PhiMiSci
Philosophy and the Mind Sciences | Vol. 6 | 2025

Objective foundations for the study of mental qualities

David Rosenthal[1,*]

Special Issue: *Structuralism in the Science of Consciousness*, edited by Sascha Benjamin Fink & Andrew Lee

Abstract

Quality spaces promise to represent mental qualities objectively. That objectivity is compromised, however, if quality spaces are constructed by subjective introspective access, which is impressionistic. But mental qualities also have a robust and objective connection to perceptual discrimination. So quality spaces can be constructed in a fully objective way by appeal to their role in perceiving. In addition, the subjective appearances of mental qualities also bear a constitutive relation to perceptual role, since each subjective appearance consists in its subjectively appearing as it does when one perceives some specific type of object. So quality spaces constructed from perceptual role can also be used to characterize those subjective appearances. And since subjective appearance depends on perceptual role, the tie mental qualities have to perceptual role is more fundamental than that with subjective appearance. Quality spaces are useful in the first instance for objectively representing discriminable stimuli, and derivatively for objectively representing the corresponding mental qualities in respect of both their perceptual roles and their subjective appearances.

Keywords: Consciousness • Just noticeable differences • Mental qualities • Perceptual discrimination • Quality space • What it is like

[1] Graduate Center, City University of New York
* Primary contact: davidrosenthal1@gmail.com

Rosenthal, D. (2025). Objective foundations for the study of mental qualities. *Philosophy and the Mind Sciences, 6.* https://doi.org/10.33735/phimisci.2025.11526

1 Why quality spaces

Conscious mental qualities are often said to pose the most intractable problem for understanding the mind. One pivotal aspect of this apparent difficulty is that it seems to many difficult, if possible at all, to give any informative description of conscious mental qualities. We know about conscious mental qualities from the first-person access we have to them, but that first-person access does not by itself seem to enable informative descriptions.

The sense that it is difficult to give an informative description of conscious mental qualities is forcefully captured by Thomas Nagel's (1974; also in 2024, pp. 1–29 and pp. 56–64) claim that no objective account of conscious mental qualities is possible. Any attempt to describe conscious mental qualities objectively would, according to Nagel, inevitably fail to capture their irreducibly subjective nature.

In that spirit Ned Block urges that the best, and perhaps only, reply to the question of what conscious mental qualities are is Louis Armstrong's famous quip about jazz: "If you gotta ask, you ain't never gonna get to know" (1978, p. 281). Block has more recently put the point by urging that "[t]he best you can do is use words to point to a phenomenon that the reader has to experience from the first person point of view" (2015, p. 47). Even those who endorse Frank Jackson's (1986) contention that one gains new factual knowledge on first consciously seeing red are decidedly reticent about what such new factual knowledge could conceivably consist in (Rosenthal, 2019).

A major challenge to this claim that conscious mental qualities cannot be informatively or objectively described rests on the observation that we can give useful descriptive conscious accounts mental qualities by appeal to the relations each quality bears to others. Describing mental qualities seems difficult only if we must do so atomically, one quality at a time. Relational properties enable richly informative descriptions.

Characterizing mental qualities relationally dispels many alleged conundra. For example, it will seem different or even impossible to describe a mental quality to somebody who has never had that qualitative experience only if one must do so atomistically. It is straightforward to describe such mental qualities to others using comparisons with other qualitative experiences.

Appeals to such relations points to the use of quality spaces to represent mental qualities. If we can characterize mental qualities by the relations each bears to others, we can construct a space that captures those relations. Each mental quality is then individuated by its unique location in a space of qualities special to the sensory modality in question, and each quality is informatively described by appeal to that unique location. Quality spaces offer hope that we can, after all, describe conscious mental qualities objectively and informatively. Indeed, it is not obvious what other strategy there might be to generate such an account.

Those who hold that conscious mental qualities resist objective, informative description might insist that such descriptions must be nonrelational, characterizing each mental qualities on its own in respect of its intrinsic nature. But there is no reason to hold that an objective account must be atomic in this way. Many things are best understood objectively in relational terms.

The idea of using quality spaces in this way by no means new. It occurs in the work of Wilfrid Sellars (e.g., 1956) and of W. V. Quine (2013, §17, pp. 82ff; 1969, pp. 123–128), as well as in scientific investigations (e.g., Schanda, 2007; Shepard, 1982). And it has been developed elsewhere in detail by me (1991, 2005, 2010, 2024) and by Austen Cark (1993, 2000).

But there has recently been a striking surge of interest in the quality-space strategy for giving an informative account of conscious mental qualities (e.g., Fink et al., 2021; Kleiner, 2024; Kleiner & Ludwig, 2024; Kob, 2023; Lee, 2021, 2024; Lyre, 2022; Malach, 2021; Prentner, 2019; Tallon-Baudry, 2022; Tsuchiya et al., 2022). Some of these proposals appeal in part to neurological findings and others make use of abstract mathematical structures. But they all invoke quality spaces to describe and explain conscious mental qualities.

These recent proposals underwrite richly informative descriptions of mental qualities. But there is an issue about whether they are fully objective. These recent proposals all construct quality spaces using introspective judgments of similarity and difference among conscious qualities. This

is also true of Clark's account (Clark, 1993; esp. Clark, 2000, p. 18), and arguably of Sellars' account as well (see Rosenthal, 2016).

But there are problems about the reliability of introspective judgments. Such judgments are less accurate than straightforward perceptual judgments about the perceptible properties that are presented to us. And in assessing the accuracy of introspective judgments there is virtually no direct and independent control about what mental quality an individual has and whether it is the same on different occasions.

These concerns, to be discussed in detail in Section 4, cast doubt on whether introspective judgments about similarities and differences among mental qualities can be fully objective. And if they are not, then quality spaces constructed by relying on them cannot be either. Those quality spaces would then constitute at best a limited corrective to the contention that mental qualities cannot be objectively and informatively described.

On Nagel's account, objective facts are those one can grasp independently of one's point of view, whereas subjective facts can be grasped only from some particular point of view. Nagel wants to avoid concerns about "the alleged privacy of experience to its possessor," and to do so he invokes degrees of similarity among points of view (Nagel, 1974, pp. 441–442). But it is unclear that the distinction construed Nagel's way can survive degrees of similarity among points of view (Rosenthal, 1983). Still, adopting Nagel's distinction between subjective and objective, Andrew Lee (2024) has convincingly argued that the purely structural relations that quality spaces fix among mental qualities are objective, since they can be understood independently of one's subjective point of view.

But objectivity of that type does not affect the concern just raised about whether introspective judgments can generate objective quality spaces. If those judgments are not reliable about the relations that hold among mental qualities, quality spaces based on such judgments will fail objectively to represent the facts about relations among mental qualities. And for present purposes we can understand objectivity simply as maximum intersubjective reliability.

The tendency to rely on introspective judgments in constructing quality spaces should not be surprising. The prevailing fashion in thinking about mental qualities is that we must understand them by way of first-person access. Any source other than first-person access would simply provide adjunct information about mental qualities, and not describe them as they truly are. In particular, any other source would fail to capture what it is like for one to have the conscious experiences that exhibit mental qualities.

Since sources other than first-person access could provide objective information about mental qualities, dismissing such other sources in effect simply channels the insistence that mental qualities cannot be objectively described. And quality spaces constructed from introspective judgments do nothing but arrange input from first-person access in a convenient format. So if quality spaces are to offer a genuine corrective to the denial that mental qualities can be objectively described, we need some way to construct them that does not rest on first-person access and introspective judgments.

And an alternative approach is readily available. Reliance on first-person access is based on the close tie that holds between mental qualities and by what it is like for one to be in the relevant qualitative states. But mental qualities exhibit another tie that is at least as robust and revealing. Mental qualities figure pivotally in perceptual functioning. We detect physical objects and events by the mental qualities they elicit in perceiving, and we discriminate the perceptible properties of those objects and events by differences among mental qualities that occur in perception. Perceptual detection and discrimination would be altogether impossible without mental qualities that correspond in appropriate ways to the physical properties we perceive.

This constitutive tie between mental qualities and perceptual functioning emerges vividly in the way we describe the mental qualities that figure in perceiving things. We describe the qualitative character of a perceptual state by appeal to the type of physical property that state enables us to discern. The mental quality that characterizes a perception of red, for example, is that mental property which enables us to pick out physically red objects. Similarly for all other perceptual mental qualities.

And the tie that mental qualities have to perceptual functioning is arguably even stronger and more fundamental than the tie mental qualities

have to consciousness. For one thing, perceptual detection and discrimination sometimes occur independently of consciousness, as revealed by priming effects and forced-choice testing. Such unconscious perceiving is not limited to laboratory experimentation; one can in everyday situations perceptually react to things one seemed subjectively not to perceive.

There are those who deny that perceiving, properly so called, can occur without being conscious (e.g., Phillips, 2018). Still, there is at least an issue about whether mental qualities always occur consciously. But there can be no serious debate about whether mental qualities enable perceptual functioning. The tie mental qualities have to perceptual functioning is arguably at least as solid and robust as the tie to what it is like.

So instead of relying on introspective judgments about conscious mental qualities to construct quality spaces, we can rely on the role mental qualities play in the perceptual detection and discrimination of physical properties. Because we have control over what perceptible stimulus properties are presented to an individual, individuating mental qualities by their perceptual role is readily testable. Constructing quality spaces from perceptual role would have a foundation that is incontrovertibly objective.

