Consciousness and Continuity

ANDREW Y. LEE

University of Toronto, Philosophy

Abstract

Let a *smooth experience* be an experience with perfectly gradual changes in phenomenal character. Consider, as examples, your visual experience of a blue sky or your auditory experience of a rising pitch. Do the phenomenal characters of smooth experiences have continuous or discrete structures? If we appeal merely to introspection, then it may seem that we should think that smooth experiences are continuous. This paper (1) uses formal tools to clarify what it means to say that an experience is continuous or discrete, and (2) develops a discrete model of the phenomenal characters of smooth experiences. As a result, I'll argue that introspection leaves open whether smooth experiences are continuous or discrete. Yet I'll also argue—perhaps surprisingly—that the discrete theory may better fit our introspective evidence.

§1 Introduction

Philosophers sometimes say that the phenomenal characters of conscious experiences have continuous structures. A famous example is from Wilfred Sellars, who considers the visual experience of someone looking at a pink ice cube:

> The manifest ice cube presents itself to us as something which is pink through and through, as a pink continuum, all the regions of which, however small, are pink (Sellars 1963: 26).[1]

In fact, nothing as exotic as a pink ice cube is needed to illustrate the idea. Consider the phenomenal characters of the following kinds of experiences, which I'll call *smooth experiences*:

---

[1] In this passage, Sellars is speaking about the contents of the experience. But it's clear, from context, that he's also intending to make a claim about the structure of visual phenomenal character. In §2, I'll say more about how ascriptions of continuity to the contents of experiences relate to ascriptions of continuity to phenomenal characters.

Smooth Experiences
a. Your color experience of a blue sky on a cloudless day.
b. Your tactile experience when your whole hand is pressed against a surface.
c. Your auditory experience of a gradually rising pitch.
d. Your thermal experience as the room temperature gradually increases.

As a contrast class, consider another set of experiences, all of which exhibit abrupt changes in phenomenal character, which I'll call *gappy* experiences:

Gappy Experiences
a. Your color experience of a patch that's red on the left and green on the right.
b. Your tactile experience when only your fingertips are touching a surface.
c. Your auditory experience of C, then F#, then A♭ (and no notes in between).
d. Your thermal experience when half your body is submerged in hot water.

Let the *continuous theory* be the view that smooth experiences are continuous, and the *discrete theory* be the view that smooth experiences are discrete. Here's the main question of this paper: Does introspection favor either of these views?

At first, it may seem that introspection favors the continuous theory. After all, it's a solid empirical fact that introspection reveals no discontinuities in smooth experiences. And it's an analytic truth that to be continuous is to lack discontinuities. One could question the inference from these premises to the intended conclusion, on the grounds that the premises make a negative claim (about what introspection doesn't reveal) while the conclusion makes a positive claim (about which view is favored by introspection). But the proponent of the argument could counterargue that there's an equivalence between the relevant negative property (lacking discontinuities) and the relevant positive property (being continuous). I'll eventually argue that this line of reasoning can be resisted, but it will take some work to reach that point.

Furthermore, the alternative view—that smooth experiences are discrete— may strike some as phenomenologically inadequate. If the phenomenal characters of smooth experiences are discrete, then it seems to follow by definition that those experiences feel discrete. But smooth experiences don't feel discrete, so we might thereby infer that smooth experiences aren't discrete. A discrete theorist could

respond by appealing to limits in our introspective capacities: perhaps smooth experiences involve changes too small to be introspectively discernible. But while I have some sympathy for this move, I suspect that some will find it dialectically unsatisfying. The continuous theorist makes a positive claim that seems to align with the phenomenology. The discrete theorist, on the other hand, makes a negative claim that seems in tension with the phenomenology.

The goal of this paper is to argue that the discrete theory is viable. I'll (**1**) develop a discrete model of smooth experiences that can adequately account for the phenomenology. By doing so, I'll (**2**) argue that introspection leaves open whether smooth experiences have continuous or discrete structures. Along the way, I'll also (**3**) clarify what it means to say that a conscious experience is continuous or discrete. Unlike prior defenses of the discrete theory,[2] my arguments won't appeal to limits in our introspective capacities. Instead, I'll develop a *structural* (as opposed to epistemic) explanation of the difference between smooth and gappy experiences. And I'll even argue—perhaps surprisingly—that those most optimistic about our introspective capacities have reason to favor the discrete theory, rather than the continuous theory.

On the discrete model that I develop, smooth experiences are *contiguous*, where this means roughly that adjacent values of one experiential domain map to adjacent values of another experiential domain. Contiguity and continuity are incompatible: any structure that is contiguous must also be discrete. The initial goal of this paper will be to explain how to apply these formal concepts to the structure of phenomenology. After doing so, I'll argue that a contiguous model of smooth experiences adequately accounts for their phenomenology.

A methodological goal of this paper is to show how formal tools can elucidate philosophical questions about the structure of smooth experiences, as well as other philosophically relevant structural properties that I'll call 'gappiness', 'adjacency', and 'contiguity'. Most prior discussions of this issue have provided only cursory glosses of continuity and discreteness. But I'll argue that thinking about the subject-matter from a more formal perspective can advance our understanding of the core philosophical issues. Furthermore, once we start to examine the

---

[2] Clark [1989] is the most explicit endorsement of the discrete theory I've found. Fara [2001], like me, contends that introspection is agnostic as to whether smooth experiences are continuous or discrete, but offers an epistemic account of smoothness.

question more systematically, we will encounter a number of complexities that will shape our understanding of the basic philosophical questions.

Towards the end of the paper, I'll discuss the implications of my arguments for some issues about introspection. There are also implications for a number of other philosophical issues, including questions about the structures of quality-spaces, the representational contents of smooth experiences, the phenomenal sorites, and the physical correlates of consciousness. These other issues will be mentioned only in passing. Relatedly, although my focus is on the structures of conscious experiences, my discussions of continuity, discreteness, and other structural properties will be applicable to analogous debates about the structure of space, time, matter, and analog representations.[3]

Here's the plan for the rest of the paper. **§2** clarifies the target question; **§3** discusses the motivations for the continuous theory; **§4** and **§5** explain what it means to say that conscious experiences are continuous; **§6** defines 'contiguity'; **§7** develops an analysis of what makes an experience smooth versus gappy; **§8** and **§9** argue that all contiguous experiences are smooth; **§10** discusses implications for introspection; **§11** concludes.

§2   The Target Question

Here's our target question: What does introspection reveal about the structures of smooth experiences (and gappy experiences)?

