomenal properties and determinable phenomenal properties with only precise determinates, precision and determinacy are independent.

The conflation between precision and determinacy is partly due to widespread ambiguities in natural language. For example, 'phenomenal red' can mean either a determinable phenomenal property that has precise determinates (e.g., phenomenal crimson, phenomenal scarlet, etc.) or a determinate phenomenal property that is imprecise. The conflation may also be partly due to the fact that the precision of mental qualities is systematically related to the determinacy of the physical qualities represented by those mental qualities. In other words, the more precise a mental quality, the more determinate the physical property represented by that mental quality: for example, a precise red experience might represent the determinate crimson whereas an imprecise red experience might represent the determinable red.[14]

It is unfortunate that precision is confused with determinacy, for the confusion masks the importance of precision for understanding the structure of experience. Developing a model of determinable phenomenal properties would not be particularly interesting, since questions about which determinable properties there could be are somewhat analogous to questions about which disjunctive properties there could be. In contrast, we will later address a number of interesting and substantive questions about the space and structure of imprecise qualities, about the implications of precision for phenomenal similarity, and about how precision interacts with the empirical investigation of consciousness.

What about the somewhat radical view that imprecise qualities involve the instantiation of a determinable without the instantiation of any of its determinates?[15] Even if this view were correct about the metaphysics, it would still be important to conceptually distinguish precision from determinacy. Otherwise, not only would there be systematic terminological ambiguities of the kind discussed above, but we

---

[14] Note that representationalism does not undermine this point, since the representing of a determinable property need not itself be a determinable property.

[15] See Wilson [2013] for a general defense of this metaphysical view.

would not even have the vocabulary to resolve those ambiguities. For example, suppose one claims that an experience instantiates the determinable property phenomenal redness. On this view, that claim could mean either that the experience is precise and instantiates one amongst many determinate phenomenal properties or it could mean that the experience is imprecise and instantiates the determinable phenomenal redness (without instantiating any of its determinates). Consequently, even those who endorse this unorthodox view should still agree that we need concepts and vocabulary that differentiate precision and determinacy.[16]

There may also be a temptation to contend that all mental qualities are precise, and that imprecision is really a matter of limitations of cognitive access. However, such a view is both introspectively and neurophysiologically implausible. For example—focusing on the case of vision—it is introspectively plausible that peripheral visual experience is less sharp and crisp than foveal visual experience, the density of receptor cells in the foveal areas of the retina is much greater than in the periphery, and there is a much larger proportion of visual cortex devoted to processing visual information from the foveal area than from the periphery. Consequently, the hypothesis that precision is merely a matter of cognitive access is undermotivated by both first-person and third-person evidence.[17]

We are now in position to see why standard models cannot capture precision. The core problem is that standard models represent qualities using points in multidimensional spaces, but imprecise qualities do not correspond to particular points in those spaces. This is because points in standard models correspond to specific values along the dimensions of the model, but imprecise qualities do not have particular values along those dimensions. Furthermore, we cannot solve the problem merely by adding an extra dimension (representing degree of precision) to existing models. Not

---

[16] Imprecision is also sometimes thought to be a matter of vagueness. For brevity, I will simply note that vagueness is typically understood as a semantic phenomenon whereas imprecision is a feature of phenomenal properties (as opposed to phenomenal terms), that terms for imprecise qualities can be sharp, and that none of the theories of precision discussed above appeal to vagueness. For an in-depth discussion of vagueness, see Williamson [1994].

[17] See Strasburger et al [2011] for a recent review of the science of peripheral vision.

only would that framework still require assigning particular values along ordinary dimensions (such as hue) to imprecise qualities, but also (as we will see in the next section) such a framework would be unable to accommodate views that allow the precision of one phenomenal dimension (such as hue) to vary independently of the precision of other phenomenal dimensions (such as brightness). The upshot is that capturing precision requires more than just tweaking existing models; instead, it requires fundamental changes to the quality-space model framework.

§ 2 | THE REGIONAL FRAMEWORK

The *regional framework* models mental qualities using *regions*, or sets of points, rather than just individual points. In what follows, I will develop the formal structure of the regional framework, address some of the technical challenges that arise in developing the framework, and explain why the framework is theoretically fruitful and philosophically consequential.

THE BASIC FORMAL STRUCTURE

Let me begin by explaining the formal structure required to specify a model in the standard framework, so that we can see the contrast with regional models. Any standard model requires a way of representing individual qualities, a way of representing the dimensions of the quality-space, and a way of representing degrees of similarity between qualities. Because of this, we can think of standard models as comprised of three elements: a set S of points (representing individual qualities), a set N of dimensions (representing the respects in which qualities can vary), and a distance metric d (where higher distance values represent lower degrees of phenomenal similarity).[18] As a result, the formal structure of any standard model is specifiable via an ordered triple: $\langle S, N, d \rangle$.

[18] Some, such as Gert [2017], have been skeptical that phenomenal similarity is best represented by a metric function. For the purposes of this paper, I will take this for granted that a metric function is needed to model a domain of mental qualities. Note that this makes the task of developing the regional framework harder, since it is not obvious how to develop a

There are three main desiderata when constructing a model in the standard framework. First, there should be one-to-one correspondence between points in the model and qualities in the targeted domain. Second, points that are more distant in the model should represent qualities that are less phenomenally similar to each other. Third, points should have distance zero just in case the qualities represented by those points are phenomenally identical. If these constraints are satisfied, then the structural properties of the model mirror the structural properties of the quality-space. But since standard models cannot capture precision structure, none of these desiderata can be fully satisfied. The challenge in what follows is to show that the analogous desiderata can be satisfied using the regional framework.

The regional framework represents experiences using *regions*, or sets of points, rather than individual points. The size of a region corresponds to the degree of imprecision of the quality represented by that region. We saw above that standard models can be specified via triples $<S, N, d>$. The regional framework requires adding more structure: regional models are specified via quintuples $<S, N, R, \mu, M>$ where S is a set of points, N is a set of dimensions, R is a set of regions, μ is a measure on S, and M is a set of metric functions (including the distance metric d over the points of S). Over the course of this section, I will explain each of these elements in detail.

Before moving on, let me explain how I will talk about imprecise qualities. At times, I will talk about similarity between regions, though this should be taken as shorthand for talk of similarity between the qualities represented by those regions. To avoid the systematic ambiguities between precision and determinacy that arise with natural language, I will denote the precision of a phenomenal property using values from 0–1 in superscript, where higher values indicate higher degrees of preci-

---

metric over regions that corresponds to phenomenal similarity. For discussion of how metric relations can be extracted from data involving only ordinal judgments, see Beals et al [1968].

sion. For example, phenomenal blue$^{.5}$ denotes a (maximally determinate) phenomenal blue quality with degree of precision 0.5. And throughout the section, I will represent imprecise qualities using diagrams like the one below:[19]

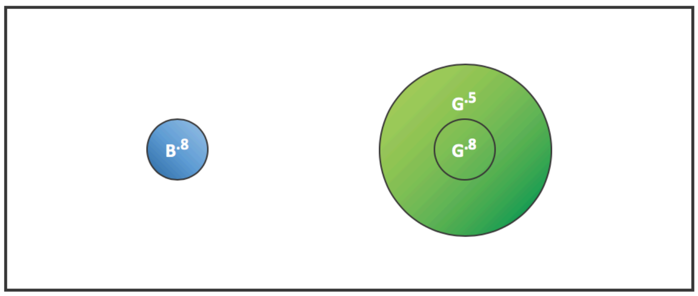

B$^{.8}$ represents a relatively precise phenomenal blue quality.

