
Consciousness & 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. Along the way, I explain why the sense of 'continuity' ascribed to smooth experiences is distinct from the sense of 'continuity' ascribed to the stream of consciousness.

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*:

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 Ab (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. Here's a simple argument for thinking so:

¹ In this passage, Sellars is speaking primarily about the content of perceptual experience. But this example is standardly interpreted as also making a claim about the character of visual phenomenology. In §2, I'll say more about how ascriptions of continuity to the contents of experience relate to ascriptions of continuity to phenomenal character.

⊥ **A Simple Argument**

P1: Introspection reveals no discontinuities in smooth experiences.

P2: To be continuous just is to lack discontinuities.

—

C: Introspection favors the continuous theory.

This argument, though simple, is compelling. P1 expresses a solid empirical fact, and P2 expresses an analytic truth. You could question the inference from the premises to the 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 simple argument is flawed, 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. Though I have some sympathy for this move, others may find it dialectically unsatisfying. The continuous theorist makes a positive claim that seems to align with the phenomenology of smooth experiences. The discrete theorist, on the other hand, makes a negative claim that seems in tension with the phenomenology.

The goals of this paper are to (1) explain how to apply some general structural concepts—especially CONTINUITY and DISCRETENESS—to conscious experiences, and to (2) develop a discrete analysis of smooth experiences that can adequately account for the phenomenology. By consequence, I'll argue that introspection leaves open whether smooth experiences have continuous or discrete structures. Unlike some defenses of the discrete theory, 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. 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 analysis 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.

To some readers, the distinction between continuity and contiguity may feel subtle. But—as I'll explain more later—the distinction between the continuous and the discrete is one of the most significant joints amongst mathematical structures. The most philosophically interesting aspect of the distinction is that continuity entails *infinite divisibility* (every finite region of a continuous structure is infinitely divisible into further subregions), while discrete structures are always *finitely divisible* (and hence always decomposable into atoms). Because of this, the dispute between the continuous theorist and the discrete theorist concerns one of the most basic questions about the structure of consciousness. It's fair—in my view—to compare the question to questions about whether consciousness is atomistic vs. holistic, whether consciousness comes in degrees or is binary, whether consciousness is multidimensional or unidimensional, and whether or not consciousness is compositional.

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'. Many discussions of consciousness and continuity provide 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. In particular, a central contribution of this paper will be to explain the connection between the mathematical definition of 'continuity'—where continuity is a property of *functions*—and the notion of continuity at stake when contrasting smooth and gappy experiences.

Here's the structure of the paper. §1 clarifies the target question; §2 develops a key argument for the continuous theory; §3 discusses the notion of a "continuous space"; §4 explains how to apply the standard mathematical definition of 'continuity' to experiences; §5 explains why the target question of this paper is structurally distinct from the question of whether the stream of consciousness is continuous; §6 defines 'contiguity'; §7 develops an analysis of what makes an experience smooth versus gappy; §8 addresses a verbal objection concerning the meaning of 'continuous'; §9 addresses an objection concerning the meaning of

'feels discrete'; §10 addresses an objection concerning the meaning of 'gappy'; and §11 discusses implications for introspection.

§1 The Target Question

Here's the 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. There's disagreement about the nature of phenomenal character. 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 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 continuous² / 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?³

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.⁴

² One could argue that an experience could represent something *as continuous* even without representing a continuum of points. In other words, one might think it's possible for an experience to represent the high-level property *being continuous* without representing the low-level properties that constitute a continuous structure. My focus will be on only the latter question.

³ If smooth experiences genuinely represent continuous structures, then the veridicality of such experiences would be beholden to fundamental physics (since it may turn out that fundamental physics is discrete). But it's plausible your experience of, say, 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 hold that the contents of smooth experiences are silent on whether their objects are continuous or discrete.

⁴ I myself favor sense-datum theory. See Lee, A. [forthcoming] for a recent defense.