The target question concerns the structure of the *phenomenal character* of smooth experiences. What phenomenal character *is* depends on which theory of consciousness one endorses. For *intentionalists*, phenomenal character is to be explained in terms of representational content. For *naïve realists*, phenomenal character is to be explained in terms of the properties of the external objects that one is perceptually aware of. For *sense-datum theorists* and *qualia theorists*, phenomenal

---

[3] For analogous debates about continuity, see Forrest [1995], van Bendegem [1995], Dummett [2000], and Builes & Teitel [2020] on space and time, Zimmerman [1996] on material objects, and Goodman [1968] and Maley [2011] on analog representations. For historical perspectives, see White [1992] on ancient theories of space, time, and motion and Fogelin [1988] on Hume and Berkeley's arguments against the infinite divisibility of space and time. For a general discussion of continuity vs. discreteness in mathematics, see Franklin [2017].

character is to be explained in terms of sense-data or qualitative states. But for any of these theories, we can ask whether phenomenal character—whatever it is—instantiates a continuous structure. If one sees a pink ice cube, does one visually represent the ice cube as continuous / is one perceptually aware of a continuous feature of the ice cube / is one's experience of the ice cube characterized by continuously structured sense-data or qualitative states?[4] I'll stay neutral on which of the above theories is correct: my arguments will be applicable to all of the above theories, at least once we translate into the relevant frameworks. I'll continue using language in such a way where I ascribe continuity and discreteness to experiences, though intentionalists and naïve realists may prefer to reinterpret these remarks as ascribing continuity or discreteness to what is presented in experience.

The target question should be distinguished from the question of whether the neural correlates of conscious experiences have continuous structures.[5] Consider, as an example, the distinction between (*a*) whether the neural correlates of conscious experiences persist continuously through time, versus (*b*) whether temporal phenomenology has a continuous structure. The former concerns the temporal structure of the neural correlates of experience; the latter concerns the structure of temporal phenomenology.[6] If we assume that the structures of conscious experiences must be isomorphic to the structures of their physical correlates, then an answer to one of these questions will constrain the answer to the other. But I

---

[4] I'm skeptical that smooth experiences represent continuity. Suppose, per reductio, that your experience of a pink ice cube represents the ice cube as continuous. Then it follows that your experience is veridical only if the ice cube is in fact continuous. But this would mean that the veridicality of smooth experiences is beholden to fundamental physics, since it may turn out that fundamental physics is discrete. Yet it seems that your experience of the pink ice cube is veridical so long as you see a pink ice cube, regardless of whether the ice cube is fundamentally continuous or discrete. Given this, I think we ought to instead hold that smooth experiences leave open whether their objects are continuous or discrete.

[5] For discussions on whether the physical correlates of experience are continuous, see Sellars [1963], Maxwell [1975], Dennett [1993], Lockwood [1993], VanRullen & Koch [2003], Blackmore [2002], Sergent & Dehaene [2004], White [2018], and Buildes [2020].

[6] For examples of views that dissociate these factors, see Lee, G [2014] and Phillips [2011].

won't appeal to any such constraints, since my principal concern is with what introspection reveals about the structure of smooth experiences.[7]

§3 The Continuous Theory

Earlier, I offered a preliminary argument in favor of the continuous theory: introspection reveals no discontinuities in smooth experiences, and to be continuous is to lack discontinuities, so introspection favors the continuous theory.

I think this basic line of reasoning captures the motivations behind many endorsements of the continuous theory. These sorts of claims occur commonly in discussions of temporal experience, such as when Dainton [2014: 130] says that temporal experiences are "infinitely divisible…phenomenal continua." They also occur in discussions of spatial experience, such as when Sellars [1963: 26] says that the phenomenal character of one's experience of a pink ice cube is as of "a pink continuum, all the regions of which, however small, are pink." Sometimes they occur in general discussions of the structures of conscious experiences, such as when Prentner [2019: 29] says that the "phenomenology of consciousness is such that it seems composed of an indefinite number of (phenomenal) parts." And even philosophers who reject the continuous theory tend to grant that such claims are appealing, such as when Clark [1989: 277] says that when looking at a sunset, "one seems to see a continuum of color" such that "between any two colored points…there seem to be other colored points."

A natural move for the discrete theorist is to say that there are limits to the grain of introspection. But the continuous theorist may worry that the discrete theorist is falling prey to a dubious analogy between introspecting an experience and perceiving a picture. A picture seen from far away might appear continuous, even though it turns out to be discrete upon close examination. But there seems no analogue of moving closer or further away in the case of introspection. This means that the standard method for explaining away the appearance of continuity is unavailable for the case of conscious experiences. From the standpoint of phenomenology, the continuous theory may seem to be on better grounds.

---

[7] Rashbrook [2013] distinguishes the following two claims: (1) the state of consciousness is continuous, and (2) the stream of experience is continuous. As far as I can tell, Rashbrook understands (1) as the claim that conscious experiences persist continuously through time, and (2) as the claim that temporal phenomenology is continuous.

In fact, the argument from introspection will strike many as more compelling than a structurally analogous argument from perception. Even if the sky appears continuous, we need not thereby believe that the sky is in fact continuous, since the way the sky perceptually appears to us can deviate from the way the sky actually is. On the other hand, if an experience of the sky appears continuous, then it's harder to dismiss the idea that the experience is in fact continuous. For perception, there's a distinction between the perceptual experience and the perceptual object. But for introspection, it seems that no analogous distinction is applicable.

For the purposes of this paper, it will be useful to focus on a more precise argument for the continuous theory. To my knowledge, the argument below has never been formulated explicitly in the philosophical literature. But I think it captures the phenomenological considerations that motivate the continuous theory, and it will take some work to appreciate how the argument goes awry:

⊥　　**The Argument for Continuity**
P1:　　Some experiences are smooth.
P2:　　Smooth experiences aren't gappy.
P3:　　An experience is either continuous or discrete.
P4:　　If an experience is discrete, then it's gappy.
——
C:　　Some experiences are continuous.

The argument is valid. Both P1 and P2 are uncontestable, since the terms 'smooth' and 'gappy' were defined ostensively: one could deny that there is any deep structural difference between smooth and gappy experiences, but one cannot deny that smooth experiences exist and that they are distinct from gappy experiences. Although P3 is false (for example, consider an experience that is locally continuous but globally discontinuous), let's restrict the quantifier to experiences that are either wholly continuous or wholly discrete. The premise I wish to challenge is P4. Call this the *discrete-implies-gappy premise*.

I'll eventually argue, contra this premise, that discrete experiences can be smooth. But before doing that, I need to first clarify what exactly it means to say that an experience is continuous or discrete.

§4 State-Spaces

To understand what it means to say that conscious experiences have continuous structures, we need to disambiguate two different interpretations of the claim:

Q1: Are the state-spaces for experiences continuous or discrete?

Q2: Are individual experiences continuous or discrete?

Since the main motivation for the continuous theory appeals to introspection, and since it's individual conscious experiences (rather than state-spaces) that are the objects of introspection, Q_2 is more directly relevant to the core aims of this paper. But to evaluate Q_2, we'll need to first understand Q_1. What, exactly, does it mean to say that the state-space for a given domain of experience is continuous?