G$^{.8}$ represents a relatively precise phenomenal green quality.

G$^{.5}$ represents a relatively imprecise phenomenal green quality.

FIGURE 1: A pictorial representation of some regions in a regional model.

The rest of this section proceeds as follows: First, I discuss the structure of the spaces of regional models, focusing mostly on the set S of points, the set N of dimensions, and the point-distance metric d. Second, I discuss the structure of imprecise qualities, focusing mostly on the set R of regions. Third, I discuss similarity with respect to qualitative character versus similarity with respect to precision, focusing mostly on the measure μ and the set M of metric functions. Along the way, I will also explain how the formal challenges of the framework interact with philosophical questions about the precision structure of mental qualities.

THE STRUCTURE OF THE SPACE

Quality-space models represent qualities using multidimensional spaces. The structure of the space is determined by the set S of points, the set N of dimensions,

[19] The box represents a quality-space, the bounded shapes represent regions, and the regions are denoted using the linguistic convention described above. Note that the colors of the regions are just meant to make it easier to interpret which mental qualities the regions represent, and that the meanings of the dimensions in the figures are not important.

and the point-distance metric d. Since these are precisely the elements that characterize standard models, many of the theoretical issues that concern these elements are familiar ones that have been addressed in prior literature on modeling mental qualities. Since my aim is focus on new issues that arise with the regional framework, I will not discuss these elements in much detail. However, there is one question concerning these elements that must be addressed here.

The question is whether the spaces in regional models should be discrete or continuous. In a discrete space there is a finite number of points between any two distinct points, whereas in a continuous space there is a continuum of points between any two distinct points. A discrete regional model would enable us to simply "export" all the points from a standard model into a regional model, meaning we would retain the same set S of points as in the standard model.[20] In contrast, a continuous regional model would require a new set S that has infinitely many points, and a new set N that has dimensions with infinitely many values.[21]

Though discrete models are somewhat simpler to construct, continuous models are arguably better because of their flexibility. A continuous model can capture qualities at arbitrary levels of precision, including even qualities with greater precision than the maximally precise qualities for human experiences. Furthermore, it is mathematically simpler in continuous spaces to specify certain kinds of formal constraints on regions, which is an advantage that will be relevant in the next subsection. For these reasons, I will assume for the rest of the paper that the models under consideration are continuous (though most of the discussion will apply also to discrete models). Note that while continuous spaces have infinitely many points, they may still be bounded, in that all points lie within a fixed distance from each other.

---

[20] I assume that standard models must have a finite number of points, since that is needed to satisfy the desideratum of one-to-one correspondence between points and qualities (at least if we assume that there are finitely many mental qualities).

[21] There is also the option of dense but non-continuous space: for example, the rational numbers are dense but not continuous. I will set aside this option, since it is not clear there is any reason to favor a dense but non-continuous space over a continuous space.

For example, the interval of real numbers from 0 to 1 is continuous, but it is also bounded by the limit points 0 and 1.

In a continuous regional model, individual points are probably best thought of as idealizations: they are the maximally specific values along the dimensions of a quality-space, even if no qualities actually correspond to regions comprised of a single point. Some might worry that this would mean that regional models have more structure than is necessary. However, scientific models often idealize, especially when doing so leads to simpler formalisms. For example, suppose that all physical objects are composed of elementary particles, that there is a finite set of elementary particles, and that all elementary particles have discrete mass values. Then there are some mass values that no physical objects could have, since no combination of elementary particles would add up to that mass value. Nevertheless, it may still be useful to represent mass using real numbers, which have continuous structures. By the same lights, it may be useful to model mental qualities using continuous spaces even if it turns out that points in the spaces are idealizations.

PERMISSIBLE REGIONS

The third element of a regional model is the set R, which has as its elements *regions*, or subsets of the S of points. Every point in a standard model is guaranteed to correspond to a region in a regional model (since there are singleton regions), but some regions in regional models do not correspond to any points in standard models (since there are non-singleton regions). Because of this, a regional model can represent all the qualities that a standard model can represent, but not vice versa.

Why is there a need for R at all? Some might think that we could simply take imprecise qualities to be represented by the subsets of S, avoiding the need to posit a whole new set R. However, we need R in order to distinguish different theories of the structure and space of imprecise qualities. As we will see, it is probably not the case that every subset of S corresponds to a possible mental quality. Consequently, we need a way of distinguishing regions, which can be any subset of S, from *permissible regions*, which are the subsets of S that are members of R and that are to be interpreted as representing possible mental qualities. To put it another way, different specifications of R correspond to different theories about the structure of imprecise qualities. Since there are as many ways of specifying R as there are sets of subsets of S, it will be useful to focus on different formal constraints on R that permit different kinds of permissible regions.

These formal constraints on R fall into two kinds. The first kind concerns *sizes*: what are the minimal and maximal sizes of permissible regions? Questions about size constraints correspond to questions about degrees of precision. For example, we might wonder whether there could be maximally imprecise qualities whose regions cover entire quality-spaces. Could there be an experience that merely represents something as colored (with no further specificity about particular colors)? Conversely, we might wonder whether there could be super-precise qualities whose regions are arbitrarily small. Could there be color qualities more precise than those characterizing maximally precise human color experiences? There is a question here of how to determine the sizes of regions in spaces with infinitely many points. For now, let us set that question aside—we will address it in the next subsection.

The second kind of constraint on R concerns *shapes*: what kinds of shapes can permissible regions take? Questions about shape constraints correspond to questions about the structure of precision. Since these questions are somewhat more difficult to formulate, it will take some space to explain what different shape constraints might look like and how they correspond to different theoretical questions about the structure of mental qualities. In what follows, I will discuss three different formal constraints on the shapes of permissible regions: CONNECTEDNESS, CONVEXITY, and UNIFORMITY. Each constraint corresponds to a different class of hypotheses about the space and structure of imprecise qualities, ranging from most permissive to most restrictive. The diagram below displays four sample regions that are permitted by different formal constraints:

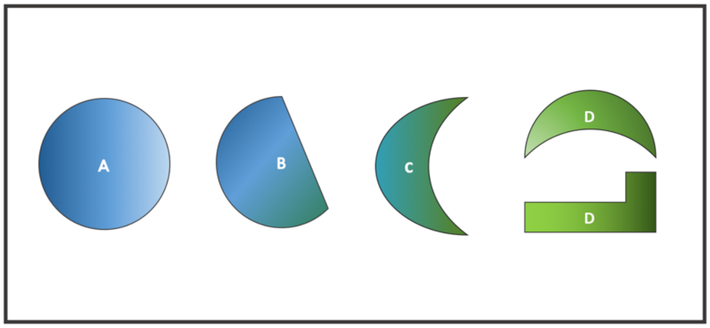

Each constraint below is more restrictive than the constraint above on the set permissible regions.