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

The target question is about the metaphysical structure of smooth experiences. This should be distinguished from representational questions about how to best model conscious experiences. Sometimes we use continuous models to represent discrete phenomena (such as when we use differential equations to model predator/prey relationships), and sometimes we use discrete models to represent continuous phenomena (such as when we use discrete scales to model temperatures). The main dialectical aim of this paper is to resist an introspective argument for the continuous theory. But this leaves open which models of conscious experiences are most useful or most explanatory.

For many people, the expression 'consciousness and continuity' will bring to mind the question of whether the stream of consciousness is continuous. But that question is structurally distinct from the target question of this paper. I'll defer the explanation for this to §5, after I've set up the basic framework for investigating the target question. For now, I'll merely note that (a) many examples of smooth experiences involve spatial experience (rather than temporal experience), and (b) even those that involve temporal experience still essentially involve another kind of phenomenology (such as auditory experience).⁷

⁵ For discussions on whether the physical correlates of experience are continuous, see Maxwell [1978], Dennett [1993], Lockwood [1993], Blackmore [2002], VanRullen & Koch [2003], Sergent & Dehaene [2004], White [2018], and Builes [2020].

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

⁷ See Rashbrook [2013] for additional conceptual clarifications about consciousness and continuity. It's worth briefly noting how some of Rashbrook's distinctions relate to the subject-matter of this paper. First, Rashbrook distinguishes (p.612–615) between ascriptions of continuity to (1) the *stream* of consciousness ("the experience had by a subject"), (2) the *state* of consciousness ("the state the subject...is in when awake"), and (3) the *objects* of experience ("what is represented by experience"). While the first notion—the stream of

§2 The Argument for Continuity

I started with a simple 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. This basic line of reasoning captures a compelling motivation for the continuous theory.⁸

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 this move implicitly assumes a dubious analogy between introspecting an experience vs. 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.⁹ 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.

consciousness—is relevant to this paper, I explain later (in §5) why the target question of this paper is structurally distinct from all three of Rashbrook's questions. Second, Rashbrook distinguishes (*p.*617) between "strict continuity" (consciousness has no gaps) and "extreme continuity" (the temporal boundaries of experience fail to be manifest in consciousness). The sense of 'continuity' I'm concerned with corresponds to strict continuity. But I'll explain more precisely what it means for consciousness to lack gaps (§7), and I'll identify an important distinction in the meaning of 'gappy' (in §9).

⁸ This sentiment has been expressed by many philosophers. For example, Dainton [2014] says that "[c]ontinuous-seeming experiences...say, an experience of an extended violin tone, or smooth movement" seem "truly atomless, and hence infinitely divisible," and Prentner [2019: 29] says that the "phenomenology of consciousness is such that it seems composed of an indefinite number of (phenomenal) parts," where each "part of a continuum is always experience itself as being structured." In fact, even those who endorse opposing views tend to concede this point about introspection. For example, Clark [1989: 277]—a discrete theorist—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." And Lockwood [1993: 277]—who is ultimately agnostic on the issue—expresses puzzlement about how "[p]articulate and discontinuous physico-chemical activity [could] yield perceived continuity" or could "just *be* a phenomenally continuous circle."

⁹ Lockwood [1993: 277] entertains the idea that "[j]ust as, with the photograph, the limited resolving power of the eyes ensures that, if we stand back sufficiently, we shall have the illusion of continuity, so we could envisage the mind, in introspection, as standing back from the underlying brain-processes-again, with consequent loss of resolution." But he ultimately says that this "picture [of introspection] is simply incoherent."

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 sophisticated argument for the continuous theory. To my knowledge, this argument below has never been formulated explicitly in the philosophical literature. But I think it's a compelling argument, 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.

§3 **Continuous 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, Q₂ is more directly relevant to the core aims of this paper. But to evaluate Q₂, we'll need to first understand Q₁. 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:




FIGURE 1: The state-space for color experience.¹⁰

The elements of the state-space are individual color qualities (red₁, green₃₄, etc.). The space is structured by a distance metric, which locates more similar color experiences closer to each other within the space, and which generates its three-dimensional structure (with dimensions corresponding to hue, saturation, and brightness). There are also state-spaces for auditory experience, olfactory experience, spatial experience, and any other feature of experience.¹¹