It may seem, however, that this type of objective foundation comes at an unacceptably high price. The issue about objectivity is whether we can give a completely objective account of the nature of conscious mental qualities. And that requires giving an objective account of mental qualities in respect of their conscious nature.

Fixing mental qualities by appeal to their perceptual role, however, achieves objectivity by setting consciousness aside. We need not appeal to the way mental qualities occur consciously to describe fully the way each mental quality enables perceptual detection and discrimination. All that is needed is an individual's response to being presented with relevant stimuli. And assuming that perceptual detection and discrimination can indeed occur unconsciously, the mental qualities that enable such unconscious perceptual functioning will themselves occur independently of consciousness.

So it may seem that quality spaces constructed from perceptual role cannot shed light on mental qualities in respect of their conscious subjectivity. And if so Nagel may be right that objectivity in describing conscious mental qualities inevitably drives out their essential subjectivity.

But this concern is groundless. Perceptual role enables an objective account of mental qualities that does not appeal to their conscious subjectivity. But that does not show we cannot build an account of that conscious subjectivity on top of the objective descriptions that perceptual role provides. Perceptual role is independent of conscious subjectivity. But perceptual role can nonetheless be the basis for an informative account of mental qualities in respect of their conscious subjectivity. One would think otherwise only if one held that we can accommodate conscious subjectivity only if we take it to be fundamental. But there is no serious reason for that contention.

Indeed, as noted a moment ago, the way we think about mental qualities in respect of their conscious subjectivity appeals to perceptual role. We describe what it is like for one to be in states with particular mental qualities by reference to the perceptual role of those mental qualities. The subjective appearance of red consists in what it is like is to see something that is red. We can fix perceptual role without appeal to consciousness, but perceptual role always figures in the way we think about perceptual consciousness.

Still, this commonsense observation is not enough. An extra explanatory step is needed to go from perceptual role to the way mental qualities appear subjectively in our stream of consciousness. If that step is itself objective, we will have an objective account of conscious mental qualities, in respect of both their perceptual role and the way they subjectively appear in consciousness. Developing that step and showing that it is objective will be the task of Sections 5–7, below.

Section 2 will start the construction of fully objective quality spaces by discussing the most fine-grained perceptual discrimination, so-called just noticeable differences between perceptible stimulus properties. Section 3, then, will show how such just noticeable differences can be used to construct such quality spaces. Section 4 will consider the disadvantages of using subjective information to construct quality spaces. Section 5 will discuss how we understand mental qualities in respect of their subjective appearances, and Section 6 will show how quality spaces constructed from

perceptual role can determine those subjective appearances. Section 7 concludes by disarming a major consideration that has led many to insist that the subjective appearances must be fundamental to any discussion of mental qualities.

2 Just noticeable differences

Quality spaces should assign each type of mental quality a distinct relative location. And we want this assignment to be objective and testable. The best procedure will rely on the perceptual role of mental qualities; indeed, that is likely the only procedure that is both objective and testable.

Mental qualities enable us to detect the presence of a stimulus and to discriminate each type of stimulus from others. But detection by itself provides little if any information about the type of mental quality. For that we must rely on perceptual discrimination.

Discriminating two stimuli requires perceiving the stimuli by way of distinct types of mental quality, one for each of the stimuli. One could not discriminate the stimuli if they elicited mental qualities of the same type. So determining what mental qualities figure in an individual's perceptual repertoire requires determining what stimuli an individual is able to discriminate.

We want, moreover, to determine the most fine-grained differences among the mental qualities an individual can have. So we must establish an individual's ability to discriminate stimuli in the most fine-grained way. We can do that by testing for just noticeable differences (JNDs). Two stimuli of distinct physical types are JND for an individual at a time if the individual can tell that they are different, though if the stimuli were physically any closer the individual would regard them as the same.

For experimental purposes, JNDs are typically established statistically. A participant counts as discriminating two stimuli if the participant indicates that they are different on some percentage of trials, sometime 50%, sometimes higher (e.g., Torgerson, 1958, pp. 132–133; see Kingdom & Prins, 2016, p. 30). Participants can judge consciously that the stimuli are distinct. But the experimental procedure can also rely on participants' perceiving the stimuli unconsciously (Torgerson, 1958, 13ff.) Any indication that stimuli are distinct, such as priming or forced-choice, works. More on unconscious discrimination toward the end of Section 3.

To establish all the JNDs that pertain to discriminable physical stimuli of a particular type, say, the colors, one could start with any stimulus and work outward with adjacent JNDs by varying the stimuli in physically minimal ways. Armed with these successive JNDs, one can construct a quality space that would represent all the discriminable stimuli by their relative location in that space. This works for every type of perceptible stimulus, such as color, pitch, timbre, odor (Young et al., 2014; cf. Martina, 2023a, 2023b), tactile pressure and texture, and taste.

Stimuli can often be physically varied in several ways to generate new JNDs, and that will sometimes require a space with more than two dimensions. As one works outward one JNDs at a time, it will eventually happen that varying a stimulus in some particular way does not result in a new JND stimulus, but rather in a stimulus that is simply not perceptible. That will determine a boundary of that space.

There are various concerns one might raise about this appeal to JNDs, though none of them undermines their use in constructing quality spaces. For one thing, no two individuals are exactly the same in discriminative ability. And the discriminative ability of each individual will, moreover, vary from one time to another, depending on many physical and psychological circumstances. One can average over individuals and over differences within a single individual. But testing for JNDs is at bottom a matter of testing an individual at a time.

This is all as it should be. We want to use the role mental qualities play in perceptual discrimination to fix those mental qualities in a fully objective way. Since each individual will objectively differ from others in discriminative ability and vary from one time to another, we can expect that individuals will vary slightly in their repertoire of mental qualities, both among themselves and from one time to another. Any objective procedure for fixing mental qualities must reflect these variations.

There are also technical issues that arise for JNDs. Having defined JNDs in terms of whether two samples match, Nelson Goodman shows that the relation of being JND will not be symmetrical; there will be cases in which one sample is JND from another but the second is not JND from the first (1977, p. 226). Indeed, the statistical nature of JNDs should by itself lead us to expect empirical asymmetries even apart from Goodman's argument. One can also expect occasional empirical failure of transitivity as well (Morrison, 2015).

But none of this detracts from the objective construction of a quality space using JNDs between stimuli. There can be empirical failures of symmetry and transitivity in any extremely fine-grained measurement of natural phenomena, and there are standard ways to adjust for such irregularities.

There is also a historical concern about JNDs. The notion of a JND originated in the 19th century with the hypothesis of Ernst Heinrich Weber and Gustav Theodor Fechner that the degree of physical difference between discriminable stimuli is a constant. Weber's law, so-called, held that it is a constant fraction; on Fechner's law the relation is logarithmic. (For a concise summary see Gescheider, 1997, ch. 1; for more detail see Laming, 1997, chs. 1-3). Construed in this way, one might well argue that there simply are no JNDs (e.g., Sanford & Halberda, 2023).

But the notion of a JND that is operative for the construction of quality spaces does not involve any such claims about the size of physical differences between JND stimuli. All that matters is that an individual indicates on a statistical basis two stimuli are distinct, and that if they were physically any closer the individual would indicate that they are indistinguishable. This construal of JNDs is reasonably standard in the literature, as reflected in the foregoing references.

We can determine JNDs for differences between any physical stimuli perceptible by some sensory modality. As already noted, that includes colors, sounds, odors, tastes, and tactile stimuli. We can also establish JNDs between physical stimuli that characteristically result in various bodily sensations, such as pains and various forms of paresthesia.

But we also sense spatial properties, such as size, shape, and spatial location, as well as temporal properties, such as temporal succession and duration. And we sense those spatial and temporal properties in connection with each distinct sensory modality,

Consider vision. Whenever we sense a particular color, there is a spatial boundary between that color and others that surround it. A color stimulus that is entirely uniform and lacks spatial boundaries, known as a Ganzfeld, typically produces perceptual disorientation and visual hallucinations (e.g., Wackermann et al., 2008). The boundaries we see between distinct colors constitute spatial properties accessible by vision, such as visible size, shape, and spatial location. Visible size and shape are straightforward; visible location must be registered relative to the boundaries of one's current visual field.

Other sensory modalities also register spatial properties, though often in a less fine-grained way than vision. Audition, olfaction (e.g., Dikeçligil et al., 2023), and gustation are all reasonably good with spatial location, but register little about size and shape, whereas tactition registers far more about size and shape than about spatial location. But each modality delivers information about these spatial properties by way of boundaries that are sensed between the relevant content properties, such as color, sound, odor, taste, and pressure, and in the case of spatial location relative to the relevant sensory field.

Each modality registers spatial properties by boundaries among its proprietary content properties. There is no way to sense spatial properties independently of some specific sensory modality. So the spatial inputs of each modality must be calibrated with the spatial input of others, a crucial learning process in very early life. This is evident from rare cases in which we seem to see something as having a different size or location from the size or location we sense it as having by touch.

Once those calibration are established, they appear seamless and fluid. And that might make it that spatial properties are sensed in some way that is independent of the individual modalities. But that is an illusion. There is no sensory access to spatial properties except by way of differences in the content properties sensed by each modality. So JNDs pertaining to spatial

stimuli must be independently established for each modality, just as we establish JNDs between colors, sounds, and other content properties.

