
Modeling Mental Qualities

* Published in the *Philosophical Review* (2021) 130 (2): 263–298.

Andrew Y. Lee | andrew.yuan.lee@gmail.com | <http://andrewyuanlee.com/>
University of Oslo, Department of Philosophy

Abstract:

The standard framework for modeling the mental qualities of conscious experiences represents them via points in geometrical spaces, where distances between points inversely correspond to degrees of phenomenal similarity. This paper argues that the standard framework is structurally inadequate and develops a new framework that is more powerful and flexible. The core problem for the standard framework is that it cannot capture *precision structure*: for example, consider the phenomenal contrast between seeing an object as crimson in foveal vision versus merely as red in peripheral vision. The solution I favor is to model mental qualities using regions, rather than points. I explain how this seemingly simple move not only provides a natural way of modeling precision, but also yields a variety of further theoretical fruits: it enables us to formulate novel hypotheses about the spaces and structures of mental qualities, formally differentiate two dimensions of phenomenal similarity, generate a probabilistic model of the phenomenal sorites and a formal measure of discriminatory grain, and acquire a new theoretical tool for the empirical investigation of consciousness. A noteworthy consequence is that the structure of the mental qualities of conscious experiences is not isomorphic to the structure of the perceptible qualities of external objects.

KEYWORDS: quality space, mental qualities, phenomenal precision, formal phenomenology, structure of experience, phenomenal sorites, phenomenal similarity, perceptual confidence

Introduction

Conscious experiences are characterized by mental qualities, such as those involved in seeing red, feeling pain, or smelling cinnamon. The standard approach to modeling mental qualities is to target a particular domain, such as color experience, and to develop a *quality-space model* for that domain. These models represent mental qualities via points in geometrical spaces, where points that lie closer in the space correspond to mental qualities that are more similar to each other. For example, in the canonical three-dimensional model of color qualities, any particular color quality 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 model of the structure of color experience.

This framework for modeling mental qualities is highly promising, since every domain of mental qualities is structured by similarity relations. While the model for color qualities is the most developed, in recent years quality-space models have also been proposed for pain, temporal experience, auditory experience, and olfactory experience. And even domains of experience that are difficult to empirically investigate, such as emotional or cognitive experience, arguably could be modeled in such a way if only we knew the relevant structural facts. If we wish to map the structures of conscious experiences, it seems a significant part of the project will consist in constructing quality-space models across different experiential domains.¹

This paper (1) argues that this standard framework is structurally inadequate, and (2) develops a new framework that is more powerful and flexible. The core limitation of the standard framework is that it overlooks what I call *precision structure*. Consider, for example, the phenomenal contrast between seeing an object in foveal vision as crimson versus seeing an object in peripheral vision merely as red (rather than as any particular shade of red): in such a case, your foveal visual experience is more precise than your peripheral visual experience. I will argue that imprecise qualities do not correspond to individual points within standard models, and so cannot be captured using the standard framework. Though this may at first

¹ See Churchland [2005] on color experience, Klincewicz [2011] on temporal experience, Kotic [2014] on pain, Renero [2014] on auditory experience, and Young, Keller, & Rosenthal [2014] on olfactory experience. See Clark [2000] and Rosenthal [2010, 2015] for general discussions of the quality-space model framework.

appear to be a minor technical challenge, solving the problem has significant philosophical and methodological ramifications.

The basic idea behind my new framework is to model mental qualities using regions (rather than points) in geometrical spaces. This seemingly simple innovation not only provides a natural way of capturing imprecise experiences, but also yields a variety of other philosophical fruits. In particular, my new framework enables us to formulate novel hypotheses about the space and structure of mental qualities, formally differentiate two dimensions of phenomenal similarity, identify the connections between precision and discriminatory grain, generate a probabilistic model of the phenomenal sorites, and deploy a new theoretical tool for the empirical investigation of conscious experiences.

A core lesson of this paper is that precision is much more significant for understanding the structure of experience than has been previously appreciated. On the picture I develop, precision is a structural feature of experience, akin to similarity and magnitude. A noteworthy consequence of the new framework is that the structure of the mental qualities of conscious experiences is fundamentally different from the structure of the perceptible qualities of external objects: only mental qualities have precision structure. This subverts the common assumption that mental quality-spaces are isomorphic to the corresponding perceptible quality-spaces.

A more general goal of this paper is to illustrate the prospects for *formal phenomenology*, or the application of formal tools to the study of conscious experiences. Conscious experiences are richly structured, yet there is relatively little work that attempts to capture that structure via formal models.² This paper aims to exhibit why such an approach is a promising means of advancing consciousness research.

§1 explains what precision is and why it poses a problem for standard models of mental qualities; §2 develops my new framework for modeling mental qualities, which I call the *regional framework*; §3 discusses applications of the regional framework to issues concerning discriminatory grain, the phenomenal sorites, empirical investigation, and perceptual confidence theories of perceptual experience.

² For some other examples of recent work applying formal tools to consciousness research, see Tononi [2007], Yoshimi [2011], Prentner [2019], and Kleiner [2020].

§1 | The Standard Framework

I will 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 and the relations between them.³ The standard approach to modeling mental qualities is to represent qualities via points in geometrical spaces. Under this approach, any quality-space model aims to represent a domain of mental qualities such 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.⁴ Consider, for example, how the three-dimensional 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 perceptible qualities of external objects (such as red).⁵ But the focus of this paper is solely on mental qualities. As I discuss later, precision is not a property of perceptible qualities, so the framework I develop is inapplicable to perceptible qualities. For the rest of the paper, whenever I use the term ‘quality’ without qualification I will mean mental qualities.

³ There is often ambiguity between talking about the formal representation of a domain of qualities versus the domain of qualities itself: for example, consider Clark [2000, p.4]’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 ‘quality-space’ to mean the domain of qualities and ‘quality-space model’ to mean the formal representation of those qualities.

⁴ Quality-space models often aim to also capture magnitude relations, such 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 brightness dimension than *b*). To simplify the exposition, I focus mainly on similarity.

⁵ See Byrne [2011] for more on the distinctions between different kinds of qualities.

And for brevity, I will use terms such as ‘hue’ rather than ‘phenomenal hue’ to designate the dimensions of mental qualities.

I take for granted that what it is like to undergo a conscious experience is partly constituted by which mental qualities are instantiated by that conscious experience. Otherwise, I will be neutral on most questions about the nature of mental qualities. In particular, I remain neutral on whether mental qualities can be instantiated even in the absence of consciousness, on the metaphysical relationship between mental qualities and perceptible qualities, and on whether mental qualities are fundamentally physical. Staying neutral on these issues ensures that the framework I develop is compatible with a wide variety of philosophical positions.

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 (such as color qualities), we must empirically investigate the relevant domain. But in order to develop an adequate framework, we must specify what kind of formal structure is required for modeling any arbitrary domain of mental qualities, regardless of how any particular model is structured. I will eventually argue that in order to capture the precision structure of mental qualities, we must make basic changes to the whole quality-space model framework (rather than just revisions to particular models). And in §3, I will discuss how empirical methods can be used to construct particular models within my new framework.

There is a diverse body of literature in both philosophy and cognitive science pertaining to the modeling of mental qualities. The relevant philosophical literature has focused mainly on questions about the relationship between mental qualities and perceptible qualities and on the nature of mental qualities.⁶ The relevant cognitive science literature has focused mainly on issues concerning the psychophysical

⁶ For some classic and contemporary philosophical texts addressing these questions, see Goodman [1954] and Clark [2000], Rosenthal [2000, 2015]. For recent discussions of the nature of phenomenal qualities, see Coates & Coleman [2015].

relations between physical stimuli and mental qualities and the challenges in measuring mental qualities.⁷ However, in both disciplines, research that directly addresses the modeling of mental qualities tends to focus on similarity structure, leaving out precision structure.

This lacuna may be partly due to the common assumption that models of mental qualities are isomorphic to models of perceptible qualities. The standard methodology for constructing a model of mental qualities is to first use data concerning perceptual discrimination judgments to construct a model of perceptible qualities, and to then extrapolate from that 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 perceptible 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 perceptible quality-spaces.⁹ I will eventually argue that this isomorphism thesis is false: models of mental qualities require more structure than models of perceptible qualities, for only mental qualities have precision structure.