Our current question is what it means to say that a state-space is continuous vs. discrete. This is trickier than it might initially seem. The expression 'continuous space' doesn't have a standard mathematical definition. Instead, in mathematics, continuity is typically understood as a property of functions (a point I'll return to in §4). And while the expression 'discrete space' *does* have a standard mathematical definition, we'll need to be careful about which kind of discrete space we attribute to the discrete theorist.¹²

¹⁰ Note that this is a simplified representation: the model captures only hue and saturation, leaving out brightness. The image is taken from Wikipedia Commons: https://commons.wikimedia.org/wiki/File:Color_circle_%28hue-sat%29.png

¹¹ See Clark [2000] for a more general discussion of state-spaces for features of experiences.

¹² See Franklin [2017] for a more general discussion of continuity versus discreteness.

The intended distinction between continuous and discrete spaces is often illustrated by contrasting \mathbb{R} —the real numbers—and \mathbb{Z} —the 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 sometimes called a ‘continuous space’; \mathbb{Z} is sometimes called a ‘discrete space’.¹³

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. Technically, infinite divisibility isn’t sufficient for continuity. For example, the set \mathbb{Q} of rational numbers is infinitely divisible (any interval of \mathbb{Q} can be divided into subintervals) 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. 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.¹⁴

A *discrete space* is usually defined as a space where each element is *isolated*, meaning that for any element x , there’s some distance δ such that 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

¹³ Note, however, that it’s also possible to make sense of continuous spaces that aren’t composed from point-sized elements (as with \mathbb{R}). For a general philosophical discussion of these kinds of structures, see Arntzenius [2008]. For a discussion in the context of consciousness, see Dainton [2014: 112].

¹⁴ You might still wonder how to formally define ‘continuous space’. It’s tempting to appeal to the notion of a *continuum*, defined as a non-empty compact connected topological space where distinct points have distinct neighborhoods. But this definition precludes \mathbb{R} (which isn’t compact because it’s unbounded) and $(0, 1)$ (which isn’t compact because it doesn’t include its boundary points), both of which are clearly amongst the target structures that we want to capture with the term. A better option might be to appeal to the notion of a *connected manifold* (with dimension ≥ 1), meaning a space that’s locally homeomorphic to Euclidean space at each point, that isn’t the union of two disjoint non-empty subsets, and—in this context—that possibly includes its boundary points.

elements and edges (where each edge connects two elements). Under this formalization, the distance between two elements is naturally measured by 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 (rather than, say, containing some continuous regions and some non-continuous regions). For those interested in the more complex cases, it will be straightforward to generalize my arguments and analyses to state-spaces with other structures. For similar reasons, I'll assume that the state-spaces under discussion consist of all metaphysically possible experiences of the relevant experiential domain, rather than merely the subset of experiences that are possible for a particular creature or a particular species.

§4 Continuous 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 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 $\text{red}_1\text{--red}_{100}$ 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:

experience α satisfies the **continuum condition** with respect to F -experience $\text{=}_{\text{def}}^{15}$

- α instantiates all and only the values within a continuous region of the state-space for F -experiences.

¹⁵ 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 $\text{=}_{\text{def}} F[\alpha]$ is continuous under d (see fn. 9).

The following two figures illustrate the idea. In FIGURE 2, a continuous region of the state-space is enclosed within the dotted line. Hence, an experience that instantiates exactly the color qualities within that region satisfies the continuum condition (with respect to color experience). In FIGURE 3, a *discontinuous* region of the state-space is enclosed within the dotted line. Hence, an experience that instantiates exactly the color qualities within that region does *not* satisfy the continuum condition (with respect to color experience).