These considerations hold for temporal duration as well. We have access to temporal succession and duration only by sensing changes in content properties accessible to each modality. So each modality delivers temporal information independently of the others, which we must learn to calibrate in very early life. And we can determine JNDs between nearly indistinguishable durations as we do with the content properties special to each modality.

Quality spaces for each type of stimulus property can be constructed using JNDs. Each location in such a quality space will represent a type of physical stimulus that is JND from those of its physically closest neighbors. And because JNDs are tested objectively, the resulting spaces are also fully objective.

3 Objectivity for mental qualities

The JNDs appealed to here are between physical stimulus properties. So the resulting quality spaces represent those stimulus properties that are discriminable for a particular individual. But as already noted, it is straightforward to extrapolate from quality spaces of discriminable stimuli to quality spaces that represent the mental qualities that enable those discriminations.

For two physical stimuli to be discriminable they must elicit different mental qualities. So once we can map all the stimulus properties an individual can discriminate by way of some sensory modality, we thereby have a map of that individual's mental qualities pertaining to that modality.

And that map represents mental qualities in respect of their most fine-grained differences, since JNDs provide the most fine-grained differences between discriminable stimuli. Differences in the stimuli mental qualities can discriminate will capture any differences subjective awareness could reveal. So the quality space of discriminable stimulus properties will also exhaustively represent all the mental qualities an individual is capable of having within a particular family of mental qualities (Rosenthal, 1991; Rosenthal, 2005, ch. 7; Rosenthal, 2010, 2022, 2024). And since the space of JND discriminations was generated objectively, the space of corresponding mental qualities is itself objective (cf. Pauen, 2017).

Each mental quality enables one to discriminate a perceptible stimulus property from its barely discriminable neighbors. So we can also see each mental quality as representing that stimulus property (Berger, 2015; Rosenthal, 2005, ch. 7, §V; Rosenthal, 2022, §14.2.2). Such representation relies on corresponding locations in respective quality spaces, and so operates differently from conceptual representation. But perceiving typically involves Interactions between conceptual representation and representation by mental qualities.

Many who endorse the use of quality spaces in giving an account of mental qualities appeal not the role of mental qualities in discriminating physical stimuli, but to subjective introspective judgments of similarity and difference between conscious mental qualities themselves. Thus Clark urges that quality spaces should be constructed in terms of "the relations of qualitative similarity among the occupants, but mention no stimuli" (2000, p. 18; cf. 2000, cf. p. 13).

But relying on introspective judgments to construct quality spaces has significant disadvantages, which undermine the objectivity of the resulting spaces. These concerns will be addressed in Section 4. But it is worth stressing here that on the current proposal, quality spaces represent similarity and difference relations between mental qualities only derivatively. The degree of similarity between two mental qualities is a function of how many JNDs there are between the stimuli those mental qualities enable one to discriminate. Even though the difference between JND stimuli is not constant, there is no measure of similarity or difference other than number of JNDs.

If one instead takes similarity of mental qualities as fundamental for constructing quality spaces, an unfounded holism about mental qualities results, on which each mental quality is individuated by its similarity relations to every other mental quality in the space. Reliance on stimulus JNDs also individuates mental qualities relationally, but the individuating relations are then highly local. Each mental quality is individuated as that

mental property which enables the discrimination of a particular stimulus from its barely discriminable neighbors.

The methodology of stimulus JNDs enables the construction of spaces of mental qualities that figure in each sensory modality. But it is sometimes argued that the modalities themselves cannot be individuated except by appeal, at least in part, to subjective phenomenology (e.g., Grice, 1962; Macpherson, 2011). If so, that might threaten the full objectivity of quality spaces constructed exclusively from stimulus JNDs.

But the JND methodology actually provides a fully objective way to individuate the sensory modalities independently of phenomenology. Two physical stimuli are JND only if they would be indistinguishable if they were physically any closer. But that cannot happen with stimuli accessible by distinct modalities. Such stimuli would always be distinguishable even if they were physically closer. Since no physical stimulus accessible by one modality can be JND from a physical stimulus accessible by another, JNDs between physical stimuli provide a way to distinguish the sensory modalities in a fully objective way (Rosenthal, 2015).

JNDs determine sets of stimuli that differ physically in ways so small that an individual cannot discriminate them. And a particular stimulus will sometimes belong to two JND sets by being is indiscriminable from the other members of two sets whose other members are mutually discriminable. In this way, JND sets will sometimes overlap. And a quality space of JND stimuli must represent all that. So such a quality space must represent sets of JND stimuli by very tiny regions of the space. To do so it will be useful to adapt the elegant machinery Lee (2021) has developed for representing regions in a quality space.

Lee uses such regions to represent the way mental qualities differ in respect of precision. The mental quality that figures in seeing a shade of color parafoveally, for example, will be less precise than the mental quality that results from seeing that same shade foveally. Variations in attention can also produce differences in such precision. A quality space of mental qualities, then, can represent a more precise quality by a smaller region. Quality spaces that represent mental qualities by points, as with Clark (1993) and others, cannot capture such differences.

Mental qualities differ in respect of both perceptual role and subjective appearance. And both perceptual role and subjective appearance can be more or less fine-grained. So the precision of mental qualities can vary both in respect of perceptual role and in respect of subjective appearance. Those two types of precision will often vary together; mental qualities that are less precise in respect of subjective appearance will typically enable less fine-grained discriminations (see Lee, 2021, §4.3). Still, the two ways of varying need not always match.

Precision in respect of perceptual role is usefully captured by the number of JND sets of stimuli that a mental quality enables one to discriminate. A mental quality that is maximally precise in respect of perceptual role will enable the discrimination of a single JND. A mental quality less precise in perceptual role will enable discrimination only of a group of JNDs. Degrees of precision in respect of the subjective appearances will be discussed in Section 6, along with other issue about the subjective appearances.

But even when mental qualities are maximally precise in both ways, JNDs consist of sets of physically distinct but indiscriminable stimuli. And regions of physical stimuli are needed to capture those sets of indiscriminable but distinct stimuli. So regions in the quality space of JND stimuli will represent the physically distinct stimuli in each JND set. And those regions should be preserved in extrapolating from the space of JND stimuli to the space of mental qualities, so as not to lose information about numbers and overlap of stimulus JNDs. But one can also construct dedicated quality spaces that use regions to represent how precise mental qualities are in particular circumstances for discriminating stimuli.

A stimulus typically elicits a mental quality whose relative location in its quality space corresponds to the relative location of the stimulus in its quality space. But not always. In the memory color effect, for example, a gray banana stimulus will elicit a yellow mental quality (Bartleson, 1960; Hansen et al., 2006). And the color mental qualities elicited in simultaneous color contrast (Soranzo, 2016) and in assimilation (Gori, 2016) also differ in hue from the physical color stimuli. None of that poses a problem for the current proposal, whose goal is not to specify why mental quality will occur in particular perceptual circumstances, but to construct quality spaces that

exhaust an individual's repertoire of mental qualities. And the JND testing used to construct quality spaces of JND stimuli can readily avoid those special perceptual effects.

Metamerism is also not a problem. Metamers are physically distinct stimuli that elicit the same mental qualities. So metameric stimuli will be JND from neighboring stimuli in exactly the same way as standard stimuli.

As noted in Section 2, there will be quality spaces for JNDs between the spatial and temporal stimuli of each of the sensory modalities. Each modality will have its distinct quality space for each of the spatial properties of size, shape, and relative location. So extrapolating from spaces of stimulus JNDs to spaces of mental qualities will determine for each modality mental qualities of size, shape, and relative location (Meehan, 2007). And since there will be a quality space of JND stimuli of temporal duration for each modality, such extrapolation will determine for each modality a space of mental qualities of duration (Klincewicz, 2011).

Clark (2000) has argued that there are no mental qualities for spatial or temporal properties (§2.4), and that spatial properties in particular must be handled altogether independently of mental qualities, by appeal to what he calls feature placing (§2.6), borrowing the term from P. F. Strawson (1954). But his argument is not convincing.

Clark urges that a "sound with the same qualities could occur at a different [apparent] place," and concludes that "apparent location is not given by the addition of other qualities to the ones already present. There seems to be nothing analysable in the quality of the tone that makes it a tone from here rather than there" (2000, p. 61). That seems right; one cannot determine the subjective location of a tone by analyzing how it sounds independently of subjective location. But that shows only that subjective location is not due to the tone's other auditory mental qualities, such as pitch, timbre, and loudness. It is irrelevant to whether there is in addition a dedicated type of spatial mental quality that subserves the sensing of subjective location.

And there is compelling reason to think that there is. A tone's subjective location is determined by the spatial contrast of that tone with other sounds and with absence of sound spread out in one's auditory field. The mental quality pertaining to subjective location is not the same type of mental quality as pitch, timbre, and loudness. But that does not show that it not a type of mental quality. And since subjective location is determined by the spatial array of various sounds in one's auditory field, it is reasonable to count subjective auditory locations as themselves qualitative in nature. Subjective auditory location is a matter of dedicated mental qualities pertaining to relative location.