⁷ For 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 an approach to geometrically modeling concepts that shares some (though not all) formal features with my framework, see Gärdenfors [2014].

⁸ Perceptible qualities (e.g. colors) should be distinguished from physical stimuli (e.g. specific wavelengths of light). The isomorphism thesis is intended to apply only to the former. See Clark [2000] for more on this distinction.

⁹ As examples, Sellars [1963, p.48] talks of an "isomorphism of acts of sense and material things," Palmer [1999, p.933] says that the isomorphism "function maps color experiences onto points in a dimensional color space such that relations among color experiences...are preserved by corresponding relations among corresponding points in space," Churchland [2007, p.119] says there is a "homomorphism" from the space of objective colors to "the internal structure of human phenomenological color space," and Rosenthal [2016, p.165] 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 individual qualities) and a distance metric (where distances between points are inversely correlated with degrees of similarity between the qualities represented by those points). In what follows, I explain why the standard framework is structurally inadequate.

Imprecise Experiences

Consider your color experience in foveal vision versus in peripheral vision. In foveal vision, you see an object as a specific shade of red, such as crimson. But in peripheral vision, you no longer see it as any specific shade of red, but instead just as red. It is not merely that you see the object as a different specific shade of red across the two cases. Instead, even if your peripheral color experience represents its object as having *some* specific shade of red or other, it leaves open *which* shade of red that might be, and it is compatible with your experience that you are seeing *any* given shade of red within a certain range.¹⁰ Speaking somewhat metaphorically, peripheral color experience is less sharp and crisp than foveal color experience. This difference in phenomenal character is a matter of what I call *precision*.¹¹

Though I focus on precision with respect to color experience across foveal and peripheral vision, there are other examples that may also be used to illustrate the phenomenon. Consider the phenomenal contrasts between your color experience of an object that is far away versus nearby, or between your spatial visual experience with vision correction lenses versus without, or between your tactile experience while touching a texture with your fingertips versus while touching a texture with your back. In each case, the former experience is more precise than the latter.

¹⁰ See Hansen, Prajecus, & Gegenfurtner [2009] for psychophysical evidence of this contrast. See Strasburger et al [2011] for a recent review of the science of peripheral vision.

¹¹ I take the term ‘precision’ from Block [2015]. A number of philosophers have used the terms ‘determinacy’ or ‘determinability’ for what I call ‘precision’, but I will later explain why ‘precision’ is better.

Precision is different in kind from familiar phenomenal properties such as hue, loudness, or painfulness. Those phenomenal properties correspond to dimensions of quality-space models. But as we will see, precision cannot be captured in the same way (at least not without making some substantive theoretical assumptions). As some initial evidence for this, consider how a visual experience could (arguably) be precise with respect to color yet imprecise with respect to shape or how qualities even across different modalities can be similar with respect to precision. Instead of thinking of precision as merely another dimension of mental quality-spaces, it is more apt to think of precision as a structural feature of experience (like similarity or magnitude). This hints at why modeling precision requires modifying the entire quality-space model framework rather than just patching up particular models.¹²

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 the representational contents of experiences. For naïve realists, it is natural to think that precision is a matter of being perceptually acquainted with more determinate properties of external objects. For qualia theorists, it is natural to think that precision is a structural property of phenomenal character somewhat akin to the resolution of an image. This paper remains neutral on these issues, and the framework I develop will be deployable by theorists across the board. For ease of explication, I will often talk of mental qualities representing perceptible qualities, but my discussion could likewise be framed in terms of other relations (such as perceptual acquaintance or causal correspondence).¹³

There is a mix of literature across both philosophy and cognitive science pertaining to imprecise qualities. In philosophy, there has been recent work examining imprecise qualities in connection with philosophical theories of perception, generic

¹² Notably, Block [2015] says that the notion of “phenomenal precision [is] obscure...we have a well-developed science of perception but very little science of the phenomenology of perception.” This paper aims to help bridge this gap.

¹³ For a general overview of theories of perception, see Crane & French [2017]. For argument against representationalism about precision, see Block [2015].

phenomenology, and the representational contents of experience.¹⁴ However, these discussions have not directly addressed how precision structure relates to similarity structure or how to integrate imprecise qualities into the quality-space model framework. In cognitive science, there is 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.¹⁵ However, these discussions tend to focus on perceptual capacities and their functional roles, rather than on the structure of the mental qualities associated with those perceptual capacities.

Precision vs. Determinability vs. Noise

Before we turn to why precision poses a problem for the standard framework, we need to first see why precision is conceptually distinct from two other phenomena: namely, determinability and noise. This will both sharpen our understanding of precision and set the stage for some of discussion in §3.

Consider first *determinability*, or the relation between determinates and determinables. 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 determinability are independent, consider first 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

¹⁴ See Block [2015] on precision's implications for the nature of perception, Cutter [2019] on precision and color, Fink [2015] and Fazekas & Overgaard [2018] on precision in relation to generic phenomenology and empirical methods for investigating consciousness, Stazicker [2011], Morrison [2016], Munton [2016], and Nanay [forthcoming] on the relation between precision and representation, and Hellie [2005] and Pelling [2008] on precision and perceptual discrimination.

¹⁵ See Intriligator & Cavanagh [2001] and Block [2012] on the relationship between visual resolution and visual attention, Anton-Erxleben & Carrasco [2013] on attention, spatial resolution and the role of receptor cell density for perceptual discrimination, Bruns et al [2014] on tactile spatial resolution, Debats et al [2012] on measuring haptic precision, Denison [2017] on precision in relation to perceptual uncertainty.

represent any specific shade of color. Conversely, consider the determinable phenomenal property that has as determinates the precise color phenomenal properties characterizing your foveal visual experiences when looking at a series of 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 they all represent specific shades of color. Since there are both determinate imprecise phenomenal properties and determinable precise phenomenal properties, precision and determinability are doubly dissociable.

Precision and determinability are liable to be confused because of systematic ambiguities in natural language. For example, 'phenomenal red' can mean either a determinable phenomenal property that has precise determinates (e.g. the disjunction of phenomenal crimson, phenomenal scarlet, etc.) or a determinate phenomenal property that is imprecise. The risk of confusing precision and determinability is also heightened by the fact that precision of mental qualities is inversely correlated with determinability of the perceptible qualities represented by those mental qualities: for example, the determinate crimson is represented by a precise red experience whereas the determinable red is represented by an imprecise red experience.¹⁶

These conceptual distinctions are worth highlighting, for taking precision to be merely a matter of determinability 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. By contrast, we will soon consider a variety of philosophically substantive questions about the space and structure of imprecise qualities, the implications of precision for phenomenal similarity, how precision relates to the nature of discriminability and the empirical investigation of conscious experiences, and why precision challenges the isomorphism thesis mentioned earlier.

At this point, some might raise the rather radical hypothesis that imprecise qualities involve the instantiation of a determinable without the instantiation of any

¹⁶ Note that representationalism does not undermine my point, since the property of representing a determinable property need not itself be a determinable property.

of its determinates.¹⁷ However, even if this view were correct about the metaphysics, it would still be important to conceptually distinguish precision from determinability. Otherwise, not only would there be systematic terminological ambiguities of the kind mentioned above, but we would also lack the conceptual distinctions needed to resolve those ambiguities. Consider, for example, the claim that an experience x instantiates the determinable phenomenal red. If we were to collapse the distinction between precision and determinability, then we would be unable to disambiguate between two interpretations of that claim: first, that x is precise and instantiates one amongst many determinate phenomenal properties, and second, that that x is imprecise and instantiates the determinable phenomenal red (without instantiating any of its determinates). No matter which view one favors about the metaphysics of precision, one ought to conceptually distinguish precision from determinability.¹⁸

Now let us turn to *noise*, or the degree of random variation within a psychophysical channel. A *psychophysical channel* is a causal pathway connecting a stimulus (such as a color chip) to a signal (such as a color experience). The noisier a channel, the less the channel's source determines its signal. A hypothesis that some might find attractive is that precision is simply the phenomenal manifestation of psychophysical noise.¹⁹ Though this is an interesting empirical hypothesis, it is important to appreciate why precision and noise are at least conceptually distinct.