A *state-space* is a structured set of the possible states that a system or object can be in. One of the most familiar examples is the state-space for color experience, which may be thought of as a three-dimensional space with hue, saturation, and brightness as dimensions, where color experiences that are more similar are located closer within the space. There are also state-spaces for auditory experience, olfactory experience, spatial experience, and any other feature of experience.[8] The state-space for color experiences has red_1, red_2, and so forth as its elements, and we might take the space to be structured by a distance metric that yields the result that red_1 is closer to red_2 than to red_{10}.

Our current question is what it means to say that a state-space is continuous or discrete. This is trickier than it may initially seem. The term 'continuous space' doesn't standardly occur within mathematics; instead, continuity is standardly understood as a property of functions (a point I'll return to in §5).[9] And while the term 'discrete space' *is* a standard mathematical term, we need to be careful about which kind of discrete space we attribute to the discrete theorist. Let me make some brief remarks to clarify these expressions.[10]

---

[8] See Clark [2000] for a more general discussion of state-spaces for features of experiences.

[9] I suspect the most natural mathematical analysis of 'continuous space' is in terms of a *connected manifold* (with dimension ≥ 1), or a space that's locally homeomorphic to Euclidean space at each point and that isn't the union of two disjoint non-empty subsets.

[10] See Franklin [2017] for more general discussion of continuity versus discreteness. Oftentimes, the distinction is illustrated by contrasting $\mathbb{R}$—the real numbers—and $\mathbb{Z}$—the

In philosophical contexts, the notion of a *continuous space* is usually treated as equivalent to the notion of an infinitely divisible space. Consider how debates about whether space, time, and matter are continuous or discrete are taken to turn on whether space, time, and matter are infinitely divisible into arbitrarily small spatial regions, temporal intervals, and material parts, or whether there are indivisible spatial, temporal, and material atoms. Given this, I'll follow convention and assume that the state-space for a domain of experiences is infinitely divisible just in case it's continuous.[11]

A *discrete space* is usually defined as a space where each element is *isolated*, meaning that for any element x, there is some distance δ where no distinct element y lies within distance δ from x. But the discrete theorist, in my view, ought to focus on a particular kind of discrete space: namely, *graphs*, or collections of elements and edges (where each edge connects two elements). Under this formalization, the distance between two elements is naturally understood as the number of edges of the shortest path connecting them. For the rest of the paper, I'll assume that discrete state-spaces are graphs. This will simplify some of the technical exposition, and will also be relevant to an objection against the discrete theory.

To streamline the discussion, I'll assume that all state-spaces under consideration are either wholly continuous or wholly discrete. For those interested in the more complex cases, it will be straightforward to generalize my arguments and analyses to state-spaces of arbitrary structure. For similar reasons, I'll assume that the state-spaces under discussion consist of all metaphysically possible

---

integers. In $\mathbb{R}$, every element is connected to every other element, there are continuum many elements between any two elements, and no element has an immediate predecessor or successor. In $\mathbb{Z}$, there are abrupt jumps from each element to the next, there are finitely many elements between any two elements, and every element has an immediate predecessor and successor. $\mathbb{R}$ is an example of a continuous structure; $\mathbb{Z}$ is an example of a discrete structure.

[11] Technically, infinite divisibility isn't sufficient for continuity: the set $\mathbb{Q}$ of rational numbers is infinitely divisible (any interval of $\mathbb{Q}$ contains multiple rationals) but discontinuous (every interval of $\mathbb{Q}$ is missing all the irrational numbers). However, questions about infinite divisibility lie at the heart of philosophical debates about continuity. To my knowledge, no contemporary philosopher has seriously argued that space or time or matter or consciousness is infinitely divisible yet discontinuous.

experiences of the relevant experiential domain, rather than merely the subset of experiences that are possible for a particular creature or a particular species.

§5   Individual Experiences

Our core question is whether smooth experiences, such as your visual experience of the blue sky on a cloudless day, have continuous or discrete phenomenal characters. To evaluate this question, we need to clarify what it means for an individual experience to be continuous versus discrete.

Let's start with a relatively trivial observation. For any individual experience α and any experiential feature F, there's a (possibly empty) set of values from the state-space for F-experiences that are instantiated by α. Suppose, for example, that you see a red gradient. Then your visual experience might instantiate values red_1–red_{100} from the state-space for color experiences. Let's call the set of values of the state-space for F-experiences that are instantiated by α the *F-values of α*.

Here's a hypothesis that's initially attractive but that turns out to be false: an experience α is continuous in feature F just in case α instantiates a continuous set of F-values, meaning α instantiates exactly the values within a continuous region of the state-space for F-experiences. In this circumstance, let's say that α satisfies the *continuum condition* with respect to F-experience.[12] In the example from above, if the state-space for color experience is continuous and if red_1–red_{100} is a continuous region of that state-space, then your experience satisfies the continuum condition with respect to color. The hypothesis says that what it is for an experience to be continuous (with respect to some feature) is for it to satisfy the continuum condition (with respect to that feature).

Now for a counterexample to the hypothesis. Suppose the state-space for color experience is continuous and suppose α instantiates every color quality within some continuous region of that state-space. However, suppose there is no systematic correspondence in α between color experience and spatial experience. You might imagine α as somewhat analogous to the kind of visual experience you have when looking at a noisy static image, such as a television screen with no signal. Even though α satisfies the continuum condition for color experience, α isn't

---

[12] FORMAL DEFINITION: Let $F[\alpha]$ be the set of F-values instantiated by α and d be the metric for the state-space for F-experiences. Then α satisfies the **continuum condition** with respect to F-experience $=_{def}$ $F[\alpha]$ is continuous under d (see fn. 9).

continuous in color experience. This example tells us something important about what it means to ascribe continuity to individuals.

In mathematics, continuity is normally understood as a property of functions, where *continuous functions* are those such that sufficiently small changes in inputs map to arbitrarily small changes in outputs. Putting it pictorially, continuous functions are those that can be drawn without lifting pen from paper, and that involve no "breaks, jumps, or wild oscillations."[13] At first, it may seem as though continuity is now being ascribed to fundamentally different kinds of things: functions by mathematicians and worldly things (such as space, time, matter, or experiences) by philosophers. However, the counterexample from the previous paragraph enables us to see how these two senses of 'continuity' come together.

If we ask whether an individual experience α is continuous, the question must be precisified as whether α is continuous in some feature F *with respect to some other feature* G. Even though α in the example above satisfies the continuum condition for color experience, it's nevertheless discontinuous in color experience with respect to spatial experience. If we reconsider the examples of smooth and gappy experiences from earlier in the paper, it's likewise easy to see which are the relevant feature F's and feature G's. Take the claim that your auditory experience of a rising pitch is continuous: it's clear that the relevant feature F is auditory experience and the relevant feature G is temporal experience.