Those considerations are independent of the methodology of stimulus JNDs of the current proposal. But that methodology for determining mental qualities supports the same conclusion. Just as sounds are discriminable in respect of pitch, timbre, and loudness, so they are also discriminable in respect of spatial location. So we can test for JND locations of sounds, and extrapolate to corresponding mental qualities of auditory location. There is no serious reason to reserve the term, 'mental quality' for content qualities, such as pitch, loudness, color, odor and the like, and not apply it also to subjective spatial properties.

As for subjective size and shape, that may be elusive for many sounds if their subjective size is tiny. But it is always subjectively greater than a geometrical point. For some sounds subjective size and shape are plain, as with an orchestra or an explosion. And the sound of a jet plane moving along a runway not only has subjective size and shape, but also subjective location that determines subjective motion along the runway.

On the current proposal, the mental qualities in an individual's repertoire are fixed by extrapolating from the quality space of JND stimulus properties. Since JNDs are determined statistically and may sometimes overlap, mental qualities are themselves determined statistically and may overlap.

Some may insist that we individuate mental qualities in an absolute way, which would preclude statistical considerations and overlapping qualities. But it is unclear what serious reason there could be for that apart from the contention that we must individuate exclusively by what it is like for one. If the role mental qualities play in perceptual discrimination figures at all in individuating them, statistical considerations and occasional overlap are both natural and unavoidable.

And there is compelling evidence that the discriminative role of mental qualities is more fine-grained than subjective appearance, so that discriminative role is necessary to do complete justice to the ways mental qualities differ. Diana Raffman (2011) reports experimental work in which participants consciously judge as identical color stimuli that differ very slightly, though in performing matching tasks the participants register those differences. The relevant mental qualities plainly figure in the matching performance, since matching is for consciously seen stimuli. Nonetheless, what it is like for the participants is that the stimuli are identical. Discriminative role is more fine-grained than subjective appearances.

In related work, Arnaud Beauny and colleagues (2020) presented participants with very brief visual stimuli in the μsec range to establish a threshold at which they could consciously detect stimuli, but not also consciously identify them. Still, forced choice enabled participants to identify these stimuli well above chance. Amerio and colleagues (2024) report similar findings. Here too the stimuli were consciously seen, but perceptual role was more fine-grained than the subjective appearances. This disparity doubtless occurs in everyday experience, as when one is subjectively aware of something in a relatively generic way though one's behavior relies on more perceptually detailed information.

Since perceptual role is more fine-grained than what it is like, relying solely on the subjective appearances cannot individuate in as fine-grained a way as perceptual discrimination enables. That undermines the alleged need for absolute individuation encouraged by exclusive reliance on the subjective appearances; such absolute individuation would not reveal all the ways in which mental qualities differ. We must rely on discriminative ability, and accept statistical individuation and occasional overlap.

In the experimental work just reviewed, participants see stimuli consciously, but make perceptual discriminations more fine-grained than how they consciously appear. Their mental qualities are all conscious, but their discriminations are more fine-grained than what is consciously available to them. So their mental qualities must themselves differ in ways that are more fine-grained than the subjective appearances. Those conscious mental qualities differ in respect of aspects that remain unconscious.

In the work reported by Raffman, participants matching performance reveal differences in color stimuli unavailable from the subjective appearances. But the participants see themselves as matching consciously, based on what it is like for them. They see their matching as driven by their conscious mental qualities. So those conscious mental qualities must have unconscious qualitative aspects that enable that matching behavior.

But discrimination also occurs in the absence of any relevant conscious mental quality, and revealed by forced choice or priming. Since conscious mental qualities plainly sometimes have unconscious qualitative aspects, it is natural to explain discrimination in the absence of any relevant conscious mental quality as due to mental qualities that remain wholly unconscious. There are unconscious qualitative aspects, but no accompanying conscious aspects.

Those who favor individuating mental qualities exclusively by subjective appearance will deny that conscious mental qualities can have unconscious qualitative aspects, or that mental qualities can fail to be conscious at all. They might insist that unconscious discrimination is due solely to subpersonal processing, which is not mental. But if subpersonal processing did explain unconscious discrimination, it could as easily explain conscious discrimination as well. Conscious mental qualities would be idle in respect of perceptual discrimination. That is hard to accept.

And the only reason to appeal to such subpersonal processing is the insistence that we individuate mental qualities exclusively by their subjective appearances, a contention wholly without serious independent support. Section 7 will seek to disarm that claim by explaining why so many f it appealing despite its lack of support. For now it enough that the tie between mental qualities and perceptual discrimination is overwhelmingly strong, so that we cannot individuate or understand mental qualities independently of that tie.

4 Alternative approaches

The current proposal individuates mental qualities in a maximally fine-grained way by initially testing for JNDs between physical stimuli. And since discriminating stimuli requires being in relevantly distinct mental qualities, we can extrapolate from a space of JND stimuli to a corresponding space of mental qualities.

One might propose that we cut out that initial step of testing stimulus properties, and directly test the mental qualities themselves for JNDs, as envisaged, for example, by Goodman (1977, 226ff.). But those results would be far less objective. We have complete experimental control over what stimulus an individual is presented with. So when we test for JNDs between stimuli, there is complete objectivity about exactly what is being tested. That enables replication of test trials; we can repeat the same stimuli and determine on a statistical basis whether an individual can discriminate them. That is also crucial for objectivity.

By contrast, we have no serious control over what mental quality is elicited by presenting a particular stimulus. Perhaps it is exactly the same on repeated presentations, but perhaps not. So we cannot reliably or objectively test for discriminability of the mental qualities themselves.

The issue about what mental quality is elicited by a stimulus not only undermines fully objective testing; it also raises a general methodological concern. We could objectively assess whether introspective judgments about mental qualities are accurate only we could determine independently of introspective reports what mental qualities actually occur. But there is no way to determine what mental qualities occur independently of introspective reports except perceptual discrimination. We have no access to mental qualities other than subjective report and perceptual role. And we cannot simply stipulate that the same stimulus always elicits exactly the same mental quality. So assessing the accuracy of introspective reports will itself rely on testing stimuli for JNDs,

There is, in addition, an issue about the reliability of subjective introspective judgments about whether two mental qualities are exactly the same. There is strong evidence that introspective judgments about what qualitative state one is in is less reliable than direct discrimination of perceptible stimuli. Participants are significantly less accurate and slower and report greater difficulty when they instructed to match colors in respect of their subjective properties than when instructed simply to match the properties of the object itself (Arend & Reeves, 1986; Cornelissen & Brenner, 1995). Discriminating physical stimuli is easier and more reliable than assessing the corresponding mental qualities.

That direct perceptual discrimination is more accurate and faster than introspective decisions and subjectively easier should not be surprising. In making perceptual decisions participants function somewhat like measuring devices that register responses to stimuli. Introspective judgments, by contrast, require additional psychological processing, making them slower, less accurate, and more demanding.

This fits well with the results reported by Raffman, Beauny et al, and Amerio et al. As those findings show, perceptual discrimination can unconsciously reveal differences between physical stimuli more fine-grained than what is available consciously. That is confirmed in elegant work by Liam Norman and colleagues (2014), which demonstrates that perceived surface colors are registered unconsciously, independently of conscious color experience. Perceptual discrimination operates in part independently of what introspection represents, and in a way more fine-grained than what is available to introspection.

Constructing a quality space from JNDs, whether based on stimuli or mental qualities, would take a huge amount of time. Nobody would in actual practice construct a quality space that way, determining one JND at a time. The JND methodology is not a recommendation for actually constructing quality spaces. It is a theoretical account of how a space could in principle be generated that represents all the mental qualities available to an individual at a time. For actual practice a shortcut in needed.

And a technique known as multidimensional scaling does allow the construction of reasonably revealing quality spaces in a workable amount of time. One starts with a number of triads of samples, such that one sample in each triad is judged to be more similar to a second than to a third. It is typically assumed that the samples are the mental qualities themselves,

but one could start with triads of physical stimuli and extrapolate to mental qualities as in Section 3. Multidimensional scaling then enables the construction a quality space of all the samples of the relevant type. The details of this procedure are not relevant here; Clark presents an admirably accessible and useful account (1993, §4.5 and Appendix).

Because the triads are defined by relative similarity among the three samples, quality spaces constructed by multidimensional scaling encourages an unqualified holism about mental qualities. Each mental quality would be fixed by relative similarity relations to all the others. Still, one could see multidimensional scaling as simply a shortcut, and think of mental qualities as fixed theoretically by the highly local relation of JNDs.

But how is relative similarity of the samples in the triads to be established? One way would be to construct relative similarity from relative numbers of JNDs. Clark argues against that, noting that distances between JND stimuli will vary (1993, p. 91). But that source of variation is arguably minor compared with any other measure of relative similarity. And vastly fewer JNDs would be needed for that procedure than relying solely on JNDs, as in Section 3.

Clark proposes that we use instead an account Quine tentatively offers of relative perceptual similarity, which appeals to stimulus generalization. Roughly, a stimulus [a] is more similar to [b] than to [c] if conditioning established for [a] generalizes to [b] but not to [c] (Quine, 1974, pp. 16–18). As Clark remarks (1993, pp. 117–118), this procedure would accommodate nonlinguistic animals, since it is independent of verbal report.

But Quine's proposal will not do. As Quine himself notes (1974, p. 18), perceptual similarity so defined is "a very disconnected relation"; it will not apply to all stimuli. More important, one will almost certainly get highly different generalization results for the different types of response that could figure in the relevant conditioning. Quine's suggestion cannot provide the relative similarity relations required for multidimensional scaling.