To illustrate with a contrast case, consider first Achilles, whose color perception system stochastically generates one of a number of precise color experiences upon detection of a color. If Achilles looks at a scarlet color chip ten times, he might first have a scarlet experience, then a crimson experience, then a vermilion experience, and so forth. Achilles' color experiences are precise yet result from noisy psychophysical processes. Conversely, consider Zagreus, who lacks foveal vision but otherwise has an extremely reliable color perception system. If Zagreus looks at a

¹⁷ See Wilson [2013] for a general defense of this metaphysical view.

¹⁸ A similar confusion occurs with imprecision and vagueness. For brevity, I will simply note that vagueness is typically understood as a property of terms or concepts whereas imprecision is a property of phenomenal properties or experiences, that terms and concepts for imprecise qualities can be sharp, and that none of the theories of precision I mentioned (p. 7) appeal to vagueness. For an overview of vagueness, see Williamson [1994].

¹⁹ See Hellie [2005] for an example of this kind of view.

scarlet color chip ten times, then he has the exact same color experience every single time, but that color experience is the same as the one you would have when looking at the scarlet chip via peripheral vision. Zagreus' color experiences are imprecise yet result from unnoisy psychophysical processes.

In psychophysics, statistics, and related fields, the term 'precision' refers to the reciprocal of the degree of variance in a set of measurements, which does in fact vary inversely with noise. But *statistical precision* (the statistical property of measurements just defined) is distinct from *phenomenal precision* (the phenomenal property of mental qualities this paper focuses on). As an analogy, imagine a gun firing a round of bullets at a target: statistical precision corresponds to how closely clustered the bullet holes are to each other, whereas phenomenal precision (inversely) corresponds to the surface area of the bullet holes. These structural differences hint at the formal treatment of precision that will come later.

Although precision and noise are conceptually distinct, it is empirically possible that imprecise mental qualities result from noisy psychophysical processes. Suppose our cognitive systems are structured so that whenever a psychophysical channel is noisy (prior to the generation of an experience), it reliably generates an imprecise mental quality (rather than stochastically generates a precise mental quality). For example, suppose you see a scarlet color chip, but the noise in your color perception system renders it uncertain whether the chip is scarlet or vermilion or crimson, so your color perception system generates a red experience. If such a hypothesis is correct, then imprecise mental qualities would be the phenomenal manifestation of noise (though as a matter of empirical fact, rather than conceptual necessity). For the purposes of this paper, I will stay neutral on this hypothesis. While resolving the issue would paint a richer picture of the psychophysical processes resulting in imprecise experiences, it would still leave open how to best model imprecise mental qualities using the quality-space model framework.

The Problem of Precision

We are now in position to see why standard models cannot capture precision. The core problem is that in the standard framework, individual qualities are represented by individual points in quality-space models, but no individual points in such models are adequate for representing imprecise qualities. Putting it another way, the standard framework takes mental qualities to be specifiable by a single

value along each of the dimensions of the model, but imprecise qualities seem to instead correspond to ranges of values along those dimensions.²⁰

It may be tempting to attempt to solve the problem by simply adding an extra dimension (representing degree of precision) to existing models. But such an approach would still require assigning individual values along ordinary dimensions (such as hue) to imprecise qualities, would be unable to capture similarity relations between imprecise mental qualities belonging to different quality-spaces, and (as we will discuss in the next section) would be unable to accommodate views that allow the precision of one dimension (such as hue) to vary independently of the precision of other dimensions (such as brightness). To capture precision, we need more than just tweaks to existing models; instead, we need structural changes to the whole framework.

These problems may make some wonder whether the standard framework was always meant to be an idealization. What if the standard framework was never even intended to capture all mental qualities? However, there is a paucity of literature in both philosophy and cognitive science on how to model imprecise qualities, and the isomorphism claims frequently advanced in discussions of quality-spaces are evidence that the standard framework has been presumed to be representationally adequate. Nevertheless, the principal aim of this paper is to build on existing research on modeling 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.

§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 develop the formal structure of the regional framework, address some of the technical challenges, and explain why the framework is theoretically fruitful and philosophically significant.

²⁰ A related idea is that more precise mental qualities have more specific contents (in that they eliminate more possibilities). This aspect of precision is not captured by the standard framework, since every mental quality simply corresponds to a single point. By contrast, it will be obvious how the regional framework does better.

The Basic Formal Structure

To explain the regional framework, it is useful to first contrast it with the standard framework. Any standard model requires a way of representing individual qualities and a way of representing degrees of similarity between qualities. In light of this, we can think of standard models as comprised of a pair of elements: a set \mathcal{S} of points (representing individual qualities) and a distance metric d over those points (where greater distances²¹ map to lower degrees of phenomenal similarity).²²

There are three main desiderata when constructing a model in the standard framework. First, points in the model should stand in one-to-one correspondence with qualities in the target quality-space. 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 structure of the model mirrors the structure 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 inversely corresponds to the degree of precision of the quality represented by that region. More specifically: mental quality x_A is more precise than x_B just in case region A (corresponding to x_A) is smaller than region B (corresponding to x_B). In other words, the ordinal structure of region sizes models the ordinal structure of degrees of precision.²³

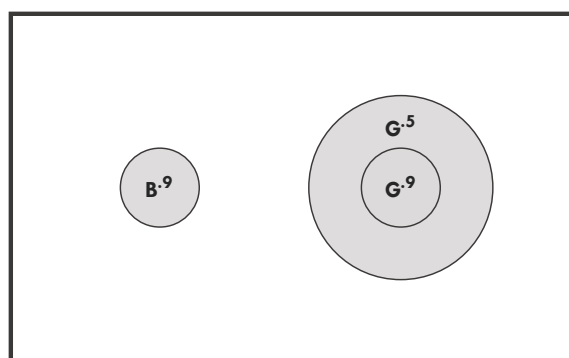
²¹ I assume that a metric is needed to capture relations of phenomenal similarity, though see Gert [2017] for an opposing view. Note that this assumption makes the task of developing the regional framework harder, since it is unobvious how to develop a metric over regions that corresponds to phenomenal similarity. For discussion of how metric structure can be extracted from ordinal judgments, see Beals et al [1968].

²² Note that the set \mathcal{S} of points and metric d suffice to determine the dimensionality of the space. See Hurevicz & Wallman [1948] on measures of dimensionality.

²³ There is a further question of whether differences or ratios between region sizes correspond to differences or ratios between degrees of precision. This depends on whether precision itself has ordinal, interval, or ratio structure, a question that I will remain neutral on in this paper. See Stevens [1946] for discussion of these different measurement scales.

We saw above that standard models can be specified with just a set \mathcal{S} of points and a distance metric d . The regional framework requires adding more structure. To specify a regional model, we need not only the set \mathcal{S} of points and the point-distance metric d , but also a set \mathcal{R} of regions (meaning a subset of the powerset of \mathcal{S}), a measure μ on \mathcal{S} , and two metrics on regions which will be formally defined later. Over the course of this section, I will explain each of these elements in detail.²⁴

A few notes on terminological conventions: I will denote regions using small-caps letters (ex: region A), perceptible qualities using lowercase letters (ex: perceptible quality a), mental qualities using lowercase x with a subscript for the corresponding region (ex: mental quality x_A), functions using lowercase script letters (ex: the measure μ) and elements of the regional framework using uppercase script letters (ex: the set \mathcal{S} of points). I will also illustrate regions using diagrams like the one below, where precision values are denoted by real numbers from 0 to 1 in superscript, with higher numbers denoting higher degrees of precision:²⁵



B^9 represents a relatively precise blue quality.

G^9 represents a relatively precise green quality.

G^5 represents a relatively imprecise 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 spaces of regional models, focusing mostly on the set \mathcal{S} of points and the point-distance metric

²⁴ Strictly speaking, specifying a regional model requires specifying only \mathcal{S} , \mathcal{R} , d , and μ . This is because the metrics on regions (*qual* and *prec*) can be recovered from d and μ .

²⁵ The box represents a space, the bounded shapes represent regions, and the regions are denoted using the linguistic convention described above. The interpretation of the dimensions in the figures is not important.

d. Second, I discuss the structure of imprecise qualities, focusing mostly on the set \mathcal{R} of regions. Third, I discuss similarity with respect to qualitative character versus similarity with respect to precision, focusing mostly on the new metrics on regions, *qual* and *prec*. Along the way, I explain how the formal framework interacts with a variety of philosophical issues.