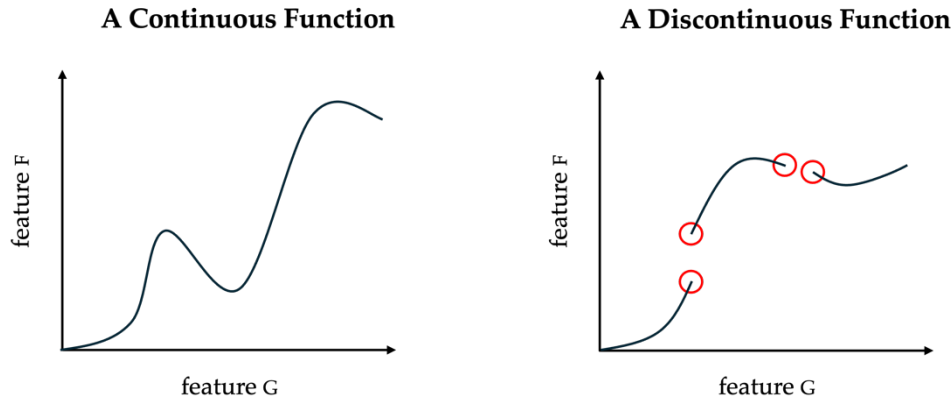
According to the current hypothesis, what it is for an experience to be continuous (with respect to some feature) is for it to satisfy the continuum condition (with respect to that feature). This hypothesis is initially attractive. But it's false.

Here's 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's no systematic correspondence in α between color experience and spatial experience. You might imagine α as similar, in the relevant respects, 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 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."¹⁶ The following figures illustrate

¹⁶ Spivak [2008: 115]. FORMAL DEFINITION: a function f with domain X is *continuous* =_{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

(the red circles mark the discontinuities, and the labeling of the axes will be explained in a moment):



This disparity between the use of ‘continuity’ across disciplines is initially puzzling. 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 to the hypothesis concerning the continuum condition 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, for example, 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.

Now we’re in position to see why the function-theoretic definition of continuity used by mathematicians is relevant to the question of whether individual experiences are continuous. To say that α is continuous in F with respect to G is to

that the most 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 instead on the “epsilon-delta” definition of ‘continuity’ (which assumes metric, rather than merely topological, structure), both because this definition will be accessible to more readers and because it’s plausible that the relevant state-spaces have metric structure.

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).¹⁷

There remains one further complication. Any function with a discrete domain trivially satisfies the mathematical definition of continuity. But this means that some experiences that intuitively ought to count as gappy will nevertheless have a continuous mapping from G-values to F-values. Suppose, for example, that the state-space for spatial experience is discrete and that α is an experience that 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 simply need the additional requirement 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 a continuous space. 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.¹⁸ And with that condition, we arrive at the following definition ('wrt' means 'with respect to'):

¹⁷ You might think that the functions for smooth experiences must not only be continuous, but also *differentiable* (meaning that the function has a derivative at every point). In fact, functions that are differentiable are called "smooth functions" in mathematics. For simplicity, I focus merely on continuity in this paper. But I think that appealing to differentiability is a promising way for the continuous theorist to model smooth experiences.

¹⁸ 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 red₁ 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.

experience α is **continuous** in feature F wrt feature G =_{def}¹⁹

- 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 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's no mapping from α 's emotional experience values to α 's gustatory experience values. Hence, we can define discreteness of individual experiences as follows:²⁰

experience α is **discrete** in feature F wrt feature G =_{def}²¹

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

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

¹⁹ FORMAL DEFINITION: Let $G[\alpha]$ be the set of G-values instantiated by α that map onto F-values. Let f_i be the F-value mapped by any $g_i \in G[\alpha]$. Then α is **continuous** in F with respect to $G =_{\text{def}}$ **(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 $d(f_a, f_n) < \epsilon$.

²⁰ A technical clarification: 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 red₁ at l_1 , red₂ at l_2 , red₃ 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 (\mathbb{Z}) is analogous to α 's G-values, f 's codomain (\mathbb{R}) is analogous to the state-space for F-experiences, and f 's image (the elements of \mathbb{R} that are outputs of f) is analogous to α 's F-values.

²¹ 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 $G =_{\text{def}}$ **(1)** $\forall x \in G[\alpha]$, $\exists \epsilon > 0 : \forall y \in G[\alpha] \setminus \{x\}$, $d(x, y) > \epsilon$, and **(2)** $G[\alpha] \neq \emptyset$.

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.²² Given this, I'll occasionally talk of *F-values* being instantiated at *G-locations*.