So that proposal also cannot enable the construction of quality spaces for nonlinguistic creatures, where verbal report is unavailable. Clark suggests in passing that confusion probabilities could help do that (1993, p. 118). But it is unlikely that the relevant confusion matrices would underwrite similarity relations robust enough to underwrite multidimensional scaling.

Since neither Quine's proposal nor confusion probabilities can establish the relative similarity needed for multidimensional scaling, the only recourse seems to be subjective judgments of similarity relations. And indeed Clark often writes as though this is what would be used (e.g., Clark, 1993, §§4.4, 4.6.1). But the unreliability of subjective introspective judgments of JNDs between conscious mental qualities would be greatly amplified in subjectively assessing of whether one sample is more similar to a second than to a third.

Some cases of relative similarity do seem intuitively quite clear; orange is plainly more similar to red than to green. But it is not obvious that the clarity of those cases carries over to relative similarity in general. Is red more similar to turquoise than to forest green? Without an understanding of what relative similarity amounts to in the general case, it is unclear even how to think about many such questions.

And absent an understanding of what relative similarity amounts to in the general case, there is no hope to apply that notion to nonlinguistic animals. By contrast, the method of stimulus JNDs readily accommodates testing nonlinguistic animals, since they can be trained to respond differentially in ways that reveal such JNDs.

Multidimensional scaling delivers highly workable quality spaces, which do extremely well for practical and illustrative purposes. But for an objective theoretical account of what mental qualities are and how to individuate them, we must rely on JNDs between discriminable stimulus properties.

5 Fixing what it is like

We know about mental qualities in two very different ways, by their role in perceiving and by how they subjectively appear in consciousness. So a satisfactory account of mental qualities must deal with both perceptual role and subjective appearance. No account is acceptable that treats mental

qualities only in respect of the subjective appearance or only in respect of perceptual role.

Individuating mental qualities solely by subjective appearance risks discounting perceptual role as at best incidental to their nature. The subjective appearances say nothing about the causal relations that mental qualities may exhibit. And causal relations are pivotal to perceptual role. In addition, if one could individuate mental qualities by appeal solely to their subjective appearances, those appearances would override any input about them from perceptual role.

Because individuating mental qualities exclusively by subjective appearance leads to sidelining perceptual role, it also encourages holding, with Nagel and others, that their mental reality of mental qualities is exhausted by how they subjectively appear in consciousness (Rosenthal, 2022). On this reductive picture, perceptual role would just provide adjunct information about mental qualities understood exclusively by subjective appearance.

And since individuating solely by subjective appearances leads to discounting perceptual role, it also encourages holding that one can conceive of mental qualities as inverted from one individual to another in ways that remain altogether undetectable. If perceptual role were irrelevant to the individuation of mental qualities, quality inversion in another individual would be undetectable.

The issue about undetectable inversion is nicely captured by the machinery of quality spaces. A quality space represents its members by the relations they bear to one another. So if a quality space has no axis of symmetry, any inversion would distort those relations. And if those relations reflect the perceptual roles of the mental qualities, any inversion would be readily detectable.

But if there were an axis of symmetry, the relations on one side of the axis would be identical with those on the other. So the space would represent the qualities on each side as identical with those on the other. Having both sides, then, would be redundant, and the space would in effect fold over on itself in respect of that axis. Undetectable interpersonal inversion would be possible only if perceptual role were irrelevant to the relations that define the quality space. For more on quality spaces and undetectable inversion, see Clark's masterful (2022).

Individuating mental qualities by subjective appearance leads to dismissing perceptual role as irrelevant to their nature. Are things symmetrical in this respect? Does individuating mental qualities by perceptual role, as on the current proposal, lead to discounting the subjective appearances? Since perceptual discrimination can occur without being conscious, individuation by perceptual role need not involve the subjective appearances in any way. And an account on which we could not accommodate the subjective appearances of conscious mental qualities might well seem significantly more unacceptable than one that takes no account of perceptual role.

But things are not symmetrical in that way. Because the subjective appearances are silent about causal relations, individuating mental qualities by subjective appearance offers no path to build an account of mental qualities in respect of their perceptual roles. By contrast, there is nothing in the individuation of mental qualities by perceptual role that forecloses an account of mental qualities in respect of their subjective appearance in consciousness. Indeed, there are considerations that point to an account of the subjective appearances that is built on an account of perceptual role.

The subjective appearances are what it is like for one. They are the way one's mental life occurs subjectively in one's stream of consciousness. What it is like is the way one's mental goings on subjectively appear to one. But for creatures with the relevant linguistic ability, sincere verbal reports of what it's like for one is the most accurate and revealing way to determine what it is like for an individual.

And we describe the subjective appearances by appeal to cases of perceiving an object with characteristic perceptible properties. We characterize what it is like for one as red when it is the type of experience typically elicited by seeing a red object. And we might say, for example, that what it is like is like seeing the yellow of a canary, or that it is closer to seeing the yellow of a lemon than that of a canary. Or instead of appealing to objects with well-known perceptible properties, one can describe what it is like by reference to objects that are perceptually available. We describe what it

is like by appeal to the perceiving of various types of object. Indeed, it is unclear how else one could describe what it is like for one.

If one rejected describing the subjective appearances by appeal to its seeming subjectively that one perceives particular things, one would be reduced to maintaining that what it is like cannot be described at all. One would contend, with Block, that "[t]he best you can do is use words to point to a phenomenon that the reader has to experience from the first person point of view" (2015, p. 47). But that simply is not the way things are. We can and typically do describe what it is like for us by saying what it subjectively appears that we are perceiving.

Such verbal reports reveal the nature of what it is like in particular cases. And because we describe what it is like in terms of perceiving particular types of object, it is inviting to expect that we can build an account of subjective appearances by on an account of perceptual roles. More in Section 6.

Verbal report of what it is like for one depend on introspective awareness. But the concerns raised in Section 4 about the accuracy of introspection are irrelevant here. Introspection purports to reveal what mental state one is in by characterizing what mental state one subjectively appears to be in. And the introspective appearances about what states one is actually in are not always fully accurate. But introspection cannot go wrong about what the subjective appearances themselves are. Since the subjective appearances constitute introspective awareness, introspection is fully accurate about those subjective appearances.

It might seem that the introspective appearances sometimes outstrip what verbal report can capture. But that impression is once again due to construing the introspective appearances atomistically, and so nonrelationally. Once one acknowledges that introspective content is relational it is plain that careful verbal reports can do justice to it.

Sincere verbal reports are the gold standard for determining what it is like for an individual. Psychological noise of one sort or another may interfere with a report's accurately reflecting one's subjective awareness. But that is relatively rare, and there are techniques that can filter verbal reports for accuracy (Dienes & Perner, 2004). Sincere verbal report is typically a fully objective indicator of what it is like for somebody.

Verbal reports are also, when available, the most objective indicator of whether a mental state is conscious. If an individual reports perceiving something, that perception is conscious; if the individual firmly denies perceiving anything, then even if guessing about the stimulus is accurate well above chance, that perception is not conscious.

Verbal report reliably reveals the subjective appearances when it is available. But the unavailability of verbal report, as with creatures that lack the relevant linguistic capacity, does not by itself show that no subjective appearances occur. Other indicators are then required; a proposal for non-linguistic creatures is suggested at the end of Section 6. But for present purposes it is good enough to rely on the clear cases, in which verbal report is available.

The use of verbal report to determine whether a perception is conscious has been challenged. Various seemingly irrelevant factors, such as attention and alertness, can affect how readily one reports perceiving a stimulus. Those factors can make one more cautious to issue a report, or more easygoing about doing so. When one is more cautious, a more intense stimulus is required to elicit such a report; when one is more easygoing a less intense stimulus will result in one's reporting. Signal-detection theory describes this in terms of how conservative or liberal one's criterion is for making a perceptual decision.

But variations in one's criterion does nothing to undermine the reliability of verbal report for indicating whether a mental state is conscious. How conservative or liberal one's signal-detection criterion is determines whether one regards oneself as perceiving a stimulus. But regard oneself as perceiving constitutes being aware of oneself as perceiving. And if one is aware of oneself as perceiving, that perception is conscious, and one will report that one perceives. Otherwise one is unaware of oneself as perceiving, and one will report that one does not perceive. So if a perception does actually occur on that occasion, it is then unconscious.

It might be tempting to think that if variation in one's criterion is due to seemingly extraneous factors, such variation that should not affect

whether one's perception is conscious. But that is simply a mistake. A perception's being conscious consists solely in its appearing subjectively to one that one perceives. So all that matters is whether one subjectively takes oneself to perceive. A more conservative criterion results in fewer perceptions being reported because it results in fewer perceptions being conscious at all; a more liberal criterion leads to more reports because more perceptions are conscious. One is inclined to issue a report if, but only if, one subjectively takes oneself to perceive. Subjective awareness by itself constitutes a perception's being conscious.

A recent terminological coinage has suggested to some that whether a perception is conscious is not a matter of subjective awareness. Following a distinction by Jim Cheesman and Philip Merikle (1984) between subjective and objective perceptual thresholds, some authors have come to distinguish subjective from objective measures of consciousness (e.g., Dienes, 2007; Timmermans & Cleeremans, 2015). On this distinction, a state is conscious on an objective measure if there is above-chance detection of perceptual information. A state is conscious on the subjective measure, by contrast, if the individual verbally reports being in that state.