The Structure of the Space

The space of any given quality-space model is determined by the set \mathcal{S} of points and the point-distance metric d . Since these are the elements that characterize standard models, most of the theoretical issues concerning these elements have already been addressed in prior work. But there is one new question that arises with the regional framework that is worth addressing here: namely, whether the spaces in regional models should be discrete or continuous. A discrete model would enable us to simply “export” all the points of the set \mathcal{S} of a standard model into a regional model.²⁶ In contrast, a continuous model would require a new set \mathcal{S} that has a continuous structure.

Though discrete models are finite, 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 even the most precise qualities characterizing human experiences. Furthermore, it is mathematically simpler to specify formal constraints on regions in continuous spaces, 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 still apply 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 (consider how the interval of real numbers from 0 to 1 is continuous but bounded by the limit points 0 and 1).

In a continuous model, individual points are probably best thought of as idealizations: they are the maximally specific values of the dimensions of a quality-space, even if it turns out that no mental qualities actually correspond to regions

²⁶ 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 within any given quality-space).

comprised of a single point. This may raise the worry that regional models have more structure than is strictly necessary. However, note that scientific models often idealize, especially when doing so leads to simpler formalisms. For example, suppose that all physical objects are composed of particles, that there are finitely many kinds of elementary particles, and that all elementary particles of the same kind have the same mass value. Then there are some mass values that no physical objects could have, since no combination of elementary particles would generate that mass value. Nevertheless, it is still useful to represent mass values 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

Any regional model must specify a set \mathcal{R} of *regions*, which are subsets of the set \mathcal{S} of points. As we will see, this new element generates new philosophical questions about the space and structure of imprecise qualities that are difficult to even formulate without the appropriate theoretical resources.

Why is there a need for \mathcal{R} at all? It may be tempting to think that imprecise qualities simply correspond to the subsets of \mathcal{S} , and that there is no need to posit a whole new set \mathcal{R} . However, \mathcal{R} is a crucial element in the regional framework, for it allows us to distinguish different theories of the structure and space of imprecise qualities. As we will see, it may not be the case that every subset of \mathcal{S} corresponds to a genuine mental quality. Consequently, we need a way of distinguishing regions, which can be any subset of \mathcal{S} , from *permissible regions*, which are the subsets of \mathcal{S} that are members of \mathcal{R} and that are to be interpreted as representing genuine mental qualities. To put it another way, different specifications of \mathcal{R} correspond to different theories of the space and structure of mental qualities.

Since there are as many ways of specifying \mathcal{R} as there are sets of subsets of \mathcal{S} , it is useful to focus on formal constraints on \mathcal{R} that permit different kinds of permissible regions. The most obvious constraints concern *sizes*: how large or small can permissible regions be? Questions about size constraints are questions about the limits of the degrees of precision. For example, we might wonder whether there are super-imprecise qualities whose regions cover entire quality-spaces or whether there are super-precise qualities whose regions correspond to a single point. I will

largely set aside questions about region sizes, though in §3 I will discuss empirical methods for investigating region sizes for particular mental qualities.

The more interesting class of constraints on \mathcal{R} concerns *shapes*: which kinds of shapes can permissible regions take? Questions about shape constraints are questions about the structure of precision. These questions are difficult to formulate linguistically, and it will take a bit of space to explain what different shape constraints look like. As examples, we will consider three different formal constraints on region shapes: CONNECTIVITY, CONVEXITY, and UNIFORMITY. Each of these constraints may be thought of as a distinct hypothesis about the space and structure of imprecise qualities. To get a feel for which kinds of shapes these formal constraints permit or exclude, consider the four sample regions in the diagram below:

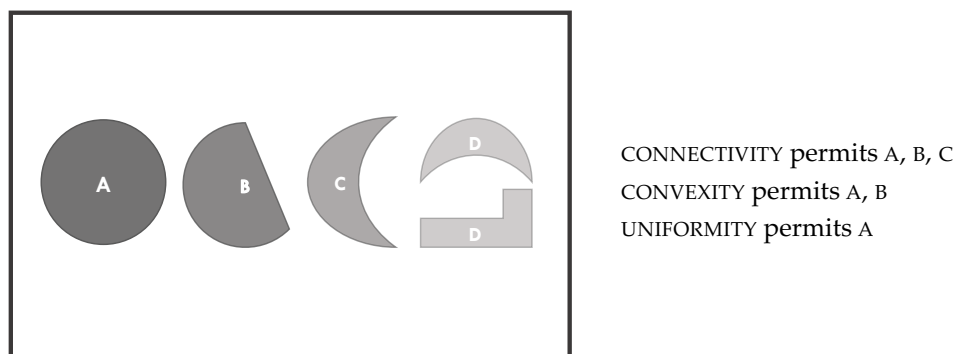


FIGURE 2: Formal constraints on permissible regions.

The most permissive of the three constraints is CONNECTIVITY, according to which all permissible regions are connected.²⁷ A region is *connected* just in case it has no discontinuities, meaning that the region is “all in one piece.” CONNECTIVITY excludes region D from FIGURE 2 but permits regions A, B, and C. More generally, CONNECTIVITY rules out mental qualities that represent arbitrary collections of perceptible qualities, such as a mental quality that represents something as either crimson or aquamarine (but nothing else). However, CONNECTIVITY still permits regions that are strangely shaped, such as region C.

²⁷ FORMAL DEFINITION: if $A \in \mathcal{R}$, then A is not the union of two disjoint open sets of \mathcal{S} , where A is open just in case $\forall a \in A, \exists \epsilon > 0$ where the open ball $B(a, \epsilon) = \{x \in \mathcal{S} \mid d(a, x) < \epsilon\} \in A$.

A stronger constraint is CONVEXITY, according to which all permissible regions are convex.²⁸ 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, any region must contain all points within the straight lines connecting its boundaries. CONVEXITY excludes regions C and D but permits regions A and B. This means that CONVEXITY is more restrictive than CONNECTIVITY but still allows irregularly shaped regions, such as region B.

An even stronger constraint is UNIFORMITY, according to which all permissible regions are balls.²⁹ A region is a *ball* just in case it includes all and only the set of points that are within a given distance from a center.³⁰ UNIFORMITY excludes regions B, C, and D, permitting only region A, meaning that the principle excludes all the irregular regions permitted by the previous constraints. But UNIFORMITY also excludes mental qualities that differ in their degree of imprecision across different dimensions, such as a color experience that is precise in hue but imprecise in brightness. This might make some worry that UNIFORMITY is too restrictive.³¹

The preceding discussion merely scratches the surface. My present aim is not to evaluate which of these formal constraints is most plausible, but instead to show how the regional framework provides us with tools that enable more rigorous and systematic theorizing about the space and structure of mental qualities. In fact, the regional framework enables us to formulate interesting hypotheses that are hard to even express without the appropriate framework. Consider how difficult it would be to demarcate these hypotheses or formulate the relevant questions using only natural language.

²⁸ FORMAL DEFINITION: if $A \in \mathcal{R}$, then $\forall a, c \in \mathcal{S}$ such that $d(a, c) > 0$, A contains all points $b \in \mathcal{S}$ such that $d(a, b) + d(b, c) = d(a, c)$.

²⁹ FORMAL DEFINITION: if $A \in \mathcal{R}$, then $\exists a \in A$ and $\epsilon > 0$ where $A = \{x \in \mathcal{S} \mid d(a, x) < \epsilon\}$.

³⁰ In continuous n -dimensional spaces, balls are usually bounded by n -dimensional spheres, though note that balls near the boundary points of a space may be non-spherical.

³¹ There are other natural constraints in between UNIFORMITY and CONVEXITY, such as the constraints that all permissible regions are (1) regular polygons or (2) ellipsoids.

Qualitative Similarity

In the standard framework, it is straightforward to define a point-distance metric d that represents the similarity relations between different mental qualities. By contrast, identifying the right similarity metric in the regional framework is much more challenging. For the regional framework, we need a metric that takes as input regions of arbitrary size and shape yet still outputs distances that systematically correspond to degrees of phenomenal similarity. More specifically, let x_A , x_B , and x_C be mental qualities and let A , B , and C be the corresponding regions: if a metric m is to serve as the similarity metric for the regional framework, it ought to satisfy the following constraints:³²

- (1) if x_A is phenomenally identical to x_B , then $m(A, B) = 0$.
- (2) if x_A is more phenomenally similar to x_B than to x_C , then $m(A, B) < m(A, C)$.