CONNECTEDNESS: A, B, C
CONVEXITY: A, B
UNIFORMITY: A

FIGURE 2: Formal constraints on permissible regions.

Let us start with the most permissive of the three constraints:

CONNECTEDNESS: All permissible regions are connected.[22]

A region is *connected* just in case there are no discontinuities in the region, or just in case any two points in the region can be connected via a continuous path of adjacent points that all lie within the region. CONNECTEDNESS excludes region D from FIGURE 2 but permits regions A, B, and C. By excluding disconnected regions, we rule out mental qualities that represent arbitrary collections of physical qualities, such as a quality that represents something as either crimson or aquamarine (but nothing else). However, CONNECTEDNESS still permits regions such as region C in FIGURE 2. For those who think that CONNECTEDNESS is too permissive, the following constraint may be more attractive:

CONVEXITY: All permissible regions are convex.[23]

---

[22] Formally: if $\mathrm{A} \in R$, then A is not the union of two disjoint open sets of S. A set $\mathrm{A} \subseteq S$ is open just in case for all $a \in \mathrm{A}$, there exists $\epsilon > 0$ such that the ball $\mathrm{B}(s, \epsilon) = \{x \in S \mid d(x, s) > \epsilon\}$ satisfies $\mathrm{B}(s, \epsilon) \subseteq S$.

[23] Formally: if $\mathrm{A} \in R$, then for any points $a, c \in S$ such that $d(a, c) > 0$, A contains all points $b \in S$ such that $d(a, b) + d(b, c) = d(a, c)$.

A region is *convex* just in case for every pair of points within the region, every point on the straight line-segment that joins the pair of points is also within the region. In other words, CONVEXITY entails that any region must contain all points within the straight lines connecting its boundaries. CONVEXITY excludes regions C and D from FIGURE 2 but permits regions A and B. However, some might think that even CONVEXITY is too permissive about the structure of imprecise qualities, in which case the following constraint may be attractive:

UNIFORMITY: All permissible regions are balls.[24]

A region is a *ball* just in case it is the set of points that are less than a certain distance from a center. For example, in three-dimensional spaces, balls will typically be perfectly spherical.[25] The intuitive idea behind UNIFORMITY is that permissible regions are always anchored on a center, and then expand outwardly by the same distance in all directions. UNIFORMITY excludes regions B, C, and D from FIGURE 2, permitting only region A. If we accept UNIFORMITY, it becomes much easier to grasp the phenomenal character of any arbitrary quality. However, UNIFORMITY excludes the possibility of qualities that differ in their degree of imprecision across different dimensions. For example, UNIFORMITY excludes the possibility of a color experience that is precise with respect to hue but imprecise with respect to brightness. Because of this, some might worry that UNIFORMITY is too restrictive.[26]

The discussion above merely scratches the surface. However, my present aim is not to evaluate which of these formal constraints is most plausible. Instead, it is to show how the regional framework enables us to formulate interesting hypotheses about the space and structure of mental qualities by providing us with formal tools

[24] Formally: if $\mathrm{A} \in R$, then there is some $s \in S$ and $\epsilon > 0$ where $\mathrm{A} = \{x \in S \mid d(x, s) > \epsilon\}$.

[25] UNIFORMITY does not entail that all regions are perfectly shaped. In particular, regions that lie near the boundary points of a space may have imperfect shapes since there may be no points in the space past a certain distance.

[26] A possible response is to require regions to be ellipsoids rather than balls.

that enable more rigorous and systematic theorizing. Notice how difficult it would be to demarcate these different theories using natural language, and how it is hard to even formulate some of these questions in the first place without the appropriate theoretical framework.

QUALITATIVE SIMILARITY

A standard model has a single metric d, which takes as input two points and outputs a distance value. The regional framework needs a metric that outputs distance values between regions. As we will see, identifying the right metric takes some work. We need a metric that can take as input regions of arbitrary size and shape, and we need the outputs of the metric to systematically correspond to degrees of phenomenal similarity, where lower distance values represent higher degrees of phenomenal similarity and where distance zero represents phenomenal identity.[27]

Before moving on to candidates for metrics, let me first introduce the *measure*, which will be used to construct some of the metrics that we will discuss. The measure μ is the fourth element of a regional model, and is a function that takes as input a subset of S and outputs a size value.[28] Since the size of a region represents its degree of imprecision, μ can be thought of as telling us how imprecise the quality represented by a region is. In discrete models, it is natural to simply take the size of a region to be the number of points in the region. However, such a measure does not work well for continuous models, since regions in continuous models will typically have infinitely many points. Instead, we need the standard mathematical measure: the Lebesgue measure, which is a generalization of the notions of length, area, vol-

---

[27] Note that these metrics will take in different kinds of elements, since the elements of S are points and the elements of R are sets of points (including even the singleton sets).

[28] Some might wonder why the measure is on the set S of points rather than the set R of regions. Since a measure takes as input a subset of a set, a measure on R would output the sizes of sets of regions, rather than the sizes of regions. Consequently, determining the size of a region in R requires measuring the corresponding subset of S.

ume, and so forth. The mathematical details of the Lebesgue measure are not particularly philosophically relevant. What is important is that the Lebesgue measure has the properties we would intuitively want a measure to have: it works in continuous spaces of arbitrary dimensionality and its size outputs correspond to the degrees of imprecision of the mental qualities represented by those regions.[29] With μ on the table, we are in a position to consider whether there is a metric function for regional models that can provide a plausible representation of phenomenal similarity.

A first pass is to default to the standard way of determining distance between sets: the least distance function, which takes the distance between two regions to be the lowest distance value between any two points in the two regions.[30] However, consider two regions A and B that are distinct but overlap (i.e., two regions that contain some but not all of the same points). Since these regions are distinct, they represent distinct qualities, such as phenomenal red and phenomenal reddish-orange. But since A and B overlap, there is a point in A that has distance zero to a point in B. As a consequence, the least distance function has the result that the distance from A to B is zero. Since A and B represent distinct qualities, and since distance zero represents phenomenal identity, we have the wrong result.

A more promising option is the average distance function, *avg*, which takes the distance between regions A and B to be the average distance from points in A to points in B. More specifically, determining the average distance between A and B requires taking a point in A, determining the average distance between that point and each point in B, repeating the procedure for every point in A, and then averaging

[29] Note that the Lebesgue measure will always output finite values in bounded spaces (even if the space is continuous). Some other attractive properties of the Lebesgue measure include additivity ($\mu(\text{A} \cup \text{B}) = \mu(\text{A}) + \mu(\text{B})$), non-negativity ($\mu(\text{A}) >= 0$ for any region A), and translation invariance. For a more detailed overview of the Lebesgue measure, see Tao [2011].