§5 The Stream of Consciousness

You might wonder whether we always need to invoke the function-theoretic definition of continuity when ascribing continuity to conscious experiences. Can't we sometimes simply ask whether an experience is continuous in feature *F*, without relativizing the question to a feature *G*?

The best candidates for such cases involve "locative features," such as spatial and temporal phenomenology. It seems sensible to simply ask whether an experience is continuous with respect to spatial or temporal phenomenology, even if we don't specify another feature. In fact, a central question in the temporal experience literature—namely, whether the stream of consciousness is continuous—is arguably a question of this form. In the remainder of this section, I'll focus on this question.

What exactly does it mean for the stream of consciousness to be continuous? This question has received good amount of attention in the philosophical literature.²³ Now, there are a number of different ways in which that question is interpreted. But one interpretation concerns whether the stream generates a "continuous space," an issue discussed in §2. In other words, if we are mapping the structure of temporal phenomenology, then we might wonder whether it's more apt to appeal to model it using an interval of \mathbb{R} or an interval of \mathbb{Z} .²⁴ I won't speculate on

²² See Clark [2000: Ch.3] on locative vs. qualitative features.

²³ For example, William James [1890: 239] who coined the term 'stream of consciousness', wrote that "[c]onsciousness...does not appear to itself be chopped up in bits...it flows." In contemporary philosophy, proponents of this view include Foster [1979] and Dainton [2000], while proponents of discrete theories include Chuard [2017] and Prosser [2017]. Other relevant discussions that stay agnostic include Lockwood [1993] and Dainton [2014].

²⁴ A natural thought is that so long as experience α satisfies the continuum condition for spatial or temporal experience, α is continuous with respect to that kind of experience. Perhaps that's right—though I think the answer isn't entirely obvious. Suppose experience α instantiates color values at spatial positions l_1 – l_{10} and l_{20} – l_{30} (with no visual experience at l_{10} – l_{20} or l_{30} – l_{40}) and sound values 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

the answer to this. Instead, I wish to close a point that was alluded to earlier: namely, that the target question of this paper is structurally distinct from the question of whether the stream of consciousness is continuous.

The main goal of this paper is to understand the structure of smooth experiences, defined as experiences that involve perfectly gradual variation in one feature (say, color or sound experience) *with respect to* another feature (say, spatial or temporal experience). By contrast, the question of whether the stream of consciousness is continuous concerns merely a single feature (temporal experience). Because of this, the function-theoretic definition of ‘continuity’ isn’t straightforwardly applicable to discussions of temporal experience.²⁵ This means the sense of ‘continuity’ principally relevant to smooth experiences is different from the sense of ‘continuity’ principally relevant to the stream of consciousness. The former concerns continuous functions; the latter concerns continuous spaces. Both questions are important, but they’re structurally distinct.

Of course, there are still connections between the questions. Whether certain kinds of smooth experiences are continuous will depend on whether the stream of consciousness is continuous. Consider your experience of a rising pitch over an interval of time. If this experience is continuous in auditory experience (feature F) with respect to temporal experience (feature G), then the experience must instantiate a continuous set of values with respect to temporal experience. That’s equivalent, in this context, to the claim that the stream of consciousness is continuous. But not all smooth experiences involve temporal experience. The opening example of this paper—your experience of the blue sky on a cloudless day—concerned spatial experience (and color experience).²⁶

Since I’ve left open whether the state-spaces for experiences are in fact continuous or discrete, my arguments won’t directly adjudicate debates about the stream of consciousness. Nevertheless, my defense of the discrete theory may still

α is the fusion of a spatially disconnected visual experience and a spatially disconnected auditory experience, is it correct to say that α is continuous in spatial experience?

²⁵ One way to apply the function-theoretic definition of ‘continuity’ is to ask whether temporal experience is continuous with respect to time. But this is still different from the target question of this paper, since it involves a mapping from a non-phenomenal feature to a phenomenal feature.