But an objective measure so construed is not a measure of whether a perception is conscious; it is a measure of perceptual role. An objective measure, so construed, is simply a measure of whether an individual has perceptually registered the stimulus. So if one took that so-called objective measure also to indicate whether a perception is conscious, perceiving would coincide with conscious perceiving. Any above-chance perceptual detection would count as conscious, even if the individual firmly denies perceiving anything (e.g., Phillips, 2018). There could then be no perceiving that fails to be conscious.

The use of the terms, 'objective' and 'subjective', to characterize those two approaches might suggest that the so-called subjective measure cannot be objective. But that is a mistake. Consciousness is a matter of how things appear subjectively. The only measure of consciousness, as against perceptual performance, is what an individual is subjectively aware of, and so willing to report. That is the so-called subjective measure. But being conscious is an entirely subjective matter. So the so-called subjective measure is a fully objective measure of subjectivity.

6 Explaining what it is like

Our conscious mental lives consist in subjective mental appearances. What it is like for one is how one's mental goings on subjectively appear to one. It is the way the conscious aspect of one's mental life appears in one's stream of consciousness. The subjective appearances are appearances of being in particular mental states.

So when there is a subjective qualitative appearance of red, for example, that subjective appearance consists of its seeming subjectively as it does when one consciously sees something red. More generally, what it is like for one consists in its seeming subjectively as it does when one perceives some particular type of object.

Such subjective appearances typically result from one's actually perceiving the relevant type of object. But actually having such a perception is not necessary. Each subjective appearance consists in its subjectively appearing as it characteristically does when one does perceive an object of the relevant type. And it can seem subjectively that way even when one is not perceiving such an object.

This account of what it is for mental qualities to exhibit a subjective appearance is reflected in the way we describe what it is like for us. As already noted, we describe what it is like by appeal to perceiving an object of the type that characteristically elicits the subjective appearance in question. We have no other cognitive tools for describing what it is like.

Conscious mental qualities have both perceptual roles and subjective appearances. Those subjective appearances consist in its subjectively appearing as it does when one perceives an object of a particular type. And perceiving an object of a particular type involves discriminating its relevant perceptible properties. A subjective appearance of red is the way it subjectively appears when one sees a red object. And seeing a red object

involves mental qualities that enable the discrimination of red stimuli from stimuli of other colors.

So one cannot describe mental qualities in respect of their subjective appearances without referring to the perceptual roles of relevant mental qualities. Any account of what it is for a mental quality to exhibit a subjective appearance requires and builds upon an account of the roles mental qualities play in perceiving.

Perceiving a stimulus of a particular type, moreover, consists in perceptually discriminating that type of stimulus from its nearby perceptible neighbors. And we describe such perceptual discrimination in terms of the JNDs that figure in constructing the quality spaces that represent mental qualities in respect of their perceptual role. Since the subjective appearances consist in its appearing subjectively as it does when one perceives a stimulus of a particular type, those quality spaces also underlie an objective account of the subjective appearances themselves.

Quality spaces represent mental qualities in respect of their perceptual roles by extrapolating from quality spaces that represent JNDs between stimulus properties that the mental qualities enable one to discriminate. And perceiving an object of a particular type is perceptually discriminating an object with relevant stimulus properties. A subjective appearance consists in its subjectively appearing as it does when one perceives an object of a particular type. So a quality space of stimulus properties that figure in perceiving that type of object enables an objective description of how it subjectively appears for any case of there being something it is like for one. The result is the use of quality spaces for objectively representing mental qualities in respect of both perceptual role and subjective appearance.

Holger Lyre has urged that subjective similarity judgments are needed because JNDs provide only "pairwise discriminations," which cannot "order our experiences" (2022, p. 8). But the subjective appearances consist in its subjectively appearing to one as it does when one perceives a particular type of stimulus property. So JNDs between types of stimulus property do after all determine how our subjective experiences are ordered.

As already noted, mental qualities are typically more fine-grained in respect of their perceptual role than they are in respect of their subjective appearances. This is to be expected given the different ways in which quality spaces figure in representing perceptual role and subjective appearances. Perceptual roles are a matter of the JNDs used in constructing the quality space; so they are as fine-grained as an individual's discriminative abilities allow.

The subjective appearances, by contrast, consist in its appearing subjectively as it does when one perceives an object that exhibits stimulus properties fixed by relative location in the quality space. But the stimulus properties that serve to individuate that object can vary somewhat, involving a region in the quality space. There will often be a range of stimulus properties that can figure in the way it subjectively appears when one perceives an object of the relevant type. By contrast, when perceptual discrimination is at issue the range of stimulus properties must be far narrower. So mental qualities will be less fine-grained in respect of their subjective appearances than in respect of their perceptual roles.

Subjective appearances consist in its subjectively appearing as it does when one perceives a particular type of object. That may seem to have echoes of a higher-order theory of consciousness, on which a mental state's being conscious consists in one's having a higher-order awareness of being in that state (Rosenthal, 2005).

But to say that subjective appearances are its appearing subjectively as it does when one perceives some type of object is not to give a theoretical account of subjective appearances. It is simply to specify what the phenomenon of subjective appearances consists in. Whatever theory of consciousness that one might adopt would have to specify what the phenomenon of subjective appearance is. And since there is no other way to specify what that phenomenon is, any theory of consciousness would have to accept the specification given here, on pain of being unable to give any specification at all. There is no appeal to a higher-order theory of consciousness.

Still, the notion of something's subjectively appearing to one is simply taken for granted in this specification of what subjective appearances are. So a complete account would also have to explain what it is for something to appear subjectively to somebody. A higher-order theory would explain

that in objective terms as a matter of one's having a higher-order awareness of something (Rosenthal, 2005). But one could instead conjoin the current account with a different explanation of what it is for something to appear subjectively to somebody. Or the current account can stand on its own, without any additional explanation. And the current account is on its own fully objective, since it specifies each type of subjective appearance by objectively constructed quality spaces.

As already noted, it can subjectively appear that one is perceiving something even if one is not perceiving that thing. We describe how it subjectively appears in terms of perceiving something, and typically the relevant perceiving does occur. But it need not do so.

And dramatic disparities between what one perceives and what one subjectively appears to perceive can emerge in connection with change blindness. Participants in change blindness report being subjectively unaware of a change, sometimes relatively salient, that occurs in a presented display. But even when a participant is subjectively unaware of a change, independent tests can demonstrate that the change was perceived, albeit unconsciously (e.g., Fernandez-Duque & Thornton, 2000; Thornton & Fernandez-Duque, 2001; replicated by Laloyaux et al., 2006).

When participants are subjectively unaware of a change, their post-change subjective awareness sometimes continues to represent the pre-change stimulus. But since post-change visual input is due to the post-change stimulus, the post-change perceptual state will reflect that post-change stimulus. So on miss trials on which participants continue subjectively to perceive the pre-change stimulus, post-change subjective awareness diverges from post-change perceptual states.

This striking disparity has been confirmed in a replication of John Grimes's (1996) change-blindness paradigm, in which the change occurs during a saccade when almost no retinal input reaches visual cortex. On some miss trials participants actually report continuing to be subjectively aware of the pre-change item (Odegaard et al., 2022). These experiments occur in highly controlled conditions. But change blindness doubtless occurs in everyday perceptual experience, with attendant but unnoticed disparities between subjective awareness and perceptual state.

Sincere verbal report when available is the gold standard for subjective appearance. But for many that creatures lack the relevant linguistic abilities it is nonetheless natural to think that there is something it is like for many of them to perceive things. Many researchers adopt a so-called marker approach (e.g., Allen & Trestman, 2016), on which behavioral and other observable similarities to humans suggest that there is something it is like for those creatures to perceive. But that strategy shows nothing unless one can also demonstrate that the observable similarities in nonlinguistic creatures are not due simply to unconscious perceiving. This concern is far from idle given the extraordinarily detailed and natural behavior of a blindsight rhesus monkey (Humphrey, 1974).

The disparities between subjective appearance and perceptual state that occur in change blindness suggest a more promising test. Change blindness has been established in nonhuman animals as different as chimpanzees and pigeons (Herbranson, 2022). And independent measures of change detection have also been studied for nonhuman animals (Wright, 2022). So one could investigate when an animal misses a change, independent tests sometimes reveal that the change was nonetheless registered perceptually. Absent attendant confusion, the best explanation would be that the animal missed the change consciously but perceived it unconsciously. And that would demonstrate that animals of that type perceive consciously as well as unconsciously. This is a strategy well worth exploring.

7 Dispelling the sense that subjective appearance is basic

An objective account of mental qualities individuates them in respect of their role in perceptual discrimination. Since mental qualities enable such discrimination, a quality space of discriminable stimuli also fixes the corresponding mental qualities.

What it is like for one to have a conscious mental quality, moreover, is for it to appear subjectively as it does when one perceives an o of the relevant type. Since we describe what it is like in terms of the types of

objects we perceive, quality spaces that represent mental qualities in respect of perceptual role will also represent the various types of subjective appearance.

Nonetheless, many still insist on understanding mental qualities exclusively by way of the subjective appearances, altogether independently of perceptual role. There is little if any substantive argument for that approach, which tends to rely only on reports of intuitions. But it is unsound to oppose argument with intuitions. Intuitions are appealing encapsulations of theoretical claims (Maćkiewicz et al., 2023; Rosenthal, 2022, 2024). And when those claims are packaged as intuitions their theoretical status remains tacit and unarticulated, so that it seems to need no support and to be immune from assessment.