[30] Formally (letting *inf*(A) mean the infimum of set A): the least distance from A to B $=$ *inf* $\{d(a, b) \mid a \in \text{A},\ b \in \text{B}\}$.

the averaged distance values. Since the average distance function is sensitive to all the points in A and B, it is an improvement over the least distance function.[31]

Nevertheless, it turns out that the average distance function is also inadequate. To see why, consider the average distance from any region A to itself. So long as A contains more than one point, there will be some pair of points a and b in A where $d(a, b) > 0$. Since the average distance between regions is just an averaging of distances between their points, and since some of the distances between points are non-zero, the average distance from a region to itself is also non-zero. However, non-zero distance values are supposed to represent phenomenally distinct qualities. In other words, we get the absurd result that all mental qualities that are not maximally precise are not maximally similar to themselves. The core problem is that the *avg* function does not differentiate between points that are shared between regions versus points that belong to only one region. When regions are disjoint, *avg* delivers intuitive results; when regions overlap, *avg* runs into problems. We need a new metric that is sensitive to this difference.

Now we can turn to my proposal, which I will call the 'qualitative similarity metric', or *qual*. The metric is motivated by a simple observation: any case involving overlapping regions can be treated as a pair of cases involving disjoint regions. To see how this works, consider the diagram below:[32]

---

[31] The average distance function is straightforward to characterize in finite spaces but is somewhat more complex for infinite spaces. For finite spaces, letting $sum_{A} f(a)$ mean summation of $f(a)$ for all $a \in A$, the average distance function can be defined as follows: $\text{avg}(A, B) = 1/(\mu(A)\mu(B)\ *\ (sum_{A}(sum_{B}(d(a,\ b))))$. However, since summations over divergent series are undefined, in infinite spaces summation must be replaced with Lebesgue integration. Then, letting $int_{A}(f(x)^{*}\mathrm{d}x)$ mean the Lebesgue integral of f on the elements of set A, the average distance function for infinite spaces can be defined as follows: $\text{avg}(A, B) = 1/(\mu(A)\mu(B)\ *\ (int_{A}(int_{B}(d(a,\ b)^{*}\mathrm{d}\mu(b))^{*}\mathrm{d}\mu(a))$. For a more comprehensive discussion of these average distance functions, see Fujita [2013].

[32] To interpret the diagram, observe that A overlaps with C and B, that B overlaps with A and D, and that C and D are disjoint.

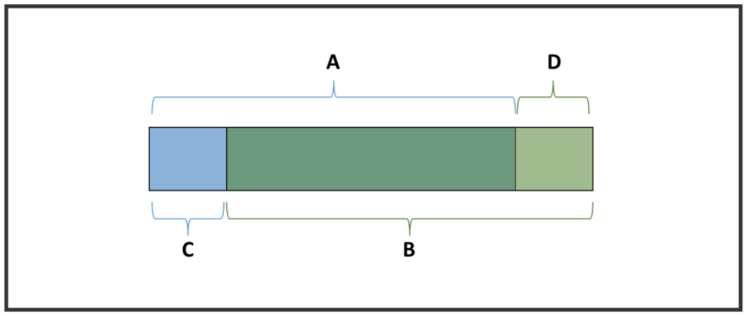

FIGURE 3: Four regions.

Suppose we wish to determine the similarity between A and B. The *qual* metric works by taking the average distance from A to the subregion of B that does not overlap with A, the average distance from B to the subregion of A that does not overlap with B, and then taking the weighted average of the two average distances (with the weighting in proportion to the relative sizes of A and B). If we apply this procedure to the diagram above, we would first find *avg*(A, D), then find *avg*(B, C), and then determine the weighted average of the two results. The *qual* metric is expressed formally in the equation below:[33]

$$\mathrm{qual}(\mathrm{A},\mathrm{B}) = \mathrm{avg}(\mathrm{A},\ \mathrm{B}\backslash\mathrm{A})\frac{\mu(\mathrm{B}\backslash\mathrm{A})}{\mu(\mathrm{A}\cup\mathrm{B})} + \mathrm{avg}(\mathrm{B},\ \mathrm{A}\backslash\mathrm{B})\frac{\mu(\mathrm{A}\backslash\mathrm{B})}{\mu(\mathrm{A}\cup\mathrm{B})}$$

And the metric is illustrated pictorially in the diagram below:

[33] Some might wonder why the numerators of *qual* are μ(B\A) and μ(A\B) rather than simply μ(B) and μ(A). The reason is that this enables *qual* to treat overlapping parts of regions differently from disjoint parts of regions, avoiding the issues we encountered with *avg*. If the numerators were μ(B) and μ(A), then the function would not entail that *qual*(A, A) = 0. For more detailed discussion of the formal properties of *qual*, see Fujita [2013].

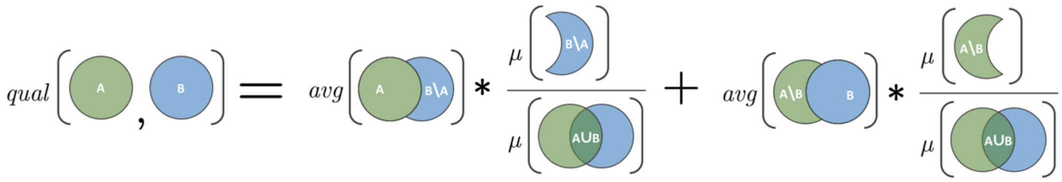

FIGURE 4: The qualitative similarity metric.

Though *qual* may appear complex at first glance, it is simple and intuitive after getting a feel for how it works. To better understand why *qual* produces the intuitively correct results, let us consider the four basic relations that two regions A and B can stand in to each other—namely, identity, disjointness, containment, and overlap—and how the metric works in each of these cases. These relationships between regions are illustrated in the diagram below:

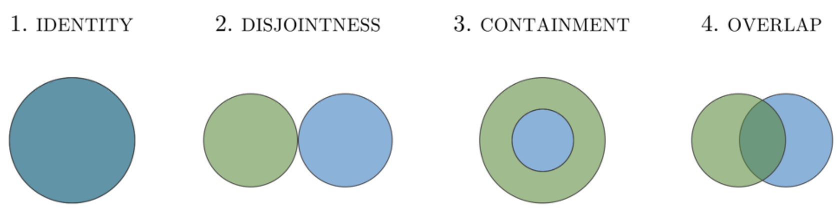

FIGURE 5: The four basic relations between two regions.