²⁶ There may still be complications, though. Suppose experience α instantiates color values at spatial positions l_1 – l_{10} and l_{20} – l_{30} (with no visual experience at l_{10} – l_{20} or l_{30} – l_{40}) and sound values 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 fusion of a spatially disconnected visual experience and a spatially disconnected auditory experience, is it correct to say that α is continuous in spatial experience?

be relevant to those debates. Just as with smooth experiences, philosophers have appealed to introspection to motivate ascriptions of continuity to the stream of consciousness. While the discrete analysis that I'll develop later is designed for smooth experiences (involving two features), it's plausible that a similar analysis could be developed for the structure of the stream of consciousness. In fact, I'll briefly note what this might look like in §7, when discussing the structural difference between smoothness and gappiness.

These points are worth emphasizing because questions connecting consciousness and continuity are sometimes automatically assumed to be questions about temporal experience. But while questions about the stream of consciousness are important, interesting, and challenging, I think that some of the most interesting questions about consciousness and continuity arise only once we start thinking in a more general way.

§6 Contiguity

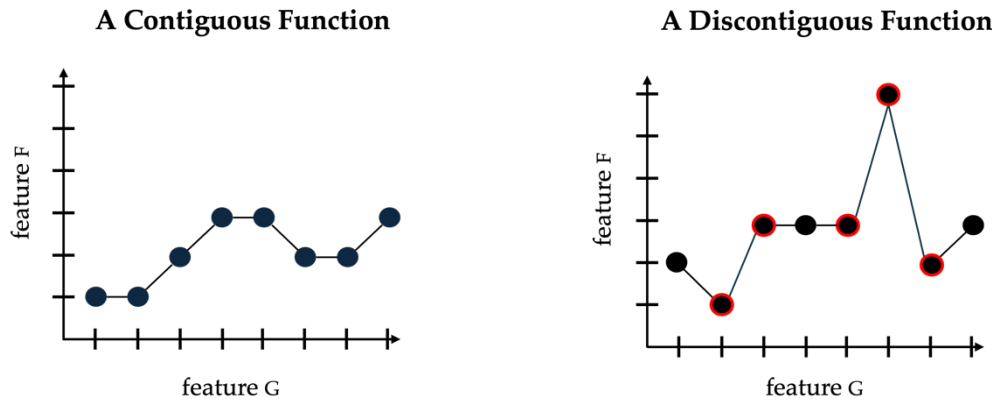
When I presented the Argument for Continuity (§2), I identified the premise that I want to reject. It was the *discrete-implies-gappy* premise, which says that if an experience is discrete, then it's gappy.

To develop my argument against the discrete-implies-gappy premise, I'll develop a discrete analysis of the structure of smooth experiences. In order to do that, I need to first define a new structural property that 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, ...)	B ₁ :	(1, 3, 5, 7, ...)
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 discontinuous because some intermediate integers are missing: for example, sequence B₁ jumps from 3 to 5. The notion of a sequence may initially seem distinct from the formal tools we have invoked so far, but sequences are really just 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). This is illustrated in the following figures:



The idea of contiguity can be generalized beyond numerical spaces. For any discrete state-space, let's say that two elements x and y are *adjacent* just in case they're connected by a single edge (assuming that the relevant spaces can be represented using collections of nodes and edges). A *path* from x to y is a sequence of edges that connects x and 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's 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:²⁷

²⁷ Lee, A. [2021] argues that the qualities of conscious experiences are best modeled by regions (rather than individual points) of metric spaces. This raises interesting questions about how to generalize the notions of continuity and contiguity to such a framework. I won't pursue this question here, but I suspect that it's possible to construct degree-theoretic analogues of continuity and contiguity that apply to regions.

experience α is **contiguous** in feature F wrt feature G =_{def}²⁸

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

Contiguity and continuity are mutually exclusive.²⁹ 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 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 *vs.* 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 A_b , 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 merely 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

²⁸ FORMAL DEFINITION: Let $G[\alpha]$ be α 's set of G -values that map onto F -values, and let f_i be the F -value mapped by any $g_i \in G[\alpha]$. Then α is **contiguous** in F with respect to G =_{def} **(1)** $G[\alpha]$ is contiguous (fn. 20), and **(2)** $\forall g_1, g_2 \in G[\alpha]$, if g_1 is adjacent to g_2 , then f_1 is adjacent to f_2 .