Here's an interesting question: are there some features for which we can simply ask whether an experience α is continuous in feature F, without relativizing the question to a feature G? The best candidates are the features associated with "locative experiences," such as spatial and temporal experience.[14] A natural thought is that so long as α satisfies the continuum condition for spatial experience, α is continuous in spatial experience. Perhaps that is right (though I think

---

[13] Spivak [2008: 115]. FORMAL DEFINITION: a function f with domain X is *continuous* $=_{\text{def}}$ for any point $a \in X$, $\forall \epsilon > 0$, $\exists \delta > 0$ such that $\forall x \in X$, if $0 < |x - a| < \delta$ then $|f(x) - f(a)| < \epsilon$. Note that a more general definition of continuity is topological, where continuous functions are those where the pre-images of open sets are open. But I've chosen to focus on the metrical definition above both because it's more accessible and because it's plausible that the relevant state-spaces have metric structure.

[14] See Clark [2000: Ch. 3] on locative experiences.

the answer isn't entirely unobvious).[15] But since addressing such questions would take us away from the main goals of this paper, I'll restrict my focus to the cases where we must ask whether an experience is continuous in F with respect to G. If there turn out to be simpler cases where we need not appeal to any feature G, then it will be clear how to generalize my arguments to those cases.

Here's another interesting question: can any feature of experience play either the F-role or the G-role? Well, it's natural to take *qualitative features* (such as color and auditory experience) to play the F-role and *locative features* (such as spatial and temporal experience) to play the G-role. It's easy to grasp what it means for an experience to be continuous in color with respect to space; it's hard to grasp what it means for an experience to be continuous in space with respect to color. I suspect this intuitive difference arises from structural differences between qualitative features and locative features.[16] Given this, I'll sometimes talk of F-*values* being instantiated at G-*locations*. However, even those who think the F-role and the G-role are restricted to particular kinds of experiences will be able to accept my arguments.

We are now in position to see why the function-theoretic notion of continuity used by mathematicians is relevant for the question of whether individual experiences are continuous. To say that α is continuous in F with respect to G is to say that the mapping from α's G-values to α's F-values is a continuous function, where this means that sufficiently small changes in α's G-values map to arbitrarily small changes in α's F-values. This enables us to precisify our initial question: the question now is whether smooth experiences are such that sufficiently small changes in one feature (such as color experience) map to arbitrarily small changes in another feature (such as spatial experience).

There remains one last complication. Any function with a discrete domain trivially satisfies the mathematical definition of continuity. But this means that some features of individual experiences that are intuitively not continuous will

---

[15] Suppose α represents colors at spatial regions l_1–l_{10} and l_{20}–l_{30} (with no visual experience at l_{10}–l_{20} or l_{30}–l_{40}) and sounds at l_{10}–l_{20} and l_{30}–l_{40} (with no auditory experience at l_1–l_{10} and l_{20}–l_{30}). Then α satisfies the continuum condition for spatial experience. But given that α is the conjunction of a spatially disconnected visual experience and a spatially disconnected auditory experience, is it correct to say that α is continuous in spatial experience?

[16] See Clark [2000: Ch.3] for discussion.

count as continuous given our present analysis. Suppose that the state-space for spatial experience is discrete and that α represents red at spatial location l_1, green at l_2, and blue at l_3. It would be bizarre to say that α is continuous in color experience with respect to spatial experience. Yet it turns out, given the definition of a continuous function, that the function mapping α's G-values to α's F-values is continuous.

Fortunately, we've already encountered the tool that's needed to solve this problem. We need simply require that α satisfies the continuum condition for *feature G*, meaning that α's G-values are exactly those within a continuous region of the state-space for G-experiences. In other words, the domain of the function from G-values to F-values must be continuous. This not only solves the technical problem described above, but also forges a neat connection between continuity of individual experiences and continuity of state-spaces: in order for α to be continuous in F with respect to G, the state-space for G-experiences must itself be continuous.[17] And with that condition, we arrive at the following definition ('*wrt*' means 'with respect to'):

experience α is **continuous** in feature F *wrt* feature G $=_{def}$[18]

- α instantiates a continuous set of G-values.
- sufficiently small changes in α's G-values map to arbitrarily small changes in α's F-values.

The remaining task is to define what it is for an individual experience α to be discrete in feature F with respect to feature G. There's no standard definition of 'discrete function', but a common characterization is that a discrete function is

---

[17] By contrast, the state-space for F-experiences needn't be continuous. Suppose that the state-space for color experience is discrete but that the state-space for spatial experience is continuous, and suppose α is an experience as of looking at a uniformly red wall, where red_1 is represented at spatial locations l_1–l_{100}. Then α is continuous in color experience with respect to spatial experience, even though the former state-space is discrete.

[18] FORMAL DEFINITION: Let $G[\alpha]$ be the set of G-values instantiated by α that map onto F-values, and let f_i be the F-value mapped by any $g_i \in G[\alpha]$. Then α is **continuous** in F with respect to G $=_{def}$ **(1)** α satisfies the continuum condition with respect to G-experience (see fn. 12), and **(2)** $\forall \epsilon > 0$ and $\forall g_a \in G[\alpha]$, $\exists \delta > 0$ such that $\forall g_n \in G[\alpha]$, if $d(g_a, g_n) < \delta$ then $d(f_a, f_n) < \epsilon$.

simply a function with a discrete domain. This turns out to be almost exactly what we need for characterizing discreteness as a property of individual experiences. The only caveat concerns an edge case: we need to require that the domain of the function be non-empty. This ensures that our definition of discreteness doesn't overgeneralize: for example, we wouldn't want α to trivially count as discrete in gustatory experience with respect to emotional experience simply because there is no mapping from α's emotional experience values to α's gustatory experience values. With this condition, we can define discreteness of individual experiences as follows: [19]

experience α is **discrete** in feature F *wrt* feature G =def[20]

- α instantiates a discrete set of G-values.
- there are some G-values in α that map to F-values in α.

But what happens if the mapping from α's G-values to F-values isn't even a function? Well, this situation will arise only if multiple F-values are associated with the same G-value. If *only* locative experiences can play the G-role (and if each locative value can be instantiated only once per experience), then *every* mapping from G-values to F-values will be a function. But suppose, to take an example, that F is spatial experience, that G is color experience, and that red_1 maps to both l_1 and l_2. Then, according to the analysis, α *isn't* continuous in F with respect to G (though it may still be true that α is continuous in color experience with respect to spatial experience). This is intuitively correct.

