

# Consciousness and Continuity

ANDREW Y. LEE

University of Oslo, Department of Philosophy

## ABSTRACT:

Consider your visual experience of a blue sky, your auditory experience of a rising pitch, or your temporal experience of an interval of time. Does introspection indicate that the phenomenal characters of such experiences are continuous (as opposed to discrete)? I first clarify what exactly it means to say that conscious experiences have continuous structures. Then I develop a discrete model of such experiences. I argue that the aforementioned experiences might be *contiguous*, where this means roughly that adjacent values of one experiential domain map to adjacent values of another experiential domain. If an experience is contiguous, then it is discrete. This means that introspection leaves open whether conscious experiences are continuous or discrete. At the end, I argue that the discrete theory does better than the continuous theory in minimizing the gap between the structure of experience and the grain of introspection.

---

## §1 Introduction

Consider the phenomenal characters of the following kinds of conscious 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.
- e. Your temporal experience over an interval of time.

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 is 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.
- e. Your visual experience when looking at a pixelated low-resolution image.

What does introspection reveal about the structures of smooth experiences and gappy experiences?<sup>1</sup> More specifically, does introspection indicate that the phenomenal characters of smooth experiences are continuous? Or does it leave open the possibility that they are discrete? 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 a natural line of reasoning: (P1) introspection reveals no discontinuities in smooth experiences, and (P2) to be continuous is to lack discontinuities, so (C) introspection favors the continuous theory. This argument feels compelling. The premises are hardly contestable: P1 is a solid empirical fact, and P2 is an analytic

---

<sup>1</sup> This question concerns the *phenomenal character* of smooth experiences, rather than (say) its physical realizers. In §2, I'll say more about how to interpret this paper's target question.

---

truth. One 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 P2 establishes an equivalence between the relevant negative property (lacking discontinuities) and the relevant positive property (being continuous).

Moreover, the alternative view — that smooth experiences are discrete — may strike some as phenomenologically inadequate. If an experience is discrete, then it should 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. I have some sympathy for this move. But I suspect some will find the move dialectically unsatisfying, at least without some supporting arguments. The continuous theorist makes a positive claim that (at least initially) seems to align with the phenomenology. The discrete theorist, on the other hand, makes a negative claim that (at least initially) seems in tension with the phenomenology.

Here are the goals of this paper. 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.

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. After presenting the basic ideas, I'll argue that a contiguous model of smooth experiences adequately accounts for their phenomenology. I won't take a stance on whether smooth experiences are in fact continuous or discrete. Instead, my goal is merely to argue that introspection itself is neutral.

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 provide only cursory glosses of

---

continuity and discreteness. Once we start to examine the question more carefully, we will encounter a number of complexities. Some of these complexities will require us to think about the technical details, but my discussion of technical points will be in service of clarifying the philosophical ideas.

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 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.<sup>2</sup>

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

---

<sup>2</sup> 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.

---

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? 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.

The point above speaks to a question about the explanatory target of the paper. A referee asks: is the question whether smooth experiences *represent* continuous structures, or whether smooth experiences *themselves* are continuous? That is, one might wonder whether the target question is intended to concern vehicles or contents. However, there's no theory-neutral answer to that question. The answer depends on how one thinks about the relationship between phenomenal character, vehicles, and contents. If one identifies phenomenal character with the contents of experiences, then the target question concerns contents. If one identifies phenomenal character with the vehicles of experiences, then the target question concerns vehicles. If one thinks that the vehicle/content distinction isn't even applicable to experiences—as certain relationalists and certain qualia theorists do—then the target question concerns neither vehicles nor contents. But whichever of these options one favors, one can still wonder whether phenomenal character—whatever one takes it to be—is continuous or discrete in structure.<sup>3</sup>

---

<sup>3</sup> A sidenote: 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.

---

The target question should be distinguished from the question of whether the neural correlates of conscious experiences have continuous structures.<sup>4</sup> 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.<sup>5</sup> 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.<sup>6</sup>

### §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 “phenomenal continua” that are infinitely divisible. 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

---

<sup>4</sup> 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 Builes [2020].

<sup>5</sup> For examples of views that dissociate these factors, see Lee, G [2014] and Phillips [2011].

<sup>6</sup> 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.

---

seems composed of an indefinite number of (phenomenal) parts.”<sup>7</sup> And even some philosophers who reject the continuous theory sometimes say 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.”

The introspective argument that conscious experiences are continuous will strike many as more compelling than a structurally analogous argument appealing to perception of external objects. 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 such distinction is applicable.

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.

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

---

<sup>7</sup> Note that for the purposes of this paper, it doesn’t matter much who intended to ascribe continuity to conscious experiences and who intended to ascribe some more imprecise property. The goals of this paper are to clarify what it means to say that conscious experiences have continuous structures, to develop a discrete model of smooth experiences, and to argue that introspection is neutral on whether smooth experiences are continuous or discrete, rather than to argue against the specific claims that have been quoted.