First, A and B are identical if and only if $qual(\text{A}, \text{B}) = 0$. This means that two qualities are represented as qualitatively identical just in case those qualities are represented by the same region. Second, whenever A and B are disjoint, $qual(\text{A}, \text{B}) = avg(\text{A}, \text{B})$. In other words, the qualitative similarity metric collapses to the average distance function if A and B do not overlap—and as observed above, the average distance function produces intuitively correct verdicts in such cases. Third, whenever A contains B, $qual(\text{A}, \text{B}) > 0$. This means that qualities that differ in precision are never represented as qualitatively identical. Moreover—presuming A contains B—

increasing the difference in size between A and B also increases their distance, meaning that the metric predicts qualitative similarity to decrease as the difference in precision increases. Fourth, whenever A and B overlap, *qual*(A, B) > 0. The more A and B overlap, the lower the distance between them. More generally, the metric entails that regions with points that are more distant are themselves more distant.

These results mean that *qual* satisfies the desiderata on a metric outlined earlier.[34] This is a significant finding: neither the least distance metric nor the average distance metric produced results that plausibly correspond to degrees of similarity, and other standard candidates for metrics on regions fare poorly as well.[35]

PRECISION SIMILARITY

Nevertheless, it turns out that there remains an aspect of phenomenal similarity that *qual* does not capture. To see this, consider again the diagram from the beginning of this section:

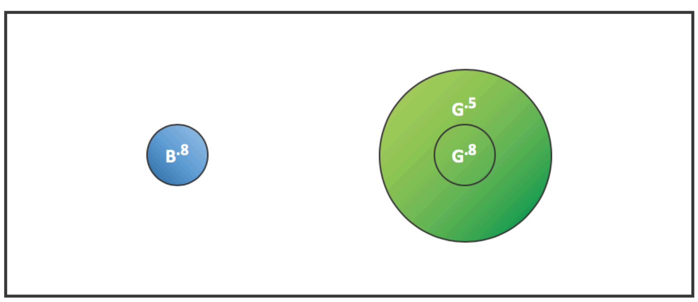

B$^{.8}$ represents a relatively precise phenomenal blue quality.

G$^{.8}$ represents a relatively precise phenomenal green quality.

G$^{.5}$ represents a relatively imprecise phenomenal green quality.

FIGURE 6: *qual* does not represent precision similarity.

---

[34] A further virtue of *qual* is that it works in origin-free spaces (i.e., spaces whose dimensions do not have zero-values). This point is important, for some dimensions of quality-spaces do not have zero-values (e.g., hue), and any framework for modeling qualities must be able to accommodate such structures.

[35] For example, the Hausdorff metric is another standard metric, but its output depends only on certain local maxima and minima points, rather than all points in the relevant regions. See Rockellar & Wets [2005] for discussion of this metric.

In this situation, *qual* predicts that the precise phenomenal blue quality is as similar to the precise phenomenal green quality as it is to the imprecise phenomenal green quality. More specifically, *qual*($B^{.8}$, $G^{.8}$) = *qual*($B^{.8}$, $G^{.5}$), since for disjoint regions *qual* works the same way as *avg*. But $B^{.8}$ is more similar to $G^{.8}$ than it is to $G^{.5}$. In particular, $B^{.8}$ and $G^{.8}$ are similar with respect to precision, whereas $B^{.8}$ and $G^{.5}$ are not. The result reveals that *qual* is not directly sensitive to the relative sizes of regions.

How should we proceed? In my view, this result is a feature rather than a flaw, for the result reveals two different dimensions of phenomenal similarity. On the one hand, two mental qualities might be similar with respect to *qualitative character*, or the aspects of phenomenal character that are characterized by phenomenal properties such as hue, loudness, and painfulness and that correspond to the dimensions of quality-space models. On the other hand, two mental qualities might be similar with respect to *precision*, which is a structural feature of experience that does not correspond to any dimensions. It is easy to get an intuitive grip on the difference between these kinds of phenomenal similarity. For example, consider the difference in kind of similarity when comparing a precise phenomenal red quality to a precise phenomenal orange quality versus when comparing a precise phenomenal red quality to an imprecise phenomenal red quality. Our analysis of *qual* indicates that it is the right metric for qualitative similarity. But it must be supplemented with another metric, which I will call *prec*, that measures precision similarity.

Since the degree of imprecision of a quality is represented by the size of a region, defining *prec* is straightforward. In particular, we can develop *prec* by appeal to the difference in size between regions. Since the order of comparison should not matter and since it does not make sense to talk about negative degrees of similarity, it is the absolute value of the difference in size that matters. Accordingly, we can express *prec* as follows:[36]

$$prec(\mathrm{A}, \mathrm{B}) = |\mu(\mathrm{A}) - \mu(\mathrm{B})|$$

With *prec*, we have a solution to the problem outlined above. The situation involved a precise phenomenal blue experience $\mathrm{B}^{.8}$, an imprecise phenomenal green experience $\mathrm{G}^{.5}$, and a precise phenomenal green experience $\mathrm{G}^{.8}$. The problem was that *qual* does not capture the precision similarity between the precise phenomenal blue experience and the precise phenomenal green experience. But *prec* is designed to capture this second dimension of similarity (without encroaching on the qualitative similarity captured by *qual*). The result is expressed formally below:[37]

$$qual(\mathrm{B}^{.8}, \mathrm{G}^{.8}) = qual(\mathrm{B}^{.8}, \mathrm{G}^{.5})$$
$$prec(\mathrm{B}^{.8}, \mathrm{G}^{.8}) > prec(\mathrm{B}^{.8}, \mathrm{G}^{.5})$$

Some might object that we still have not identified the overall phenomenal similarity metric that represents phenomenal similarity simpliciter. I am inclined to think that there is no objective fact of the matter about how to compare the kinds of phenomenal similarity tracked by *qual* and *prec*. However, the regional framework itself is designed to accommodate a wide range of views. For those inclined to think

[36] Technically, this makes *prec* a pseudometric, rather than a metric, since any metric d must satisfy the condition that if $d(a, b) = 0$ then $a = b$. For ease of explication I will continue referring to *prec* as a metric.

[37] The *prec* metric commits to a view where precision has at least interval structure, meaning we can make sense of ordering relations on differences in precision. A further question is whether precision has ratio structure, where we can make sense of multiplicative ratios on precision values. I do not have a view on this, though it is worth noting that the regional framework has the resources to model the ratio structure of precision (if such a view is correct), since the outputs of μ (e.g., real numbers) presumably have ratio structure. For discussion of these different measurement scales, see Stevens [1946].

there are objective facts about overall similarity between mental qualities, we could always search for a more general metric that captures both qualitative and precision similarity.[38] However, whether or not there are objective facts about overall phenomenal similarity, it is plausible that *qual* and *prec* track two natural phenomenal kinds.