Before moving forward, let me first mention the *measure*, which will be used (alongside the point-distance metric d) to construct our candidates for metrics on regions. The measure μ takes as input a subset of \mathcal{S} and outputs a size value, which represents the degree of imprecision of the mental quality represented by that region.³³ In discrete models, it is natural to simply take the size of a region to be the number of points in that region. However, such a measure does not work well in continuous models, since basically all regions in continuous models have infinitely many points. Instead, we can appeal to the standard mathematical measure: the Lebesgue measure, which is a generalization of the notions of length, area, volume, and so forth. The mathematical details of the Lebesgue measure are not really philosophically relevant. What is important is that the Lebesgue measure has the prop-

³² Actually, I will later argue that these constraints ought to be relativized to different kinds of phenomenal similarity. But this point does not matter for the moment.

³³ Why is the measure on the set \mathcal{S} of points rather than the set \mathcal{R} of regions? Since a measure takes as input a subset of a set, a measure on \mathcal{R} would output the sizes of sets of regions (rather than the sizes of regions). Consequently, determining the size of a region in \mathcal{R} requires measuring the corresponding subset of \mathcal{S} .

erties we would intuitively want a measure to have: in particular, it produces intuitive size values in continuous spaces of arbitrary dimensionality.³⁴ And with μ on the table, we are in position to consider candidates for metrics on regions.

A first pass is to turn to the default way of determining distance between regions: the least distance metric, which takes the distance between regions A and B to be the lowest distance value between any pair of points between A and B.³⁵ However, suppose that A and B are distinct but overlap (meaning that A and B contain some but not all of the same points). Since A and B 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 metric 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 candidate is the average distance metric, *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, *avg*(A, B) takes a point in A, determines the average distance between that point and all the points in B, repeats the procedure for every other point in A, and then averages the averaged distance values.³⁶ Since the average distance metric is sensitive to all the points in A and B, it is an improvement over the least distance metric. Yet *avg* is also inadequate. 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 $d(a, b) > 0$. As a consequence, the average distance from a region to itself must be non-zero. But non-zero distance values represent phenomenally distinct qualities. This means we get the absurd result that all mental qualities that are not maximally precise are not maximally similar to

³⁴ See Tao [2011] for an overview of the Lebesgue measure.

³⁵ FORMAL DEFINITION: Let $\inf(A)$ denote the infimum of set A. Then the least distance from A to B = $\inf\{d(a, b) \mid a \in A, b \in B\}$.

³⁶ FORMAL DEFINITION: The formal definition of *avg* differs in discrete versus continuous spaces. For discrete spaces, $\text{avg}(A, B) = \sum_{a \in A} \sum_{b \in B} d(a, b) \div \mu(A)\mu(B)$. But since summations over divergent series are undefined, in continuous spaces summation must be replaced with integration, where $\text{avg}(A, B) = \int_A (\int_B d(a, b) d\mu(b)) d\mu(a) \div \mu(A)\mu(B)$. For a more comprehensive discussion of these average distance metrics, see Fujita [2013].

themselves. The source of the problem is that *avg* 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. But 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 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 it works, consider first an example illustrated by the diagram below:

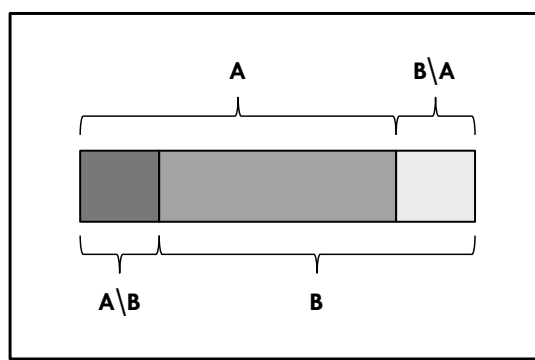


FIGURE 3: A partially overlaps with B.

Suppose we wish to determine the similarity between the mental qualities represented by A and B. The *qual* metric works by first taking the average distance from A to the subregion of B that does not overlap with A, and then taking 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). In other words, we find $avg(A, B \setminus A)$, then find $avg(B, A \setminus B)$, and then find the weighted average of those two results. The *qual* metric is illustrated pictorially in the diagram below:

$$qual \left[\begin{array}{c} \text{A} \\ \text{B} \end{array} \right] = \frac{avg \left[\begin{array}{c} \text{A} \cap \text{B} \\ \text{B} \setminus \text{A} \end{array} \right] \times \mu \left(\begin{array}{c} \text{B} \setminus \text{A} \end{array} \right) + avg \left[\begin{array}{c} \text{A} \setminus \text{B} \\ \text{A} \cap \text{B} \end{array} \right] \times \mu \left(\begin{array}{c} \text{A} \setminus \text{B} \end{array} \right)}{\mu \left[\begin{array}{c} \text{A} \cup \text{B} \end{array} \right]}$$

FIGURE 4: The qualitative similarity metric.

And the metric is expressed formally in the following equation:³⁷

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

To verify that *qual* produces the correct results, we can check its behavior across the different situations that can occur between two regions: namely, identity, disjointness, containment, and partial overlap:

1. If A is **identical** to B, then $qual(A, B) = 0$.
2. If A and B are **disjoint**, then $qual(A, B) = avg(A, B)$.
3. If A **contains** B, then $qual(A, B) > 0$.
4. if A and B partially **overlap**, then $qual(A, B) > 0$.

These results are exactly what we want out of a metric capturing phenomenal similarity. The first result means that whenever two regions are identical, they represent the corresponding mental qualities as phenomenally identical. The second result means that whenever two regions are disjoint, *qual* collapses to *avg*—and as noted earlier, *avg* produces intuitively correct verdicts in cases involving disjoint

³⁷ Why are the numerators $\mu(B \setminus A)$ and $\mu(A \setminus B)$ rather than simply $\mu(B)$ and $\mu(A)$? 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 $\mu(B)$ and $\mu(A)$, then the metric would no longer entail that $qual(A, A) = 0$. For more detailed discussion of the formal properties of *qual*, see Fujita [2013].

regions. The third result means that if region A contains region B, then their corresponding qualities are not represented as phenomenally identical. Moreover—presuming A contains B—increasing the difference in size between A and B also increases their distance, meaning that *qual* predicts that phenomenal similarity decreases as the difference in precision increases. Finally, the fourth result means that if A only partially overlaps with B, then their corresponding qualities are not represented as phenomenally identical. In fact, the more A and B overlap, the lower the distance between them, with the distance approaching zero as degree of overlap approaches identity.

These observations are evidence that *qual* satisfies the desiderata on a metric that were outlined earlier. And that 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.³⁸ As a bonus, *qual* even works in spaces whose dimensions are not linear orders. This point is important, for any framework for modeling mental qualities must be able to accommodate dimensions such as hue.

Precision Similarity

In spite of its virtues, there remains an aspect of phenomenal similarity that *qual* does not capture. Consider again the diagram from the start of this section:

³⁸ 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.

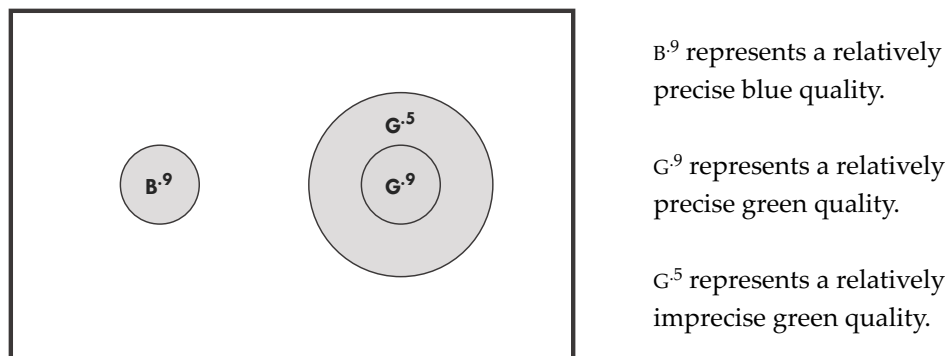


FIGURE 5: $qual(B^9, G^9) = qual(B^9, G^5)$.