---

[19] A technical remark: α can be discrete in F with respect to G even if the state-spaces for both F-experiences and G-experiences are continuous. Suppose, for example, that both color experience and spatial experience have continuous state-spaces, and that α instantiates red_1 at l_1, red_2 at l_2, red_3 at l_3, and no other color values or spatial values. Then α is discrete in color experience with respect to spatial experience, even though the relevant state-spaces are continuous. As an analogy, consider how any function f from $\mathbb{Z}$ to $\mathbb{R}$ is discrete, even though $\mathbb{R}$ is continuous. In this case, f's domain (i.e., $\mathbb{Z}$) is analogous to α's G-values, f's codomain (i.e., $\mathbb{R}$) is analogous to the state-space for F-experiences, and f's image (i.e., the elements of $\mathbb{R}$ that are outputs of f) is analogous to α's F-values.

[20] FORMAL DEFINITION: Let $G[\alpha]$ be the set of α's G-values that map to F-values. α is **discrete** in F with respect to G =def **(1)** $\forall x \in G[\alpha], \exists \epsilon > 0 : \forall y \in G[\alpha] \setminus \{x\}, d(x, y) > \epsilon$, and **(2)** $G[\alpha] \neq \emptyset$.

Now that we have these basic concepts in place, we can turn to my discrete model of the structure of smooth experiences.

§6   Contiguity

To develop my argument against the discrete-implies-gappy premise, I'll first define a new structural property, which I'll call *contiguity*. Contiguity is, basically, a matter of adjacent values of one domain mapping to adjacent values of another domain. Consider a contrast between two kinds of sequences of integers—the A-sequences below are contiguous, while the B-sequences are not:

A_1:	$(1, 2, 3, 4, \ldots)$	B_1:	$(1, 3, 5, 7, \ldots)$
A_2:	$(1, 1, 1, 1)$	B_2:	$(1, 1, 1, 3)$
A_3:	$(3, 2, 1, 0, 1, 2, 3)$	B_3:	$(3, 0, 2, 1, 3, 1, 2)$

Let's say two integers a and b are *adjacent* just in case either $a = b$ or $a = b \pm 1$. The A-sequences are contiguous because every subsequent integer is adjacent to its predecessor in the sequence. The B-sequences are discontiguous because some intermediate integers are missing: for example, sequence B_1 jumps from 3 to 5. The notion of a sequence may initially seem distinct from the formal tools we have invoked so far, but notice that sequences are simply functions in disguise: any sequence may be thought of as a function whose domain is the natural numbers and whose image is the values of the sequence. If we think of sequences in this way, then it's easy to see that for contiguous sequences of integers, adjacent G-values (the indices of the sequence) map to adjacent F-values (the integer at a given index).

The notion of contiguity can be generalized. In a discrete state-space, two elements x and y are *adjacent* just in case they are connected by a single edge. If there is a sequence of edges that connects x and y, then there is a path from x to y. Let's say a *contiguous region* of a state-space is a region R where for any two elements x and y of R, there is a path wholly within R connecting x and y. Put another way, contiguous regions are those where we can move from any value to any other via a sequence of adjacency pairs. Finally, let's say experience α *instantiates a contiguous set* of G-values just in case α instantiates all and only the elements within a contiguous region of the state-space for G-experiences. With these definitions in place, we can define contiguity in a way that parallels the prior definition of continuity:

experience α is **contiguous** in feature F *wrt* feature G $=_{def}$[21]

- α instantiates a contiguous set of G-values.
- adjacency in α's G-values corresponds to adjacency in α's F-values.

In the examples below, α_1 and α_2 are contiguous (and β_1 and β_2 are discontiguous) in color experience with respect to spatial experience (assuming that the state-spaces for both color experience and spatial experience are discrete):[22]

α_1:   an experience of a spectrum of color, representing red_1–red_{50} at locations l_1–l_{50} (respectively).

α_2:   an experience of a homogenous field of color, representing red_1 at locations l_1–l_{50}.

β_1:   an experience of a series of colors with every other color skipped over, representing red_1 at location l_1, red_3 at l_2, red_5 at l_3, and so forth.

β_2:   an experience of a spectrum of colors with a gap, representing red_1–red_{50} (except for red_{37}) at locations l_1–l_{50} (except for l_{37}).

Contiguity and continuity are mutually exclusive.[23] To be contiguous, an experience must instantiate a contiguous region of the state-space for G-experiences. But that requires the state-space for G-experiences to be discrete, since there are no adjacent elements in continuous spaces. To be continuous, an experience must instantiate a continuous region of the state-space for G-experiences. But that

---

[21] FORMAL DEFINITION: Let G[α] be α's set of G-values that map onto F-values, and let f_i be the F-value mapped by any $g_i \in$ G[α]. Then α is **contiguous** in F with respect to G $=_{def}$ **(1)** G[α] is contiguous (fn. 20), and **(2)** $\forall g_1, g_2 \in$ G[α], if g_1 is adjacent to g_2, then f_1 is adjacent to f_2.

[22] Lee, A [2021] argues that the qualities of experiences correspond to regions, rather than individual elements, in their respective state-spaces. This complicates the analyses of both continuity and contiguity, since it's not obvious how to generalize those notions to such a framework. But I suspect it will be possible to construct degree-theoretic analogues of continuity and contiguity over regions.

[23] The sole exception is the degenerate case where α instantiates only one G-value and one F-value: such an experience is both continuous and contiguous in F with respect to G.

requires the state-space for G-experiences to be continuous, since the elements in discrete spaces are isolated from one another.

Now for one of my central claims: any experience that is contiguous is smooth. If that claim is true, then the discrete-implies-gappy premise is false, and the Argument for Continuity is unsound.

§7   The Analysis of Smoothness

To argue that contiguous experiences are smooth, I need to first develop an analysis of what makes any experience smooth versus gappy. Let's start with gappiness. As a reminder, gappy experiences include the visual experience you have when looking at a pixelated low-resolution screen, your auditory experience of C, then F#, then Ab, and your tactile experience when your fingertips are spread out and pressed against a surface. What do such experiences have in common?

It would be inadequate to say that all gappy experiences are discrete. Although that claim may in fact be true, it fails as an analysis of gappiness. For the discrete theorist, such a claim is trivial, since all experiences are discrete. For the continuous theorist, such a claim is false, at least assuming that the relevant state-spaces are continuous. A continuous theorist is likely to think instead that gappy experiences are globally discontinuous but locally continuous. As an analogy, consider $\mathbb{R} \setminus \mathbb{Z}$, the set of real numbers minus the set of integers, which is gappy (it's missing all the integers) but not discrete (its elements aren't all isolated from one another).