²⁹ 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 .

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 again 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 §1. 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.³⁰

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

³⁰ Note that these analyses are stated 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.

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.

In §5, I mentioned the idea that we can sometimes ask whether an experience is continuous with respect to just a single feature G, when G is a kind of locative experience. There are natural ways to generalize my analyses of smoothness and gappiness to such cases. In particular, we might take experience α to be **gappy** wrt feature G just in case (1) α contains non-adjacent locations g_1 and g_n , and (2) α is missing some intermediate G-locations. Similarly, we might take experience α to be **smooth** wrt feature G just in case (1) α instantiates non-adjacent locations g_1 and g_n , and (2) α contains all intermediate G-locations.

A point worth highlighting is that my analysis of smoothness is *structural*, rather than epistemic. By contrast, those who have expressed sympathy towards the discrete theory usually account for smooth experiences by appealing to limits in our introspective capacities. For example, Clark [1989] contends that smooth experiences might be composed of many tiny discrete units (like pixels) that are too granular to be individually noticeable. Similarly, Fara [2001: 927] says that we cannot be sure that a moving object “does not really look as if it is moving *discontinuously*, looking to take very tiny discrete jumps” because those steps might be “too slight...for us to notice.” The idea is that smooth experiences might merely be 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.

The remainder of the paper discusses some objections.

§8 The Verbal Objection

Here's the first objection:

The Verbal Objection: The way this paper uses the term ‘continuous’ differs from the way other philosophers have used the term ‘continuous’. Other philosophers meant something looser, such as ‘either continuous or contiguous’ or ‘strikes one as continuous’.

I’ll make a few points in response.

First—as noted earlier in the paper—there’s textual evidence of other philosophers using the term ‘continuous’ in a sense that entails infinite divisibility. Furthermore, the “apparent continuity” of consciousness has often been thought to be in tension with the apparent discreteness of its physical correlates. But if ‘continuous’ in the looser ways described above, then continuity and discreteness would actually be compatible with each other.

Second, we should distinguish speaker meaning from semantic meaning. Casual ascriptions of continuity often aren’t sensitive to the sorts of distinctions that have been drawn in this paper. But that doesn’t mean that we should default to a non-standard interpretation of ‘continuous’ when interpreting those claims. As an analogy, if I were to say ‘Tomatoes are vegetables’, then my claim is false, even if I know next to nothing about the standard definition of ‘vegetable’. Similarly, if a philosopher says ‘Conscious experiences are continuous’ (and provides no explicit specification of what they mean by ‘continuous’), then we should interpret ‘continuous’ in the standard sense when evaluating their claim. There are, of course, authors who use ‘continuous’ in a looser way. But as Fara [2001: 923] notes, “just because it may be convenient to describe a change as apparently continuous” does not mean “that it really is that way.”

Third, the sense of ‘continuity’ defined in this paper respects established terminological standards in science, mathematics, and other areas of philosophy. If we were to use the term ‘continuous’ in a non-standard way, then we would multiply senses without much gain. Suppose, for example, that we were to use ‘continuous’ to mean ‘continuous or contiguous’. Then even many paradigmatically discrete structures (such as \mathbb{Z}) would count as continuous, questions about continuity would no longer have direct implications for questions about infinite divisibility, and there would be a confusing disparity between the sense of ‘continuity’ invoked in discussions of consciousness *vs.* the sense of ‘continuity’ invoked by mathematicians, scientists, metaphysicians, and philosophers of science.

§9 The Feels-Discrete Objection

Here’s the second 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's a reading of P1 that's true and a reading of P2 that's 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 .³¹

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. Now, perhaps you can come to know that your visual experience is prime-colored if you introspect carefully enough. But while your visual experience strikes you as colored, it doesn't strike you as prime-colored (nor as not 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

³¹ 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 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.

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 may itself 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 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 red₁₇ and the adjacent pixel red₅₄. It may be tempting to think of contiguous values in a state-space as structurally analogous. But consider the fact that you're able to notice the discontinuities across adjacent pixels only because you're able to perceptually represent the color values that are missing between pixels—that is, some of the values between red₁₇ and red₅₄. 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.