Recall that when regions are disjoint, *qual* works the same way as *avg*. Because of this, $qual(B^9, G^9) = qual(B^9, G^5)$, meaning that *qual* predicts the precise blue quality to be as similar to the precise green quality as to the imprecise green quality. But B^9 is more similar to G^9 than it is to G^5 . After all, B^9 and G^9 are similar with respect to precision, whereas B^9 and G^5 are not. The heart of the issue is that *qual* is not directly sensitive to the relative sizes of regions, even though relative size seems to track one aspect of phenomenal similarity. How should we proceed in light of this result?

My view is that this result is a feature rather than a flaw, for we are now in position to formally distinguish 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 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 I have argued does not correspond to any particular dimension. It is easy to get an intuitive grip on the difference between these two kinds of phenomenal similarity. Consider the difference between comparing a precise phenomenal red quality to a precise phenomenal orange quality versus between comparing a precise phenomenal red quality to an imprecise phenomenal red quality. The 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.

The core desideratum for *prec* is that the distance between regions A and B should be greater than that between A and C just in case the precision similarity between qualities x_A and x_B is less than that between x_A and x_C . This criterion leads to two natural options for defining *prec*: (1) in terms of absolute difference (i.e. the absolute value of the difference in size between A and B), or (2) in terms of absolute ratio (i.e. the ratio in size between A and B, where the numerator is the size of the smaller region). Though both options satisfy the above desideratum, I think the absolute ratio measure is somewhat more attractive. If we were to adopt the absolute difference measure, then we would get the result that very precise qualities could differ only marginally in precision similarity (since the absolute difference between small regions will never be large) while very imprecise qualities will often differ greatly in precision similarity (since the absolute difference between large regions will often be large). But on the contrary, it seems plausible that there can be as much precision dissimilarity between precise qualities as between imprecise qualities.

The formal definition of *prec* is straightforward. Let $\min(A, B)$ be the size of the smaller region between A and B and $\max(A, B)$ be the size of the larger region between A and B. Then (assuming that the absolute ratio option is correct):

$$prec(A, B) = \frac{\min(A, B)}{\max(A, B)}$$

And here is the diagram for *prec*:

$$prec \left[\left(\text{A} \right) , \left(\text{B} \right) \right] = \frac{\min \left[\left(\text{A} \right) , \left(\text{B} \right) \right]}{\max \left[\left(\text{A} \right) , \left(\text{B} \right) \right]}$$

FIGURE 6: The precision similarity metric.

With *prec*, we have a solution to the problem from earlier. The situation involved a precise blue quality B^9 , a precise green quality G^9 , and an imprecise green quality G^5 . We noted that *qual* does not capture the precision similarity between the

precise blue experience and the precise green experience. But *prec* is designed to capture this second dimension of similarity (without encroaching on the qualitative similarity captured by *qual*).³⁹ The result is expressed formally below:

$$\begin{aligned} qual(B^9, G^9) &= qual(B^9, G^5) \\ prec(B^9, G^9) &> prec(B^9, G^5) \end{aligned}$$

It is worth briefly addressing a technical point about the significance of *prec* for the structure of precision. Since the formulation of *prec* appeals to absolute ratio, it may be tempting to infer that precision itself has ratio structure, meaning we can make sense of ratios (rather than merely orderings or differences) between precision values. However, it is possible to accept the absolute ratio formula for *prec* while denying that precision itself has ratio structure. On my proposal, the precision similarity between x_A and x_B is greater than that between x_A and x_C just in case the absolute ratio of A to B is greater than that of A to C. The appeal to ratios occurs only on the formal side of the biconditional: there is no invocation of ratios between precision values. In other words, while *prec* appeals to ratios between region sizes to model degrees of precision similarity, it does not require (though is compatible with) ascribing ratio structure to precision itself. As an analogous point, consider how even though regions are built out of points, imprecise qualities need not themselves be thought of as being built out of precise qualities.

The Regional Framework

On the way I have developed the regional framework, there are distinct metrics for qualitative similarity versus precision similarity. But some may wonder whether there is an overall phenomenal similarity metric that captures phenomenal similarity simpliciter. Speaking for myself, I am skeptical that there is an objective fact of the matter about how to compare the kinds of phenomenal similarity tracked by *qual* and *prec*. But those who think otherwise could always develop a more general metric that captures both qualitative and precision similarity. In fact, an obvious approach would be to sum the outputs of *qual* and *prec*, with a weighting to scale

³⁹ Strictly speaking, *prec* is a pseudometric (rather than a metric) since any metric m must satisfy the condition that $\mu(A, B) = 0$ just in case $A = B$.

their values relative to each other. However, even if there is an objective metric for overall phenomenal similarity, it remains plausible that *qual* and *prec* track two natural kinds of phenomenal similarity.

The arguments that I have made in this section have appealed to largely theoretical considerations. But some may wonder whether it even makes sense to develop a formal framework for modeling mental qualities before empirically investigating those mental qualities. It is true that in order to determine how particular mental qualities map to particular regions within a particular model, we must empirically investigate the similarity relations between those mental qualities. But in order to construct a model using that empirical data, we need a general framework for mapping collections of data to formal structures. In other words, empirical investigation of precision requires having a theoretical framework for interpreting those empirical results, and developing that theoretical framework requires the kinds of arguments I have made in this paper.

§3 | Applications

In this final section, I explain how the regional framework sheds light on the connection between precision and discriminatory grain, the phenomenal sorites, the empirical investigation of mental qualities, and the view that perceptual experiences have probabilistic contents. These discussions will be brief, but they will still illustrate some of the power and potential of the regional framework.

Discriminatory Grain

In the quality-space literature, it is standardly taken for granted that we can investigate mental qualities via their functional roles. Consider how judgments of similarity or distinctness between perceptible qualities are taken to be evidence of phenomenal similarity or phenomenal distinctness between the corresponding mental qualities. Given this, it is natural to ask whether there is a distinctive functional role associated with precision. This section explains and explores the following conjecture: precision correlates with discriminatory grain.⁴⁰

⁴⁰ See Hellie [2005] and Pelling [2008] for similar philosophical analyses of discriminability and imprecise experiences. This section may be thought of as building on this prior work by providing a formal treatment of the issues using the regional framework.

As a first pass, think of *discriminatory grain* as the number of discriminations a subject can make over a set of physical objects using a particular perceptual capacity.⁴¹ If Achilles can make 36 discriminations over 50 color chips but Zagreus can make only 9 discriminations over the same 50 color chips, then Achilles' color discrimination capacities are finer than Zagreus'. This initial gloss elicits the intuitive connection between precision and discriminatory grain, but it also raises some questions. Suppose that a and b are perceptible qualities, that x_A and x_B are the corresponding mental qualities, and that A and B are the regions for x_A and x_B . How should we think about discriminability when A and B overlap?

Let us say that a is *strongly discriminable* from b just in case A is disjoint from B , and that a is *weakly discriminable* from b just in case A partially overlaps with B . If a and b are strongly discriminable, then the subject can be sure (taking their experience at face value) that a is distinct from b . If a and b are only weakly discriminable, then the subject cannot be sure (solely on the basis of their experience) whether a and b are distinct. Nevertheless, weak discriminability still entails that the way a looks is distinct from the way b looks (since A and B are distinct regions and distinct regions represent distinct qualities). This distinction between strong and weak discriminability is useful for understanding the sense in which perceptual *indiscriminability* is (or is not) transitive. Let us say that a and b are *indiscriminable* (in either sense) just in case they are not discriminable. Then strong indiscriminability is non-transitive while weak indiscriminability is transitive.