Here's a more promising hypothesis: gappy experiences are experiences that are missing intermediate F-values at intermediate G-locations. This hypothesis is intuitive. Your gappy visual experience is missing intermediate color values between adjacent pixels, your gappy auditory experience is missing C#, D, and D# at the relevant times, and your gappy tactile experience is missing tactile sensations for the spatial locations between your fingertips. Here's a general statement of this idea (the first condition precludes degenerate cases, where an experience instantiates only a single G-value, from counting as gappy):

experience α is **gappy** in feature F *wrt* feature G $\leftrightarrow$
- α instantiates non-adjacent values g_1 and g_2 mapping to f_1 and f_2.
- α is missing intermediate F-values at intermediate G-locations.

The natural corollary hypothesis is that smooth experiences are those where every intermediate G-value maps to an intermediate F-value. Here's a general statement of this idea (as before, the first condition precludes degenerate experiences that instantiate only a single G-value from counting as smooth):

experience α is **smooth** in feature F *wrt* feature G $\leftrightarrow$

- α instantiates non-adjacent values g_1 and g_2 mapping to f_1 and f_2.
- α instantiates all intermediate F-values at intermediate G-locations.

To precisify this, we can appeal to the notion of a *path*. In a continuous space, a path is a continuous function from the interval [0, 1] to elements in the space. In a discrete space, a path is a sequence of connected edges (or, equivalently, of pairs of adjacent elements). Either way, here's what it is for an experience α to instantiate all intermediate F-values at intermediate G-locations: every path-connected pair of α's G-values maps to a path-connected pair of α's F-values. If that condition isn't satisfied, then α is missing some intermediate values at intermediate locations. These analyses correctly categorize the smooth and gappy experiences mentioned in §2. And since both the continuous theorist and the discrete theorist can make sense of the idea of missing intermediate values, these analyses are neutral between the continuous theory and the discrete theory.[24]

Now for an important result: *both* continuity and contiguity satisfy the analysis of smoothness. If the relevant state-spaces are continuous, then smoothness is a matter of continuity, since in these cases it's all and only continuous experiences that instantiate all intermediate F-values at intermediate G-locations. If the relevant state-spaces are discrete, then smoothness is a matter of contiguity, since in these cases it's all and only contiguous experiences that instantiate all intermediate F-values at intermediate G-locations. But contiguous experiences are discrete. This means that discrete experiences can be smooth, which means that the discrete-implies-gappy premise is false, which means that the Argument for Continuity is unsound.

---

[24] Notice that these analyses are expressed as biconditionals, rather than as definitions. The terms 'smooth' and 'gappy' were defined by ostension to cases; the analyses are substantive hypotheses about what the experiences we labeled as 'smooth' and 'gappy' have in common.

A point worth highlighting is that my analysis of smoothness is *structural*, rather than epistemic. Those who have expressed sympathy towards the discrete theory usually account for smooth experiences by appealing to limits in our introspective capacities. In particular, one might think that smooth experiences are those where differences in phenomenal character between adjacent values are too small for subjects to introspectively discern. Although I'm sympathetic to introspective limitations playing a role in explaining why a given experience strikes its subject as smooth, I think that smoothness itself is best understood in structural terms. An epistemic analysis like the one above yields the counterintuitive result that whether an experience is smooth or gappy is subject-relative, and that idealized subjects with perfect introspective capacities cannot have smooth experiences. By adopting a structural analysis of smoothness, we avoid those sorts of results.

In the remainder of the paper, I'll discuss some objections to my arguments.

§8   The Feels-Discrete Objection

Here's the first objection:

⊥   **The Feels-Discrete Objection**
P1:   If an experience is discrete, then it feels discrete.
P2:   If an experience feels discrete, then it's gappy.
—
C: If an experience is discrete, then it's gappy.

The argument is obviously valid, and each premise seems compelling on its own. But the argument equivocates. Although there is a reading of P1 that is true and a reading of P2 that is true, there is no univocal precisification of 'feels discrete' that renders both premises true. The equivocation is between two senses of 'feels φ': (**1**) an experience *phenomenally* feels φ just in case it has phenomenal character φ, and (**2**) an experience *epistemically* feels φ just in case it strikes its subject as φ.[25]

---

[25] Some candidates for what it is for an experience α to strike one as φ include one being disposed to believe that α is φ (see Werner [2014]), one having the intuition that α is φ (see Bengson [2015]), or α having presentational phenomenology that φ (see Chudnoff [2012]). Note that if α does *not* strike one as φ, then that doesn't necessarily mean that α strikes one as not φ, nor that one isn't in a position to know that α is φ. And striking one as φ might

Here's an example designed to disentangle the two senses. Let a visual experience be *colored* just in case it instantiates some color qualities, and *prime-colored* just in case it instantiates a prime number of distinct color qualities. Suppose your visual experience instantiates 743 distinct color qualities. Then your visual experience is both colored and prime-colored. But while your visual experience strikes you as colored, it doesn't strike you as prime-colored (nor as not prime-colored). Perhaps you can know that your visual experience is prime-colored if you introspect carefully enough. But that doesn't mean that the visual experience strikes you as prime-colored. As an analogy, consider how you can know that π is a transcendental number if you think carefully enough, even though π doesn't strike you as a transcendental number. The sentence 'your visual experience feels colored' is true under both readings of 'feels' but the sentence 'your visual experience feels prime-colored' sentence is false under the epistemic reading.

P1 says that if an experience is discrete, then it feels discrete; P2 says that if an experience feels discrete, then it's gappy. If 'feels discrete' is interpreted in the phenomenal sense, then P1 is true but P2 is unobvious. If 'feels discrete' is interpreted in the epistemic sense, then P2 is true but P1 is unobvious. Here's the missing premise needed to secure the objection: if an experience phenomenally feels discrete, then that experience epistemically feels discrete. In other words, the feels-discrete objection tacitly appeals to the assumption that any experience that is discrete must strike its subject as discrete. Since P1 and P2 are each plausible after the disambiguations mentioned above, this missing premise in effect says that if an experience is discrete, then it's gappy. Yet this is exactly the discrete-implies-gappy premise, which I've already argued against.

The appeal of the feels-discrete objection might come from our tendency to assume that what it's like to have a discrete experience is structurally similar to what it's like to imagine an experience as discrete. If you're asked to imagine an experience as discrete, then you might imaginatively represent the discrete experience using a mental image of a pixelated image, where individual pixels correspond to the discrete units of the target experience. That imaginative experience itself may well be gappy. But just because the experience of imagining an experience as discrete is gappy doesn't mean that the discrete experience that is imagined

---

well be a matter of degree, in which case I'll assume that sentences of the form 'α (epistemically) feels φ' are true just in case α strikes one as φ to a sufficiently high degree.

is itself gappy. That inference would conflate the structure of the vehicle used to represent a target experience with the structure of the target experience itself. The gappiness is a feature of the imaginative experience, rather than of the experience imagined.