At this point, some might wonder whether it even makes sense to identify these similarity metrics before empirically investigating the similarity relations between mental qualities. This is a natural worry, but it confuses the empirical project of mapping mental qualities onto formal models with the theoretical project of specifying the formal framework in the first place. In order to determine how particular mental qualities map to particular regions within a model, we must empirically investigate the similarity relations amongst that domain of mental qualities and use the data to constrain the mapping. But to do that, we need to know how that empirical data is supposed to constrain the formal structure of the model. To put it another way, using empirical methods to develop models presupposes that we already have similarity metrics in place, and identifying those similarity metrics requires invoking general theoretical considerations of the kind discussed above.

§ 3 | APPLICATIONS AND EXTENSIONS

In this last section, I will explain how empirical methods can be used to construct regional models for particular domains of mental qualities and how the framework can be extended to accommodate probabilistic contents and phenomenal vivacity. There is not enough space to discuss either of these issues in detail, but I will provide a general outline of how to apply and extend the regional framework.

EMPIRICAL METHODOLOGY

Let us assume that we start with a standard model that captures a set of maximally precise mental qualities. The initial step is to convert that standard model into a regional model representing those same qualities. This requires mapping points

---

[38] An obvious way of yielding an overall similarity metric is to simply add the outputs of *qual* and *prec*, perhaps with a weighting to scale their values relative to each other.

in the standard model to regions in the regional model such that there is one-to-one correspondence between points in the former and regions in the latter, such that the similarity and magnitude relations are preserved, and such that the boundaries of the space remain the same. The basic procedure for this conversion is relatively straightforward, though there is a puzzle about how to determine when two regions overlap. Let us set that issue aside for the moment—we will return to it soon.

After converting the standard model to a regional model, the challenge is to identify mental qualities that are not maximally precise and to map them onto regions, with the constraints that the distance values (outputted by *qual* and *prec* between regions) correspond to the degrees of qualitative and precision similarity (between the qualities represented by those regions). To put it another way, for any particular mental quality, we need a method for identifying its corresponding region's location, size, and shape. The easiest of these tasks is location, for the methodology used to determine locations of points in standard models generalizes to determining locations of regions in regional models. Because of this, I will focus below on the challenges concerning region sizes and region shapes.[39]

To determine the size of a region for a particular mental quality, we need to partition perceptual capacities into subclasses, where subclasses are individuated by their fineness of grain of perceptual discrimination abilities. For example, since color discrimination is more coarse-grained outside of the center of visual field, color perception might be divided into subclasses corresponding to different angles subtended from the center of the visual field. More precisely, these subclasses can be individuated by distances between just-noticeably different physical stimuli, where two physical stimuli that are just-noticeably different for a more fine-grained perceptual capacity subclass will be indiscriminable for a more coarse-grained perceptual capacity

[39] I will take for granted the modest assumption that more precise mental qualities are associated with finer-grained perceptual discrimination capacities. This assumption is crucial for justifying the relevance of perceptual discrimination data to the modeling mental qualities. Note that we need assume only an evidential (as opposed to constitutive) relationship between precision and fineness of grain of perceptual discrimination.

subclass. Then, given the hypothesis that precision is related to fineness of grain of perceptual discrimination, each set of just-noticeably different stimuli identified by each perceptual capacity subclass will correspond to a set of mental qualities at a different degree of precision. In other words, this method enables us to identify different classes of mental qualities corresponding to regions of different sizes.

The remaining challenge is to devise a method for identifying region shapes. This procedure is more complex, and requires us to invoke a principle connecting overlapping regions with perceptual roles: in particular, if two mental qualities are represented by regions that overlap, then the representational contents of those qualities overlap.[40] By overlap of representational contents between qualities A and B, I mean that some of the ways A represents the world to be are also ways that B represents the world to be. In light of this, it is natural to take the degree of overlap between the region representing A and the region representing B to correspond to the probability of a subject judging the physical quality represented by A to be the same as the physical quality represented by B.[41] In other words, the more overlap between the regions representing A and B, the more likely a subject is to judge that the objects represented by A and B have the same physical qualities. Once this principle is on the table, we not only have a method for identifying region shapes but also a way of modeling the intransitivity of perceptual discrimination.

Let me start with the intransitivity of perceptual discrimination. Intransitivity concerns cases when a subject perceptually judges that *a* is the same as *b*, and that *b* is the same as *c*, but that *a* is distinct from *c*. With the regional framework, intransitivity can be modeled as cases involving partial overlap between regions A and B and partial overlap between B and C (but not necessarily partial overlap between A and C). More specifically, the principle connecting overlapping regions to

---

[40] These claims concerning representational contents can be translated into claims that are consistent with a naïve realist framework. In particular, it is natural for naïve realists to take mental qualities represented by overlapping regions to be constituted by perceptual relations to determinable properties (of external objects) that share some determinates.

[41] These principles linking mental qualities to perceptual roles must be ceteris paribus clauses, since there may be noise in the process from stimulus to judgment.

perceptual roles entails that greater overlap between A and B or between B and C increases the chance of intransitivity and that greater overlap between A and C decreases the chance of intransitivity. Then the probability of an intransitivity judgment will be related to the overlap between regions in the following way (where ~ expresses proportionality and *overlap* is a function outputting degree of overlap):

$$p(\text{intransitivity: A, B, C}) \sim \tfrac{1}{2}(\mathit{overlap}(\text{A, B}) + \mathit{overlap}(\text{B, C})) - \mathit{overlap}(\text{A, C})$$

This provides a systematic way of using intransitivity data to refine regional models, as well as using developed regional models to predict the likelihood of intransitivity. In contrast, those operating within the standard framework often attempt to analyze away intransitivity data, since there is no natural way of modeling intransitivity cases using the standard framework. An important theoretical result is that even the mental qualities captured by standard models are best represented by regions, since intransitivity cases occur even for the maximally fine-grained perceptual capacities of normal human subjects.[42]

By similar lights, we have a method for identifying region shapes. The core idea is that the principle connecting overlapping regions and perceptual roles results in different predictions for cases involving overlapping versus disjoint regions. As a result, we can use asymmetries in perceptual discrimination data to triangulate the shape of a region. The most straightforward way of doing this is to examine perceptual judgments involving an imprecise quality and other precise qualities. This would require setting up experimental conditions such that subjects are able to make perceptual judgments about pairs of physical stimuli that they perceive using different perceptual capacity subclasses. For example, this might require subjects to make similarity judgments about color chips that they see via the center of their visual field versus color chips that they see via the periphery of their visual field. If a subject

---

[42] For a more detailed defense and discussion of this explanation of the intransitivity of perceptual discrimination, see Hellie [2005].

judges two physical stimuli to be the same, then that indicates that the region representing the mental quality associated with the fine-grained perceptual capacity subclass at least partially overlaps the region representing the mental quality associated with the coarse-grained perceptual capacity subclass. By collecting enough data to identify the degrees of overlap between the region representing the imprecise quality and the regions representing the precise qualities, this method enables us to approximate the shape of the region for any given mental quality.

FROM REGIONS TO FIELDS

Though the regional framework is a significant improvement over the standard framework, it is not the end of the story. I will briefly outline a promising way of extending the regional framework: namely, by moving from regions to fields.