§10 The Discontinuity Objection

Here's the third objection:

⊥ The Discontinuity Objection

- P1:** If an experience is discrete, then it's discontinuous.
P2: If an experience is discontinuous, then it contains gaps.
P3: 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 sense of 'gappy' at play in this paper.

Here's my analysis from earlier of what it is for an 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 start of the paper? Every example that I've listed of a gappy experience involves an experience that's missing intermediate values in 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's clearly still a sense in which we can talk about some physical things containing gaps and other physical things lacking gaps. A wall that's half black and half white is gappy (in color with respect to space), while a wall that's 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.

§11 Introspection

At the start of the paper, I presented a simple argument for thinking that introspection favors the continuous theory. My main goal, over the course of the paper, has been to argue that the discrete theory is compatible with our introspective evidence. 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 what introspection reveals.

In fact, if the continuous theory is correct, then we can introspect exactly 0% of the qualities of our experiences! This is because our introspective capacities are finite: within any experience, one can introspectively discriminate only a finite number of distinct qualities. But the continuous theory entails that any non-homogenous smooth experiences (meaning smooth experiences that instantiate

multiple F-values) 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, the discrete theory entails that non-homogenous smooth experiences 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.

This point 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.³² 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.

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

If it's true that introspection leaves open whether smooth experiences are continuous or discrete, then we can also infer the falsity of the principle often known as LUMINOSITY: if one's experience is F (where F is a phenomenal property), then one can know via introspection that one's experience is F. In fact, the question of whether our experiences are continuous vs. discrete strikes me as one of the most compelling cases against LUMINOSITY. This is not only because of the points I've made concerning contiguity, but also because of general considerations about the finitude of introspection. Suppose—as seems plausible—that necessarily, any

³² See Lee, A. [2019] for a defense of non-introspectable microphenomenal properties. See Schwitzgebel [2008] for more general arguments about the limits of introspection.

subject can have only finite introspective capacities. Then it's likewise plausible that necessarily, any subject's introspective evidence leaves open whether their experiences are continuous or discrete. Hence, LUMINOSITY may fail to hold not only for humans, but for any conscious subjects whatsoever.

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.

A goal of this paper has been to demonstrate the utility of *mathematical phenomenology*, or the application of mathematical tools to the study of conscious

experiences. To pursue this fruitfully, we need to think carefully about how the mathematical concepts connect to the phenomenology. But—in my opinion—the identification of those connections is one of the most promising and exciting lines of inquiry in consciousness research.

References

- Arntzenius, Frank (2008). Gunk, Topology and Measure. In Dean Zimmerman, *Oxford Studies in Metaphysics: Volume 4*. Oxford, GB: Oxford University Press UK.
- 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.
- Chuard, Philippe, 2017, "The Snapshot Conception of Temporal Experience", in Ian Phillips (ed.), *The Routledge Handbook of Philosophy of Temporal Experience*.
- 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.
- 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.
- Hurevicz, W. & Wallman, G. (1948) *Dimension theory*, Princeton University Press.

-
- James, William, 1890, *The Principles of Psychology*, 2 vols., (American Science—Advanced Course), New York: H. Holt and company.
- Lee, Andrew Y. (2019). The Microstructure of Experience. *Journal of the American Philosophical Association* 5 (3):286-305.
- Lee, Andrew Y. (2021). Modeling Mental Qualities. *The Philosophical Review* 130 (2):263-209.
- Lee, Andrew Y. (forthcoming). A Theory of Sense-Data. *Analytic Philosophy*.
- 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.
- Maxwell, Grover (1978). Unity of consciousness and mind-brain identity. In John C. Eccles (ed.), *Mind and Brain*. Paragon House.
- Phillips, Ian (2011). Indiscriminability and experience of change. *Philosophical Quarterly* 61 (245):808 - 827.
- Prentner, Robert (2019). Consciousness and topologically structured phenomenal spaces. *Consciousness and Cognition* 70:25-38.
- Prosser, Simon (2016). *Experiencing Time*, Oxford: Oxford University Press. doi:10.1093/acprof:oso/9780198748946.001.0001
- 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.

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