Casting a claim as an intuition converts it into an academic conversation stopper, blocking any argument or explanation. Recognizing intuitions as disguised theoretical claims forces them back into the space of explanation and reasoning. And these methodological considerations to one side, there are in every case good reasons to dismiss specific intuitions that construe mental qualities atomically, and so as independent of perceptual role (Weisberg, 2023; Rosenthal, 2010, §I; Rosenthal, 2022, §14.1.1).

Still, it would be useful to understand the widespread appeal of the view that mental qualities must be understood by way of their subjective appearances. And that appeal arguably is due to a mistake about how to apply an important feature of scientific theorizing to the special case of mental qualities. Exposing that mistake may help disarm the tendency to take the subjective appearances of mental qualities as fundamental.

It is generally accepted that we must be able to describe physical reality in mathematical terms; as Galileo put it, "the universe […] is written in mathematical language" (Galilei, 2008, p. 183). This has led some to infer that physical reality contains no properties of color, sound, and the like. And those who hold that those properties cannot occur in physical reality typically recast colors, sounds, and so forth as merely mental properties, in effect relocating them to the mind.

But there is no reason to think that properties such as colors and sounds resist mathematical description unless one conceives of those properties as they subjectively appear to us. It is only the way those properties appear subjectively that rules out mathematical description. Conceived as they appear subjectively, such properties need somehow to sidestep Galilean strictures about physical reality. Relocating those properties as exclusively mental seems the natural way to do that.

This relocation move is often simply taken for granted as obviously correct. But it has a striking consequence that is typically unnoticed. It is only because one conceives of colors, sounds, and so forth as they subjectively appearance to us that one relocates them as exclusively mental properties. So when they are relocated, they will still be conceived of in respect of their subjective appearances. It is this relocation move that makes it seem inevitable that we understand mental qualities in terms of their subjective appearances.

But the relocation move is simply a mistake (Rosenthal, 2005, ch. 6). We do not ordinarily conceive of colors, sounds, and so forth as they subjectively appear to us. That is plain because we do not ordinarily conceive of those properties as occurring only when they are perceived. And when those properties occur unperceived, they have no subjective appearances. Physical colors, sounds, and so forth exist independently of being perceived, and independently of subjective appearance. So there is no obstacle to conceiving of them as physical properties that readily admit standard mathematical treatment. There is no problem for the relocation move to solve.

Since there is no reason to relocate properties such as physical colors to the mind, there is also no reason to conceive of the mental qualities that correspond to those physical properties primarily in respect of their subjective appearances. We can conceive of them initially in respect of their perceptual roles, and use those perceptual roles to explain the subjective appearances. They are mental properties that sometimes occur consciously and sometimes not.

That clears the way for a fully objective account of mental qualities, which we can see as a final step in the Galilean scientific revolution.

Acknowledgments

This is a significantly extended version of a presentation in the Symposium on Mathematical Spaces for Conscious Experiences at the 2023 meeting of the Association for the Scientific Study of Consciousness. I am very grateful to two anonymous reviewers for useful suggestions.

References

Allen, C., & Trestman, M. (2016). Animal consciousness. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford encyclopedia of philosophy*. https://plato.stanford.edu/archives/win2020/entries/consciousness-animal/

Amerio, P., Michel, M., Goerttler, S., Peters, M. A. K., & Cleeremans, A. (2024). Unconscious perception of vernier offsets. *Open Mind*, *8*, 739–765. https://doi.org/10.1162/opmi_a_00145

Arend, L., & Reeves, A. (1986). Simultaneous color constancy. *Journal of the Optical Society of America A*, *3*(10), 1743. https://doi.org/10.1364/JOSAA.3.001743

Bartleson, C. J. (1960). Memory colors of familiar objects. *Journal of the Optical Society of America*, *50*(1), 73. https://doi.org/10.1364/JOSA.50.000073

Beauny, A., De Heering, A., Muñoz Moldes, S., Martin, J.-R., De Beir, A., & Cleeremans, A. (2020). Unconscious categorization of sub-millisecond complex images (M. E. Król, Ed.). *PLOS ONE*, *15*(8), e0236467. https://doi.org/10.1371/journal.pone.0236467

Berger, J. (2015). The sensory content of perceptual experience. *Pacific Philosophical Quarterly*, *96*(4), 446–468. https://doi.org/10.1111/papq.12110

Block, N. (1978). Troubles with functionalism. In C. W. Savage (Ed.), *Minnesota Studies in the Philosophy of Science* (pp. 261–325, Vol. 9). University of Minnesota Press.

Block, N. (2015). The puzzle of perceptual precision. In T. Metzinger & J. M. Windt (Eds.), *Open MIND* (Vol. 5(T)). MIND Group. https://doi.org/10.15502/9783958570726

Cheesman, J., & Merikle, P. M. (1984). Priming with and without awareness. *Perception & Psychophysics*, *36*(4), 387–395. https://doi.org/10.3758/BF03202793

Clark, A. (1993). *Sensory qualities*. Clarendon Press.

Clark, A. (2000). *A theory of sentience*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198238515.001.0001

Clark, A. (2022). Qualitative character and conceivable inversions. In J. Weisberg (Ed.), *Qualitative Consciousness* (1st ed., pp. 159–178). Cambridge University Press. https://doi.org/10.1017/9781108768085.012

Cornelissen, F. W., & Brenner, E. (1995). Simultaneous colour constancy revisited: An analysis of viewing strategies. *Vision Research*, *35*(17), 2431–2448. https://doi.org/10.1016/0042-6989(94)00318-1

Dienes, Z. (2007). Subjective measures of unconscious knowledge. In R. Banerjee & B. K. Chakrabarti (Eds.), *Models of brain and mind: Physical, computational and psychological approaches* (pp. 49–64). Elsevier. https://doi.org/10.1016/S0079-6123(07)68005-4

Dienes, Z., & Perner, J. (2004). Assumptions of a subjective measure of consciousness: Three mappings. In R. J. Gennaro (Ed.), *Higher-order theories of consciousness: An anthology* (pp. 173–199). John Benjamins. https://doi.org/10.1075/aicr.56.10die

Dikeçligil, G. N., Yang, A. I., Sanghani, N., Lucas, T., Chen, H. I., Davis, K. A., & Gottfried, J. A. (2023). Odor representations from the two nostrils are temporally segregated in human piriform cortex. *Current Biology*, *33*(24), 5275–5287.e5. https://doi.org/10.1016/j.cub.2023.10.021

Fernandez-Duque, D., & Thornton, I. M. (2000). Change detection without awareness: Do explicit reports underestimate the representation of change in the visual system? *Visual Cognition*, *7*(1–3), 323–344. https://doi.org/10.1080/135062800394838

Fink, S. B., Kob, L., & Lyre, H. (2021). A structural constraint on neural correlates of consciousness. *Philosophy and the Mind Sciences*, *2*. https://doi.org/10.33735/phimisci.2021.79

Galilei, G. (2008). *The essential Galileo* (1st ed). Hackett Publishing.

Gescheider, G. A. (1997). *Psychophysics: The fundamentals* (3rd ed.). Lawrence Erlbaum Associates.

Goodman, N. (1977). *The structure of appearance*. Springer. https://doi.org/10.1007/978-94-010-1184-6

Gori, S. (2016). Assimilation. In M. R. Luo (Ed.), *Encyclopedia of Color Science and Technology* (pp. 57–60). Springer. https://doi.org/10.1007/978-1-4419-8071-7_272

Grice, H. P. (1962). Some remarks about the senses. In R. J. Butler (Ed.), *Analytical philosophy, first series*. Oxford University Press.

Grimes, J. (1996). On the failure to detect changes in scenes across saccades. In A. Elizabeth (Ed.), *Perception* (pp. 89–110). Oxford University Press.

Hansen, T., Olkkonen, M., Walter, S., & Gegenfurtner, K. R. (2006). Memory modulates color appearance. *Nature Neuroscience*, *9*(11), 1367–1368. https://doi.org/10.1038/nn1794

Herbranson, W. T. (2022). Change blindness. In J. Vonk & T. K. Shackelford (Eds.), *Encyclopedia of Animal Cognition and Behavior* (pp. 1279–1282). Springer. https://doi.org/10.1007/978-3-319-55065-7_1358

Humphrey, N. K. (1974). Vision in a monkey without striate cortex: A case study. *Perception*, *3*(3), 241–255. https://doi.org/10.1068/p030241

Jackson, F. (1986). What Mary didn't know. *The Journal of Philosophy*, *83*(5), 291. https://doi.org/10.2307/2026143

Kingdom, F. A. A., & Prins, N. (2016). *Psychophysics: A practical introduction* (2nd ed.). Elsevier Academic Press.