Let me begin by stating two motivations for extending the regional framework. First, the regional framework cannot accommodate views that take mental qualities to have probabilistic contents.[43] For example, some think that a mental quality could represent an object as having a 50% chance of being red, a 25% chance of being orange, and a 25% chance of being yellow (where we may suppose that red, orange, and yellow each correspond to regions of the same size). Second, the regional framework cannot represent phenomenal vivacity.[44] For example, it is plausible that perceptual experiences can be more vivid than imaginative experiences, even holding fixed their qualitative and precision properties. The limitation of the regional framework is that regions in the regional framework are "flat," in that they do not assign

[43] See Morrison [2016] for argument for the view that experiences have probabilistic contents. See Denison [2017] for a counterargument against the probabilistic contents view.

[44] See Fazekas & Overgaard [2017] for discussion of the difference between precision and vivacity, and Kind [2017] for a more general discussion of vivacity. Kind concludes that the concept of vivacity is too unclear to be philosophically useful, where the main motivating consideration is that vivacity cannot be characterized in more basic phenomenal terms. However, that consideration may instead to be taken to suggest that vivacity is a basic structural feature of experience, just like precision.

different weightings to different points. Because of this, the regional framework does not have enough structure to model either probabilistic contents or vivacity.

However, both of these phenomena can be captured using a *field framework.* A *field* on a space of points assigns a value to every point in that space. In other words, a field may be thought of as a function from points to values. A region may be thought of as a special case of a field, where the region assigns to each point either 0 (if the point is outside the region) or 1 (if the point is inside the region). However, fields can assign a broader range of values (where a natural constraint is that values are real numbers between and including 0 and 1). The field framework has all of the structure of the regional framework, as well as additional structure due to assigning values to points. As a result, the field framework is another natural step in the evolution of the quality-space model framework: the representation of mental qualities progresses from points to regions to fields.

The field framework can model both probabilistic contents and phenomenal vivacity. The probabilistic contents of a quality would correspond to the distribution of values assigned to points by the field representing that quality. For example, a mental quality representing an object as 50% chance red, 25% chance orange, and 25% chance yellow would correspond to a field that assigns higher values in the red region of the quality-space than in the orange and yellow regions. The phenomenal vivacity of a quality would correspond to the integral of the field that quality, or the value we get by adding up all the values assigned to points by that field. More specifically, if quality A is more vivid than quality B, then the integral of the field representing A would be higher than the integral of the field representing B.[45]

---

[45] Some might notice that in the field framework, the measure μ represents vivacity rather than precision. This is because the measure takes in different information across the two frameworks. In the field framework, the measure takes into account not only the size of a region but also the values assigned to points, which are not directly relevant to the precision of a region. In contrast, in the regional framework, both region size and standard deviation are ways of measuring precision similarity. In my view, there is no objective fact of the matter about which of these is a better metric for measuring precision similarity: both region size and standard deviation track precision using different but related mathematical methods.

As with the regional framework, there is a challenge of identifying the right similarity metrics. In particular, the field framework would likely need three similarity metrics, representing qualitative similarity, precision similarity, and vivacity similarity. There is not enough room to adequately explain why the particular metrics I favor track the relevant kinds of similarity. Instead, I will simply state what I believe is the best way of developing each metric. First, the qualitative similarity metric would enrich *qual* by assigning weights to each distance value that figures into the calculation, with the weights corresponding to the average of the values of each pair of points. Second, the precision similarity metric would be determined by absolute differences in weighted standard deviations, where the standard deviation of a field would be determined by the dispersion of its points from its center (with the center being the point that minimizes average distance and the weighting determined by the values assigned to points by the field). Third, the vivacity similarity metric would be determined by absolute differences between the integrals of fields.

The field framework merits more discussion, but I hope to have said enough to illustrate the core idea. Moreover, I hope it is evident that the regional framework not only improves upon the standard framework, but also provides the basis for more sophisticated frameworks with richer structures.

CONCLUSION

To recap, a regional model can be specified via a quintuple $<S, N, R, \mu, M>$, where S is a set of points, N is a set of dimensions, R is a set of regions, μ is a measure, and M is a set of metric functions comprised of a point-distance function d, a qualitative similarity function *qual*, and a precision similarity function *prec*.

The regional framework is more powerful and flexible than the standard framework. Its additional structure enables us to formulate novel hypotheses about the space and structure of mental qualities, provides a unified way of modeling similarity and precision structure, and establishes a theoretical framework for the empirical investigation of imprecise qualities. On top of that, the regional framework also provides the foundation for further innovations of the quality-space model framework: in particular, for a field framework that provides ways of modeling probabilistic contents and phenomenal vivacity.

After all this, some might object that those who have written about modeling mental qualities implicitly had the regional framework in mind all along. However, there is a paucity of literature in both philosophy and cognitive science on how to model imprecise qualities, and the isomorphism claims often advanced in discussions of quality-space models is evidence that the regional framework is not taken for granted. Nevertheless, the principal aim of this paper is to build on existing research on the modeling of mental qualities. Whether or not the regional framework has been implicitly assumed, it has certainly not been explicitly developed, and its implications for the structure of experience have not been widely appreciated.

An upshot of this paper is that the structure of the mental qualities of conscious experiences is fundamentally different from the structure of the physical qualities of external objects. Whereas both mental qualities and physical qualities have similarity and magnitude structure, only mental qualities have precision structure. Because of this, the isomorphism thesis discussed at the beginning of this paper is false. Since the standard framework is adequate for representing similarity and magnitude structure, the framework suffices for modeling physical qualities. But to capture the precision structure of mental qualities, we cannot merely refurbish existing models. Instead, we must reshape the entire quality-space model framework.

On a final note, this paper has aimed to exhibit the utility of *formal phenomenology,* or the application of formal tools to the study of consciousness. In philosophy and cognitive science, theoretical work on consciousness has tended to focus on the mind-body problem, the nature of perception, the neural correlates of consciousness, and psychophysical relations. These issues are important and interesting, but there is also a distinct project of modeling the structures of conscious experiences themselves that can be developed without resolving these long-standing theoretical disputes. I believe that progress in such a project is a promising means of advancing our understanding of consciousness.

REFERENCES

Anton-Erxleben, K., Herrmann, K. & Carrasco, M. (2013). Independent effects of adaptation and attention on perceived speed. Psychological Science, 24 (2), 150-159. 10.1177/0956797612449178

Beals, Richard ; Krantz, David H. & Tversky, Amos (1968). Foundations of multidimensional scaling. *Psychological Review* 75 (2):127-142.

Block, Ned (2012). The Grain of Vision and the Grain of Attention. *Thought: A Journal of Philosophy* 1 (3):170-184.

Block, Ned (2015 a). The Puzzle of Perceptual Precision. Open Mind.

Block, Ned (2015 b). Solely Generic Phenomenology. Open MIND 2015.

Brown, Richard (ed.) (2013). Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience. Springer Studies in Brain and Mind.