A similar point applies to the objection that if the distances between contiguous values are sufficiently high, then contiguous experiences would feel discrete. I suspect that this objection is motivated by an inadequate analogy. If you look at the individual pixels in an image, you may notice discontinuities in color when moving from one pixel to the next: perhaps one pixel is red$_{17}$ and the adjacent pixel red$_{54}$. It may be tempting to think of contiguous values in a state-space as structurally analogous. But observe that you are able to notice the discontinuities across adjacent pixels only because you are able to perceptually represent the color values that are missing between pixels—that is, some of the values between red$_{17}$ and red$_{54}$. If you weren't able to perceptually represent any of those missing color values, then it's unobvious that the pixels would look discontinuous to you. To notice a gap, one's cognitive system must be sensitive to the values that would fill in the gap. If there are no such values, then one cannot notice the gap.

§9   The Discontinuity Objection

Here's the second objection:

⊥   **The Discontinuity Objection**
P1:   If an experience is discrete, then it's discontinuous.
P2:   If an experience is discontinuous, then it contains gaps.
P2:   If an experience contains gaps, then it's gappy.
—
C: If an experience is discrete, then it's gappy.

As before, the argument is valid, and each premise seems compelling: P1 and P2 seem to follow from the definitions of 'discrete' and 'discontinuous', and P3 sounds tautological. The force of the argument turns on how exactly we interpret the expression 'contains gaps', and whether it in fact entails the our sense of 'gappy'.

Here's my analysis from earlier of what it is for experience to be gappy: experience α is gappy in feature F with respect to feature G just in case some

intermediate F-values are missing at intermediate G-locations. Since contiguous experiences map adjacent values to adjacent values, and since there are no intermediate values between adjacent values, contiguous experiences aren't gappy. Supposing we interpret P2 as a tautology, it follows that P1 false: just because an experience is discrete doesn't mean it contains gaps, since discrete experiences need not be missing any intermediate values.

Now, there are of course other ways of defining 'gappy', and some of these alternate definitions would render P1 true. Let's call the analysis of 'gappy' I favor the *missing-values definition*. The natural alternative is the *discontinuity definition*, which says that α is gappy in F with respect to G just in case α is discontinuous in F with respect to G. These two definitions yield different diagnoses of contiguous structures: only the discontinuity definition says that contiguous structures are gappy. It's obvious that we are now in the vicinity of a verbal dispute. But we can avoid that trap by focusing on the substantive question: which definition best captures the class of experiences I originally labeled 'gappy'? I'll argue that the discontinuity definition yields plausible results only if we already presume that the relevant structures are continuous.

Suppose experience α is contiguous in color experience with respect to spatial experience and that red_1 and red_2 are adjacent color values instantiated by α. Is α gappy? In other words, does α belong to the same class of experiences as those that were labeled 'gappy' at the beginning of the paper? In every example we have encountered of a gappy experience, the experience is missing intermediate values along the relevant state-spaces. But α doesn't have this feature since there are no color experiences between red_1 and red_2. This is reason to think that α isn't the kind of experience that would strike its subject as gappy, which is evidence that α isn't gappy. To hold otherwise, one would have to say that α is gappy even though it's impossible to "fill in" that gap. Since only the missing values definition classifies α as gappy, we thereby have reason to favor it over the discontinuity definition.

The advantage of the missing-values definition becomes more obvious when we consider non-experiential gaps. Suppose the physical world turns out to be fundamentally discrete. Then, according to the discontinuity definition, every physical structure contains gaps (since no physical structures are continuous). But there is clearly still a sense in which we can talk about some physical things containing gaps and other physical things lacking gaps. A wall that is half black and

half white is gappy (in color with respect to space), while a wall that is uniformly black is not; a flickering light is gappy (in light with respect to time), while a constant light is not. The missing-values definition works whether the target structures are continuous or discrete; the discontinuity definition does not.

Here's one more way of illustrating the point. Suppose you have non-zero credence that the discrete theory is true but also adopt the discontinuity definition. Since the discontinuity definition entails that all discrete experiences are gappy, it follows that any credence you have in the discrete theory should also give you credence that every experience—including all the experiences we labeled 'smooth'—are in fact gappy. But if we were to classify *every* experience as gappy, then we will have lost sight of the initial explanandum of this paper: namely, explaining the difference between the class of experiences I called 'smooth' and the class of experiences I called 'gappy'. Moreover, we would still be able to invent new terms to distinguish between smooth and gappy experiences and once again ask what differentiates the two classes. On the other hand, if we adopt the missing-values definition, then we can distinguish between smooth and gappy experiences, no matter which theory turns out to be true.

§10   Introspection

At the beginning of the paper, I noted that it's tempting to think that introspection favors the continuous theory over the discrete theory. My main goal, over the course of the paper, has been to argue that the discrete theory is compatible with our introspective evidence.[26] I'll now argue—perhaps surprisingly—that introspection might actually favor the discrete theory over the continuous theory.

To begin, note that the discrete theory is more conservative than the continuous theory in extrapolating beyond our introspective evidence. This is because the discrete theory is compatible with ascribing no more structure to experience than what's needed to account for the introspective data, while the continuous theory entails that there are infinitely many more elements of our experiences than

---

[26] Let 'LUMINOSITY' be the following principle: if one's experience is F (where F is a phenomenal property), then one can know via introspection that one's experience is F). If introspection leaves open whether smooth experiences are continuous or discrete, then LUMINOSITY is false.

what introspection reveals. In fact, if the continuous theory is correct, then we can introspect exactly 0% of our phenomenal qualities!

To see why, start by noting that our introspective capacities are finite: for any given experience, one can introspectively discriminate only a finite number of distinct qualities within that experience. If the continuous theory is correct, then non-homogenous smooth experiences (meaning smooth experiences that instantiate multiple F-values) must instantiate infinitely many distinct qualities. This is because if the state-space for F-experiences is continuous, then between any two F-values there are infinitely many other F-values. By contrast, if the discrete theory is correct, then non-homogenous smooth experiences need instantiate only a finite number of distinct qualities. This is because if the state-space for F-experiences is discrete, then between any two F-values there are only finitely many other F-values.

The point above is dialectically significant because it undercuts a potential criticism of the discrete theory. Most who have favored the discrete theory have appealed to the idea that smooth experiences involve changes in phenomenal character that are too small to be introspectively discernible. This move might be criticized on the grounds that it requires ascribing more structure to our experience than what introspection reveals. But it turns out that the continuous theory faces the same cost. In fact, the continuous theory does worse than the discrete theory on this dimension.

Furthermore, the discrete theorist could contend that every adjacent value in a smooth experience is introspectible, but that the experience nevertheless feels smooth because those values are contiguous. In other words, if F_1 and F_2 are adjacent values (and are contiguous with respect to some instantiated G-locations), then they might be introspectively discernible yet still belong to a smooth experience. Because of this, the discrete theorist needn't necessarily ascribe more structure to our experiences than what's introspectively discernible.