Kleiner, J. (2024). Towards a structural turn in consciousness science. *Consciousness and Cognition*, *119*, 103653. https://doi.org/10.1016/j.concog.2024.103653

Kleiner, J., & Ludwig, T. (2024). What is a mathematical structure of conscious experience? *Synthese*, *203*(3), 89. https://doi.org/10.1007/s11229-024-04503-4

Klincewicz, M. (2011). Quality space model of temporal perception. In A. Vatakis, A. Esposito, M. Giagkou, F. Cummins, & G. Papadelis (Eds.), *Multidisciplinary aspects of time and time perception* (pp. 230–245). Springer. https://doi.org/10.1007/978-3-642-21478-3_18

Kob, L. (2023). Exploring the role of structuralist methodology in the neuroscience of consciousness: A defense and analysis. *Neuroscience of Consciousness*, *2023*(1), niad011. https://doi.org/10.1093/nc/niad011

Laloyaux, C., Destrebecqz, A., & Cleeremans, A. (2006). Implicit change identification: A replication of Fernandez-Duque and Thornton (2003). *Journal of Experimental Psychology: Human Perception and Performance*, *32*(6), 1366–1379. https://doi.org/10.1037/0096-1523.32.6.1366

Laming, D. (1997). *The measurement of sensation*. Oxford University Press.

Lee, A. Y. (2021). Modeling mental qualities. *The Philosophical Review*, *130*(2), 263–298. https://doi.org/10.1215/00318108-8809919

Lee, A. Y. (2024). Objective phenomenology. *Erkenntnis*, *89*(3), 1197–1216. https://doi.org/10.1007/s10670-022-00576-0

Lyre, H. (2022). Neurophenomenal structuralism: A philosophical agenda for a structuralist neuroscience of consciousness. *Neuroscience of Consciousness*, *2022*(1), niac012. https://doi.org/10.1093/nc/niac012

Maćkiewicz, B., Kuś, K., & Hensel, W. M. (2023). The influence of philosophical training on the evaluation of philosophical cases: A controlled longitudinal study. *Synthese*, *202*(4), 113. https://doi.org/10.1007/s11229-023-04316-x

Macpherson, F. (2011). Individuating the senses. In F. Macpherson (Ed.), *The senses: Classic and contemporary philosophical perspectives* (pp. 3–43). Oxford University Press.

Malach, R. (2021). Local neuronal relational structures underlying the contents of human conscious experience. *Neuroscience of Consciousness*, *2021*(2), niab028. https://doi.org/10.1093/nc/niab028

Martina, G. (2023a). How we talk about smells. *Mind & Language*, *38*(4), 1041–1058. https://doi.org/10.1111/mila.12440

Martina, G. (2023b). Smell identification and the role of labels. *Philosophical Psychology*, 1–25. https://doi.org/10.1080/09515089.2023.2255227

Meehan, D. B. (2007). *The qualitative character of spatial perception*. City University of New York Graduate Center. https://academicworks.cuny.edu/cgi/viewcontent.cgi?article=2433&context=gc_etds

Morrison, J. (2015). Anti-atomism about color representation. *Noûs*, *49*(1), 94–122. https://doi.org/10.1111/nous.12018

Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, *83*(4), 435. https://doi.org/10.2307/2183914

Nagel, T. (2024). *What is it like to be a bat?* Oxford University Press. https://doi.org/10.1093/oso/9780197752791.001.0001

Norman, L. J., Akins, K., Heywood, C. A., & Kentridge, R. W. (2014). Color constancy for an unseen surface. *Current Biology*, *24*(23), 2822–2826. https://doi.org/10.1016/j.cub.2014.10.009

Odegaard, B., Lee, A., Sans, A., Lee, I., Ng, L., Haun, A., Chesney, D., Rosenthal, D., & Fallon, F. (2022). Detecting changes in visual scenes during saccades: Replicating and extending John Grimes's experiments [Poster.]. *Vision Sciences Society Virtual Meeting*.

Pauen, M. (2017). The functional mapping hypothesis. *Topoi*, *36*(1), 107–118. https://doi.org/10.1007/s11245-015-9302-y

Phillips, I. (2018). Unconscious perception reconsidered. *Analytic Philosophy*, *59*(4), 471–514. https://doi.org/10.1111/phib.12135

Prentner, R. (2019). Consciousness and topologically structured phenomenal spaces. *Consciousness and Cognition*, *70*, 25–38. https://doi.org/10.1016/j.concog.2019.02.002

Quine, W. V. (1969). *Ontological relativity and other essays*. Columbia University Press. https://doi.org/10.7312/quin92204

Quine, W. V. (1974). *The roots of reference*. Open Court.

Quine, W. V. (2013). *Word and object*. The MIT Press. https://doi.org/10.7551/mitpress/9636.001.0001

Raffman, D. (2011). Vagueness and observationality. In G. Ronzitti (Ed.), *Vagueness: A Guide* (pp. 107–121). Springer. https://doi.org/10.1007/978-94-007-0375-9_5

Rosenthal, D. (1983). Reductionism and knowledge. In L. S. Cauman, I. Levi, C. Parsons, & R. Schwartz (Eds.), *How many questions? Essays in honor of Sidney Morgenbesser* (pp. 276–300). Hackett Publishing Company.

Rosenthal, D. (1991). The independence of consciousness and sensory quality. *Philosophical Issues*, *1*, 15. https://doi.org/10.2307/1522921

Rosenthal, D. (2005). *Consciousness and mind* (1. publ). Clarendon Press.

Rosenthal, D. (2010). How to think about mental qualities. *Philosophical Issues*, *20*(1), 368–393. https://doi.org/10.1111/j.1533-6077.2010.00190.x

Rosenthal, D. (2015). Quality spaces and sensory modalities. In P. Coates & S. Coleman (Eds.), *Phenomenal Qualities* (pp. 33–65). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198712718.003.0002

Rosenthal, D. (2016). Quality spaces, relocation, and grain. In J. R. O'Shea (Ed.), *Sellars and his Legacy* (pp. 149–185). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198766872.003.0009

Rosenthal, D. (2019). There's nothing about Mary. In S. Coleman (Ed.), *The Knowledge Argument* (1st ed., pp. 32–61). Cambridge University Press. Retrieved October 26, 2024, from https://www.cambridge.org/core/product/identifier/9781316494134%23CN-bp-2/type/book_part

Rosenthal, D. (2022). Mental appearance and mental reality. In J. Weisberg (Ed.), *Qualitative consciousness: Themes from the philosophy of David Rosenthal* (pp. 243–271). Cambridge University Press.

Rosenthal, D. (2024). Methodological considerations for the study of mental qualities. In J. Hvorecký, T. Marvan, & M. Polák (Eds.), *Conscious and unconscious mentality: Examining their nature, similarities, and differences* (pp. 91–111). Routledge.

Sanford, E. M., & Halberda, J. (2023). A shared intuitive (mis)understanding of psychophysical law leads both novices and educated students to believe in a just noticeable difference (JND). *Open Mind*, *7*, 785–801. https://doi.org/10.1162/opmi_a_00108

Schanda, J. (2007). CEI colorimetry [OCLC: ocm77256743]. In J. Schanda (Ed.), *Colorimetry: Understanding the CIE system* (pp. 25–78). Wiley.

Sellars, W. (1956). Empiricism and the philosophy of mind. In H. Feigl & M. Scriven (Eds.), *Minnesota Studies in the Philosophy of Science* (pp. 253–329, Vol. 1). University of Minnesota Press.

Shepard, R. N. (1982). Geometrical approximations to the structure of musical pitch. *Psychological Review*, *89*(4), 305–333. https://doi.org/10.1037/0033-295X.89.4.305

Soranzo, A. (2016). Simultaneous color contrast. In M. R. Luo (Ed.), *Encyclopedia of Color Science and Technology* (pp. 1149–1152). Springer. https://doi.org/10.1007/978-1-4419-8071-7_268

Strawson, P. F. (1954). Particular and general. *Proceedings of the Aristotelian Society*, *54*(1), 233–260. https://doi.org/10.1093/aristotelian/54.1.233

Tallon-Baudry, C. (2022). The topological space of subjective experience. *Trends in Cognitive Sciences*, *26*(12), 1068–1069. https://doi.org/10.1016/j.tics.2022.09.002

Thornton, I., & Fernandez-Duque, D. (2001). An implicit measure of undetected change. *Spatial Vision*, *14*(1), 21–44. https://doi.org/10.1163/156856801741341

Timmermans, B., & Cleeremans, A. (2015). How can we measure awareness? An overview of current methods. In M. Overgaard (Ed.), *Behavioral Methods in Consciousness Research* (pp. 21–46). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199688890.003.0003

Torgerson, W. S. (1958). *Theory and methods of scaling*. Wiley.

Tsuchiya, N., Phillips, S., & Saigo, H. (2022). Enriched category as a model of qualia structure based on similarity judgements. *Consciousness and Cognition*, *101*, 103319. https://doi.org/10.1016/j.concog.2022.103319

Wackermann, J., Putz, P., & Allefeld, C. (2008). Ganzfeld-induced hallucinatory experience, its phenomenology and cerebral electrophysiology. *Cortex*, *44*(10), 1364–1378. https://doi.org/10.1016/j.cortex.2007.05.003

Weisberg, J. (2023). *Explanatory optimism about the hard problem of consciousness* (1st ed.). Routledge. https://doi.org/10.4324/9781003411581

Wright, A. A. (2022). Change detection. In J. Vonk & T. K. Shackelford (Eds.), *Encyclopedia of Animal Cognition and Behavior* (pp. 1282–1292). Springer. https://doi.org/10.1007/978-3-319-55065-7_1590

Young, B. D., Keller, A., & Rosenthal, D. (2014). Quality-space theory in olfaction. *Frontiers in Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.00001