Bruns, P., Camargo, C. J., Campanella, H., Esteve, J., Dinse, H. R., & Röder, B. (2014). Tactile acuity charts: a reliable measure of spatial acuity. PloS one, 9(2), e87384. doi:10.1371/journal.pone.0087384

Byrne, Alex (2011). Sensory qualities, sensible qualities, sensational qualities. In Brian McLaughlin, Ansgar Beckermann & Sven Walter (eds.), *The Oxford Handbook of Philosophy of Mind*. Oxford: Oxford University Press.

Chasid, Alon (2013). A Case against Representationalism. *Iyyun* 62 (1):29-42.

Clark, Austen (1985). A Physicalist Theory of Qualia. *The Monist* 68 (4):491-506.

Clark, Austen (1993). *Sensory Qualities*. Oxford University Press UK.

Clark, Austen (2000). Quality space. In Austen Clark (ed.), *A Theory of Sentience*. Oxford University Press.

Crane, Tim and French, Craig, "The Problem of Perception", The Stanford Encyclopedia of Philosophy (Spring 2017 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2017/entries/perception-problem/>.

Cutter, Brian (2019). Indeterminate perception and colour relationism. *Analysis* 79 (1):25-34.

Denison, Rachel N. (2017). Precision, Not Confidence, Describes the Uncertainty of Perceptual Experience: Comment on John Morrison's "Perceptual Confidence". *Analytic Philosophy* 58 (1):58-70.

Gescheider, G. A. (1997). Psychophysics: The fundamentals (3rd ed.). Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers.

Fazekas, Peter & Overgaard, Morten (2018). A Multi-Factor Account of Degrees of Awareness. *Cognitive Science* 42 (6):1833-1859.

Fechner, G., 1860, *Elements of Psychophysics*, H.E. Adler (trans.), New York: Holt, Reinhart & Winston, 1966.

Fink, Sascha Benjamin (2015). Phenomenal precision and Some Possible Pitfalls – A Commentary on Ned Block. Open MIND.

Fujita, O. Japan J. Indust. Appl. Math. (2013) 30: 1. https://doi.org/10.1007/s13160-012-0089-6

Gärdenfors, Peter (2014). *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. MIT Press.

Gert, Joshua (2017) Quality spaces: Mental and physical, Philosophical Psychology, 30:5, 525-544, DOI: 10.1080/09515089.2017.1295303

Gescheider, G. A. (1997). Psychophysics: The fundamentals (3rd ed.). Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers.

Goodman, Nelson (1951). *The Structure of Appearance*. Harvard University Press.

Hellie, Benj (2005). Noise and perceptual indiscriminability. *Mind* 114 (455):481-508.

Hilbert, David R. & Byrne, Alex (2008). Basic sensible qualities and the structure of appearance. *Philosophical Issues* 18 (1):385-405.

Horst, Steven (2005). Phenomenology and psychophysics. *Phenomenology and the Cognitive Sciences* 4 (1):1-21.

Intriligator, J., and P. Cavanagh. "The Spatial Resolution of Visual Attention." Cognitive Psychology 43 (2001): 171–216.

Kind, Amy (2017). Imaginative Vividness. *Journal of the American Philosophical Association* 3 (1):32-50.

Klincewicz M. (2011) Quality Space Model of Temporal Perception. In: Vatakis A., Esposito A., Giagkou M., Cummins F., Papadelis G. (eds) Multidisciplinary Aspects of Time and Time Perception. Lecture Notes in Computer Science, vol 6789. Springer, Berlin, Heidelberg

Kostic, Daniel (2012). The vagueness constraint and the quality space for pain. *Philosophical Psychology* 25 (6):929-939.

Logvinenko, A. D. (2015). The geometric structure of color. *Journal of Vision, 15*(1). Article ID 16.

Luce, R. D., & Krumhansl, C. L. (1988). Measurement, scaling, and psychophysics. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Stevens' handbook of experimental psychology: Perception and motivation; Learning and cognition* (pp. 3-74). Oxford, England: John Wiley & Sons.

Lycan, William, "Representational Theories of Consciousness", The Stanford Encyclopedia of Philosophy (Summer 2015 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2015/entries/consciousness-representational/>.

Morrison, John (2016). Perceptual Confidence. Analytic Philosophy 57 (1):15-48.

Munsell, Albert Henry (1905). *A Color Notation*. Boston, MA: Munsell Color Company.

Murray, David J. (1993). A perspective for viewing the history of psychophysics. Behavioral and Brain Sciences 16 (1):115.

Renero, Adriana (2014). Consciousness and Mental Qualities for Auditory Sensations. *Journal of Consciousness Studies* 21 (9-10):179-204.

Rockafellar, R. Tyrrell; Wets, Roger J-B (2005). *Variational Analysis*. Springer-Verlag. p. 117. ISBN 3-540-62772-3.

Rosenthal, David (2010), "How to Think about Mental Qualities," Philosophical Issues: Philosophy of Mind, 20 (October): 368-393.

Rosenthal, David (2015), "Quality Spaces and Sensory Modalities," in The Nature of Phenomenal Qualities: Sense, Perception, and Consciousness, ed. Paul Coates and Sam Coleman, Oxford: Oxford University Press, pp. 33-65.

Rosenthal, D. M. (2016). Quality spaces, relocation, and grain. In J. R. O'Shea (Ed.), Wilfrid Sellars and his legacy. Oxford: Oxford University Press.

Sellars, W. (1963). Science, perception and reality. London: Routledge.

Soteriou, Matthew, "The Disjunctive Theory of Perception", The Stanford Encyclopedia of Philosophy (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2016/entries/perception-disjunctive/>.

Stazicker, James (2011). Attention, Visual Consciousness and Indeterminacy. *Mind and Language* 26 (2):156-184.

Stevens, S.S. (1946). On the Theory of Scales of Measurement. *Science* 103, 677-680.

Strasburger, Hans; Rentschler, Ingo; Jüttner, Martin (2011). "Peripheral vision and pattern recognition: A review". *Journal of Vision*. 11 (5): 13.

Tao, Terence (2011). *An Introduction to Measure Theory Vol. 126*. Amer Mathematical Society.

Thomas, Nigel J.T., "Mental Imagery", The Stanford Encyclopedia of Philosophy (Spring 2018 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2018/entries/mental-imagery/>.

Tononi, Giulio (2007). The information integration theory of consciousness. In Max Velmans & Susan Schneider (eds.), *The Blackwell Companion to Consciousness*. Blackwell. pp. 287--299.

Williamson, Timothy (1994). *Vagueness*. Routledge.

Wilson, Jessica M. (2013). A Determinable-Based Account of Metaphysical Indeterminacy. *Inquiry: An Interdisciplinary Journal of Philosophy* 56 (4):359-385.

Young, Benjamin D., Keller, Andreas, and Rosenthal, David (2014), 'Quality Space Theory in Olfaction', in Frontiers in Consciousness Research.