This is a striking turnaround. At the start of the paper, it seemed that introspection was on the side of the continuous theorist. I've now argued that it's discrete theorist who is able to minimize the gap between the structures of experience and the grain of introspection. This doesn't necessarily mean that the discrete theory has the theoretical advantage. But it does mean that those optimistic about the scope of introspective knowledge ought to refrain from the temptation to conclude

that smooth experiences are continuous, and ought to instead take seriously the hypothesis that smooth experiences are discrete.

Conclusion

I started with a contrast between smooth and gappy experiences. The continuous theorist says that smooth experiences are continuous; the discrete theorist says smooth experiences are discrete. I argued that what it is for an experience α to be continuous in feature F with respect to feature G is for α to instantiate a continuous set of G-values and for sufficiently small changes in α's G-values to map to arbitrarily small changes in α's F-values. Then I presented the Argument for Continuity, where the key premise claimed that if an experience is discrete, then it's gappy.

From there, my goal was to explain why smooth experiences need not be continuous. To do this, I first defined 'contiguity', where α is contiguous in F with respect to G just in case α instantiates a contiguous set of G-values and adjacency in α's G-values maps to adjacency in α's F-values. Then I argued that α is smooth in F with respect to G just in case α isn't missing any intermediate F-values at the relevant G-locations. This definition of smoothness is satisfied by both continuous and contiguous experiences. Since contiguous experiences are discrete, it follows that some discrete experiences are smooth, contradicting the Argument for Continuity. This means that both the discrete theory and the continuous theory can explain the phenomenological differences between smooth and gappy experiences.

For some readers, there may remain the residual feeling that the discrete theory cannot do justice to our phenomenology. To ensure that our intuitions are clear, it's worth briefly reiterating some of the upshots from earlier. First, 'continuous' doesn't merely mean 'smooth'—instead, it denotes the structural property that was defined earlier and that's deployed in mathematics, science, and other areas of philosophy. Second, the fact that an experience is discretely structured doesn't automatically entail that the experience will strike its subject as discretely structured. Third, what it's like for one to imagine an experience as discrete need not be what it's like to actually undergo the discrete experience that is imagined. Fourth, all known cases of an experience that strikes its subject as discrete are cases where the experience is missing some intermediate values. Speaking for myself,

once I recognize these facts, I lose any intuition that the discrete theory cannot be phenomenologically adequate.[†]

References

Bayne, Tim ; Hohwy, Jakob & Owen, Adrian M. (2016). Are There Levels of Consciousness? *Trends in Cognitive Sciences* 20 (6):405-413.

Bengson, John (2015). The Intellectual Given. *Mind* 124 (495):707-760.

Blackmore, S. (2002), 'There Is No Stream of Consciousness', Journal of Consciousness Studies, 9: 17–28.

Builes, David (2020). Derivatives and Consciousness. Journal of Consciousness Studies 27 (9-10):87-103.

Builes, David & Teitel, Trevor (2020). A puzzle about rates of change. Philosophical Studies 177 (10):3155-3169.

Byrne, Alex (2005). Introspection. *Philosophical Topics* 33 (1):79-104.

Chudnoff, Elijah (2012). Presentational Phenomenology. In Miguens & Preyer (eds.), Consciousness and Subjectivity. Ontos Verlag.

Clark, Austen (1989). The particulate instantiation of homogeneous pink. *Synthese* 80 (August):277-304.

Clark, Austen (2000). *A Theory of Sentience*. Oxford University Press.

Dainton, B. (2006), *Stream of Consciousness: Unity and Continuity in Conscious Experience*. London: Routledge.

Dainton, Barry (2014). 'The Phenomenal Continuum'. In V. Arstila & D. Lloyd, eds., pp. 101–37.

Dennett, D. C. (1993), Consciousness Explained. London: Penguin.

Dummett, Michael (2000). Is time a continuum of instants? *Philosophy* 75 (4):497-515.

Fara, Delia Graff (2001). Phenomenal continua and the sorites. Mind 110 (440):905-935.

Forrest, Peter (1995). Is space-time discrete or continuous? — An empirical question. *Synthese* 103 (3):327--354.

Franklin, James (2017). Discrete and continuous: a fundamental dichotomy in mathematics. *Journal of Humanistic Mathematics* 7 (2):355-378.

Goodman, Nelson (1968). *Languages of Art*. Bobbs-Merrill.

Hurevicz, W. & Wallman, G. (1948) Dimension theory, Princeton University Press.

Kammerer, François (2016). The hardest aspect of the illusion problem - and how to solve it. *Journal of Consciousness Studies* 23 (11-12):124-139.

Lee, Andrew Y. (2021). Modeling Mental Qualities. *Philosophical Review*.

Lee, Geoffrey (2014). Temporal Experience and the Temporal Structure of Experience. *Philosophers' Imprint* 14.

Lockwood, Michael (1993). The grain problem. In Howard M. Robinson (ed.), Objections to Physicalism. Oxford University Press. pp. 271-291.

Maley, Corey J. (2011). Analog and digital, continuous and discrete. *Philosophical Studies* 155 (1):117-131.

Maxwell, Grover (1978). Unity of consciousness and mind-brain identity. In John C. Eccles (ed.), Mind and Brain. Paragon House.

Prentner, Robert (2019). Consciousness and topologically structured phenomenal spaces. Consciousness and Cognition 70:25-38.

Rashbrook, Oliver (2013). The Continuity of Consciousness. European Journal of Philosophy 21 (4):611-640.

Schwitzgebel, Eric (2006). The unreliability of naive introspection. Philosophical Review 117 (2):245-273.

Sellars W. 1963, 'Philosophy and the Scientific Image of Man', in Science, Perception, and Reality, Routledge & Kegan Paul, London, pp. 1-40.

Sergent, Claire & Dehaene, Stanislas (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science* 15 (11):720-728.

Spivak, Michael (1967, 4th ed. 2008). Calculus.

Tye, M. (2003), *Consciousness and Persons*. Cambridge, MA: MIT Press.

van Bendegem, Jean Paul (1995). In Defence of Discrete Space and Time. *Logique Et Analyse* 38 (150-1):127-150.

VanRullen, Rufin & Koch, Christof (2003). Is perception discrete or continuous? *Trends in Cognitive Sciences* 7 (5):207-213.

Werner, Preston J. (2014). Seemings: still dispositions to believe. *Synthese* 191 (8):1-14.

White, Michael J. (1992). *The Continuous and the Discrete: Ancient Physical Theories From a Contemporary Perspective*. Oxford University Press.

White, Peter A. (2018). Is conscious perception a series of discrete temporal frames? Consciousness and Cognition 60:98-126.

Wright, Crispin (1975). On the coherence of vague predicates. Synthese 30 (3-4):325--65.s