Now can use the regional framework to formulate a more rigorous measure of discriminatory grain. The *discriminatory grain* of a perceptual capacity is a function of the size of the maximally large set of disjoint permissible regions associated with that perceptual capacity. Putting it another way, discriminatory grain is a maximization of strong discriminations conditional on a minimization (i.e. zero) of weak discriminations. In the example from earlier, Achilles' color discriminatory capacities involve precise color experiences corresponding to relatively small regions, meaning that there is a relatively large maximal set of disjoint permissible regions

⁴¹ Other factors that may influence discriminatory grain include environment and noise. For simplicity, I assume for the rest of the paper that measures of discriminatory grain are always relativized to optimal environments and factor out noise.

associated with Achilles' color perception capacities. By contrast, Zagreus' color discriminatory capacities involve imprecise color experiences corresponding to relatively large regions, meaning that there is a relatively small maximal set of disjoint permissible regions associated with Zagreus' color perception capacities. To picture this, imagine taking the set of regions associated with a perceptual capacity and "filling up" the corresponding space as tightly as possible, as illustrated below:⁴²




FIGURE 7: A has greater discriminatory grain than B.

In what follows, I will explain how this analysis of the connections between precision and discriminatory grain yields insight into the puzzle of the phenomenal sorites and the empirical investigation of mental qualities.

The Phenomenal Sorites

A *phenomenal sorites* case occurs when a subject cannot discriminate (on the basis of their perceptual experiences) a from b or b from c yet can discriminate a from c .⁴³ A natural hypothesis is that phenomenal sorites cases involve situations where there is partial overlap between A and B, partial overlap between B and C, and non-maximal overlap between A and C. Equivalently, these are situations where a and b are not strongly discriminable, b and c are not strongly discriminable, and a and c are at least weakly discriminable. In these situations, the subject cannot be sure on

⁴² FORMAL DEFINITION: Let x be the set of regions for the mental qualities associated with perceptual capacity x . Let $\mathcal{P}(x)$ be the powerset of x . Then the *discriminatory grain* of $x = \max\{C : C \in \mathcal{P}(x) \text{ and } \forall A, B \in C (A \cap B = \emptyset)\}$.

⁴³ See Fara [2001], Hellie [2005], and Pelling [2008] on phenomenal sorites cases.

the basis of their experience whether a is distinct from b , or whether b is distinct from c , but can be sure on the basis of their experience that a is distinct from c .

In fact, the regional framework not only enables us to identify the general conditions for when a phenomenal sorites might occur, but also provides the basis for developing a probabilistic model of those conditions. In particular, it is natural to think that the probability of a phenomenal sorites occurring depends on the degrees of overlap between the relevant regions. To develop this idea, we need to first define degree of overlap:

$$\text{overlap}(A, B) = \frac{\mu(A \cap B)}{\frac{1}{2}(\mu(A) + \mu(B))}$$

This definition of *overlap* yields a value of 1 just in case A and B are identical and a value of 0 just in case A and B are disjoint. And with *overlap*, we can identify some principles connecting the degrees of overlap between regions to the likelihood of a phenomenal sorites case occurring with the experiences represented by those regions. In particular, let $p(A, B, C)$ be a function that is intended to capture the probability that the subject judges (on the basis of experiences corresponding to regions A, B, and C) that $a = b$ and $b = c$ but $a \neq c$. Now consider the following principles, which are plausible constraints for developing the probability function p :

- (1) the greater $\text{overlap}(A, B)$ and $\text{overlap}(B, C)$, the greater $p(A, B, C)$.
- (2) if either $\text{overlap}(A, B) = 0$ or $\text{overlap}(B, C) = 0$, then $p(A, B, C) = 0$.
- (3) the greater $\text{overlap}(A, C)$, the smaller $p(A, B, C)$.
- (4) if $\text{overlap}(A, C) = 1$, then $p(A, B, C) = 0$.
- (5) $0 \leq p(A, B, C) \leq 1$.

These constraints are all satisfied by the formula below, which uses degree of overlap to model the probability that a set of pairwise discriminatory judgments (of a , b , and c) based off of mental qualities represented by regions A, B, and C yields a phenomenal sorites case:

$$p(A, B, C) \simeq \text{overlap}(A, B) \times \text{overlap}(B, C) \times (1 - \text{overlap}(A, C))$$

The symbol ' \simeq ' is intentionally ambiguous, for there is a question of whether we must also accommodate other factors that influence perceptual judgments, such as basing abilities, epistemic norms, psychophysical noise, and so forth. If it is possible to abstract away from such factors to isolate the relationship between precision and discriminability, then the formula may capture a linear relationship. If such abstractions are unfeasible, or if the relationship between precision and discriminability is non-linear even after abstraction, then the formula may capture only a monotonic relationship. But even in the latter case, the regional framework still gives us a powerful formal tool for understanding the phenomenal sorites (and more generally, the relationship between precision and discriminability).

A noteworthy result is that even the mental qualities captured by standard models are best represented by regions, since phenomenal sorites cases occur even for the maximally fine discriminatory capacities of normal humans. This illustrates why the regional framework is a general framework for modeling all mental qualities, rather than a specialized tool for dealing with a particular kind of mental quality. In light of this, I think the regional framework is better thought of as a successor (rather than merely a supplement) to the standard framework.

Empirical Investigation

How can we use the regional framework to construct models of particular quality-spaces? Suppose we start with a standard model that captures a set of precise mental qualities. The initial step is to convert that standard model into a regional model representing those same qualities. This requires mapping points in the standard model to regions in the regional model such that (1) every point in the former is mapped to a distinct region in the latter, (2) similarity relations are preserved, and (3) the boundaries and dimensions of the space are preserved. The basic procedure for this conversion is relatively straightforward, though there is a question of how to determine when two regions overlap. Let us set that issue aside for now — we will return to it in a moment.

After converting the standard model to a regional model, the challenge is to identify qualities with differing degrees of precision and map them onto regions, with the constraint that distance values (outputted by *qual* and *prec*) correspond to degrees of qualitative and precision similarity (between the relevant mental qualities). To do this, we need methods for identifying the region location, region size,

and region shape for any given quality. Since the methodology used to determine locations of points in standard models generalizes to determining locations of regions in regional models, let us focus on region sizes and region shapes.

Let the *target region* (or quality) be the region (or quality) that we wish to investigate. Previously, we noted that coarser discriminatory capacities are correlated with more imprecise qualities, and more imprecise qualities are represented by larger regions. The greater the overlap between two regions, the greater the likelihood that a subject will fail to discriminate two stimuli perceived using the mental qualities represented by those regions. These connections between precision and discriminatory grain provide a basis for using empirical methods to approximate region sizes and triangulate region shapes.

A simple but limited approach is to partition perceptual capacities into subclasses, where subclasses are individuated by their discriminatory grain. For example, since color discrimination is coarser outside of the center of the visual field, color perception might be divided into subclasses determined by angular distance from the center of the visual field. Given the connections between precision and discriminatory grain, these perceptual capacity subclasses will correspond to equivalence classes of mental qualities, where each equivalence class consists of the set of mental qualities within a given quality-space at a given degree of precision. This provides a way of approximating region sizes, though it still leaves open questions about region shapes.

A more methodical but also more complex approach is to appeal to pairwise comparisons between perceptual stimuli. Suppose we wish to identify the size and shape of a target region A (representing target quality x_a , which is induced by the perception of stimulus a). This approach would require subjects to make pairwise discriminations between stimulus a and a series of other stimuli b_1 – b_n that induce other mental qualities (of the same quality-space). If the subject is often unsure whether a is distinct from the other stimulus b_i , then region A is likely large. If the subject is often sure that a is distinct from b_i , then A is likely small. The less frequently the subject judges that a is distinct from b_i , the greater the expected degree of overlap between A and the region representing b_i .

The most straightforward experimental procedures would involve inducing experiences that instantiate both the target quality (held fixed across trials) and a series of precise qualities (varying across trials). Consider, for example, a procedure

where subjects make perceptual judgments about a color chip they see via foveal vision (corresponding to stimulus a from above) versus a series of color chips they see via peripheral vision (corresponding to stimuli b_1 – b_n from above). If a subject judges the chips within a trial to be the same color, then that is evidence that the regions representing the mental qualities by which the subject perceives those chips at least partially overlap. Since more precise mental qualities are represented by smaller regions, generating these pairwise comparisons with precise qualities enables sharper identifications of the size and shapes of the target regions.

From Regions to Fields

According to the *perceptual confidence* theory, perceptual experiences have probabilistic contents.⁴⁴ I will be neutral on whether the perceptual confidence theory is correct, but I will briefly mention how the regional framework can be naturally extended to accommodate it.

Let us start with the basic idea behind the perceptual confidence theory. Suppose a subject perceives an object via a color experience, and let $p(a)$ be the probability (according to the content of that color quality) that the object has color a . On the perceptual confidence theory, it is possible that $p(\text{blue}) = .2$, $p(\text{teal}) = .6$, and $p(\text{green}) = .2$, even when blue, teal, and green correspond to regions of the same size. Since these probabilistic structures can vary even when the relevant regions are held fixed, the regional framework does not have the formal structure needed to model this kind of situation. The limitation is due to the fact that regions are “flat,” in that they do not assign different weights to different points.

Nevertheless, the perceptual confidence theory can be accommodated by moving to a *field framework*. A *field* on a space of points is an assignment of values to every point in the space.⁴⁵ Equivalently, a field is 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 (say) either \top (if the point is inside the region) or \perp (if the point is outside

⁴⁴ See Morrison [2016] and Munton [2016] for argument in favor of the perceptual confidence theory. See Denison [2017] and Nanay [forthcoming] for arguments against.

⁴⁵ This notion of ‘field’ should be distinguished from the algebraic notion, where a *field* is an algebraic structure that permits addition, subtraction, multiplication, and division.

the region). But fields can assign a broader range of values (where a natural constraint for the perceptual confidence theory is that the integral of the field must equal 1). Under the field framework, the precision structure of a mental quality would be represented by the structure of the field for that quality. For example, the color quality mentioned above would be represented by a field that assigns higher values in the teal region of the quality-space than in the blue and green regions.⁴⁶

As with the regional framework, there are novel challenges in identifying which formal constraints on fields are most plausible, in developing the right similarity metrics on fields, and in constructing particular models. Addressing these challenges is beyond the scope of this paper. But if the perceptual confidence theory is correct, then the evolution of the quality-space model framework progresses from points to regions to fields.⁴⁷

Conclusion

A regional model can be specified via a tuple: $\langle \mathcal{S}, d, \mathcal{R}, \mu, qual, prec \rangle$, where \mathcal{S} is a set of points, d is the point-distance metric, \mathcal{R} is a set of regions (specifying which regions correspond to possible mental qualities), μ is a measure on \mathcal{S} (specifying degrees of precision), $qual$ is a metric on regions (capturing qualitative similarity), and $prec$ is a metric on regions (capturing precision similarity). The structure of the regional framework makes it more powerful and flexible than the standard framework, and enables us to formulate novel hypotheses about the space and structure of mental qualities, formally differentiate two dimensions of phenomenal similarity, identify the connections between precision and discriminatory grain, generate a probabilistic model of the phenomenal sorites, and deploy a new theoretical tool for the empirical investigation of consciousness.

⁴⁶ I appeal to fields rather than probability distributions because the relevant probability distributions would range over perceptible qualities rather than mental qualities. Given this, I think it is better to think of mental qualities themselves as modeled by fields, even if their contents can be modeled by probability distributions over perceptible qualities.

⁴⁷ A more speculative suggestion: the field framework might also be used to model *phenomenal vivacity* (consider the phenomenal contrast between perceptual and imaginative experiences), where more vivid mental qualities would correspond to fields with greater integrals.

A striking consequence is that mental quality-spaces are not isomorphic to perceptible quality-spaces. Whereas both mental qualities and perceptible qualities have similarity structure, only mental qualities have precision structure. At first, this conclusion may strike some as puzzling. It may be tempting to think that perceptible qualities *must* have precision structure, since “perceptible precision” may simply be defined as whatever property of perceptible qualities systematically corresponds to phenomenal precision. But the problem with this line of reasoning is that the perceptible correlate of phenomenal precision is determinability: imprecise mental qualities correspond to more determinable perceptible qualities. However, mental qualities *also* have determinability structure, and (as we saw in §1) determinability and precision are mutually dissociable. Therefore, the fact that there is a perceptible correlate of precision does not vindicate the isomorphism thesis. And as far as I can see, there are no other credible candidates for anything playing the role of perceptible precision. This means that mental qualities have strictly more structure than perceptible qualities. And to capture that structure, we need to transition from the standard framework to the regional framework.

For those undertaking future investigations of mental qualities, there is good reason to adopt the regional framework from the outset. By doing so, one attains increased power and flexibility with little added cost. Though I focused earlier on how a standard model can be transformed into a regional model, there is no methodological advantage to starting with a standard model and subsequently converting it. In fact, the very procedures deployed to construct standard models can likewise be deployed to construct regional models. And as we saw earlier, even the mental qualities captured by standard models are often better captured by regions, since even those mental qualities give rise to phenomenal sorites cases.

A more general goal of this paper has been to exhibit the prospects for formal phenomenology, or the application of formal tools to the study of conscious experiences. By formally modeling conscious experiences, we not only sharpen our understanding of how conscious experiences are structured, but also progress our understanding of consciousness without needing to resolve long-standing theoretical disputes about the mind-body problem, the nature of perception, or the physical correlates of consciousness. In my view, this kind of project is one of the most promising ways of moving consciousness research towards a systematic science.

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). The Puzzle of Perceptual Precision. *Open 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.
- Churchland, Paul (2005). Chimerical colors: Some phenomenological predictions from cognitive neuroscience. *Philosophical Psychology* 18 (5):527-560
- Churchland, Paul (2007). On the reality (and diversity) of objective colors: How color - qualia space is a map of reflectance - profile space. *Philosophy of Science* 74 (2):119-149.
- Clark, Austen (2000). *A Theory of Sentience*. Oxford University Press.
- Coates, Paul & Coleman, Sam (eds.) (2015). *Phenomenal Qualities: Sense, Perception, and Consciousness*. Oxford University Press UK.

-
- 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.
- Debats, N. B., Kingma, I., Beek, P. J., & Smeets, J. B. (2012). Moving the weber fraction: the perceptual precision for moment of inertia increases with exploration force. *PloS one*, 7(9), e42941. doi:10.1371/journal.pone.0042941
- 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.
- Fara, Delia Graff (2001). Phenomenal continua and the sorites. *Mind* 110 (440):905-935.
- Fazekas, Peter & Overgaard, Morten (2018). A Multi-Factor Account of Degrees of Awareness. *Cognitive Science* 42 (6):1833-1859.
- 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.
- Hansen, Thorsten, Pracejus, Lars, & Gegenfurtner, Karl R. (2009). Color perception in the intermediate periphery of the visual field. *Journal of Vision* 2009;9(4):26. doi: <https://doi.org/10.1167/9.4.26>.
- 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.
- Hurevicz, W. & Wallman, G. (1948) *Dimension theory*, Princeton University Press.
- 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.
- Kleiner, Johannes. 2020. Mathematical Models of Consciousness. Forthcoming in *Entropy*. arXiv preprint [arXiv:1907.03223](https://arxiv.org/abs/1907.03223).
- Kliniewicz 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., & Kruschke, J. K. (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.
- Morrison, John (2016). Perceptual Confidence. *Analytic Philosophy* 57 (1):15-48.
- Munton, Jessie (2016). Visual Confidences and Direct Perceptual Justification. *Philosophical Topics* 44 (2):301-326.
- Murray, David J. (1993). A perspective for viewing the history of psychophysics. *Behavioral and Brain Sciences* 16 (1):115.
- Nanay, Bence (forthcoming). Perceiving Indeterminately. *Thought: A Journal of Philosophy*.
- Palmer, Stephen E. (1999). Color, consciousness, and the isomorphism constraint. *Behavioral and Brain Sciences* 22 (6):923-943.
- Pelling, C. (2008). Exactness, Inexactness, and the Non-Transitivity of Perceptual Indiscernibility. *Synthese*, 164(2), 289–312. JSTOR.
- Prentner, Robert (2019). Consciousness and topologically structured phenomenal spaces. *Consciousness and Cognition* 70:25-38.
- Renner, 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, David (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.
- Stazicker, James (2011). Attention, Visual Consciousness and Indeterminability. *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 Indeterminability. *Inquiry: An Interdisciplinary Journal of Philosophy* 56 (4):359-385.
- Yoshimi, Jeffrey. (2011). Phenomenology and Connectionism. *Frontiers in psychology*. 2. 288. 10.3389/fpsyg.2011.00288.

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