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 uncontested, 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<sub>2</sub> 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.<sup>8</sup> At an abstract level, a state-space is a set of elements and a distance metric that captures the similarity relations between those elements.<sup>9</sup> The state-space for color experiences has  $red_1$ ,  $red_2$ , and so forth as its elements, and its metric yields the result that  $red_1$  is closer to  $red_2$  than to  $red_{10}$ .

The distinction between continuous and discrete spaces is often illustrated by appeal to a contrast between the set  $\mathbb{R}$  of real numbers and the set  $\mathbb{Z}$  of integers.<sup>10</sup> In  $\mathbb{R}$ , every element is perfectly 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 a paradigm of a continuous space;  $\mathbb{Z}$  is a paradigm of a discrete space.<sup>11</sup>

In philosophical contexts, questions about continuity are often treated as equivalent to questions about infinite divisibility. Consider how debates about whether space, time, and matter are continuous are taken to turn on whether space, time, and matter are infinitely divisible into arbitrarily small spatial regions, temporal intervals, and material parts. Strictly speaking, infinite divisibility is weaker than 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 irrationals). But for our purposes, the distinction between continuity and infinite

---

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

<sup>9</sup> I'll assume that all the state-spaces under consideration are metric spaces, since a metric is arguably needed to capture the similarity relations within any domain of experiences.

<sup>10</sup> I'll assume the standard metrics for  $\mathbb{R}$ ,  $\mathbb{Z}$ , and every other set I discuss.

<sup>11</sup> See Franklin [2017] for more general discussion of continuity versus discreteness.

divisibility doesn't really matter. To my knowledge, no contemporary philosopher has seriously argued that space or time or matter or consciousness is infinitely divisible yet not continuous. Given this, I'll follow orthodoxy and assume that the state-space for a domain of experiences is infinitely divisible just in case it's continuous.

The distinction between continuous and discrete state-spaces can be formally characterized. Let  $d$  be the metric telling us the distance between any two elements of the state-space under consideration. A state-space is *discrete* just in case each of its elements is isolated, meaning that for any element  $x$  there is some distance  $\delta$  where no distinct element  $y$  is such that  $d(x, y) < \delta$ . A state-space is *continuous* just in case it's connected, where this means roughly that the space is infinitely divisible with no discontinuities between any pair of its elements.<sup>12</sup> To simplify 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 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  $\alpha$  and any experiential feature  $F$ , there is a (possibly empty) set of values from the state-space for  $F$ -experiences that are instantiated by  $\alpha$ . Suppose, for example, that

---

<sup>12</sup> More precisely, a state-space  $X$  is *continuous* =<sub>def</sub>  $X$  isn't the union of two open subsets. The terminology is a bit awkward here. In mathematics, continuity is a property of functions, rather than spaces (see §5); the above definition is the definition for a connected space. However, philosophers (and other theorists) oftentimes talk about continuous spaces, and the natural interpretation is that a continuous space is a connected space.

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  $\alpha$  the *F-values of  $\alpha$* .

Here's a hypothesis that's initially attractive but that turns out to be false: an experience  $\alpha$  is continuous in feature F just in case  $\alpha$  instantiates a continuous set of F-values, meaning  $\alpha$  instantiates exactly the values within a continuous region of the state-space for F-experiences. In this circumstance, let's say that  $\alpha$  satisfies the *continuum condition* with respect to F-experience.<sup>13</sup> In the example from above, if the state-space for color experience is continuous and if  $\text{red}_1$ – $\text{red}_{100}$  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  $\alpha$  instantiates every color quality within some continuous region of that state-space. However, suppose there is no systematic correspondence in  $\alpha$  between color experience and spatial experience. You might imagine  $\alpha$  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  $\alpha$  satisfies the continuum condition for color experience,  $\alpha$  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."<sup>14</sup> At first, it may seem as though

---

<sup>13</sup> FORMAL DEFINITION: Let  $F[\alpha]$  be the set of F-values instantiated by  $\alpha$  and  $d$  be the metric for the state-space for F-experiences. Then  $\alpha$  satisfies the **continuum condition** with respect to F-experience =<sub>def</sub>  $F[\alpha]$  is continuous under  $d$  (see fn. 9).

<sup>14</sup> Spivak [2008: 115]. FORMAL DEFINITION: a function  $f$  with domain  $X$  is *continuous* =<sub>def</sub> 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 this definition appeals to metric notions. A more general definition of continuity is topological, where continuous functions are those where the pre-images of open sets are open. I've

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  $\alpha$  is continuous, the question must be precisified as whether  $\alpha$  is continuous in some feature  $F$  *with respect to some other feature*  $G$ . Even though  $\alpha$  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  $\alpha$  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.<sup>15</sup> A natural thought is that so long as  $\alpha$  satisfies the continuum condition for spatial experience,  $\alpha$  is continuous in spatial experience. Perhaps that is right (though I think the answer isn't entirely unobvious).<sup>16</sup> 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

---

chosen to focus on metric definitions because they are more accessible and because the choice of metrical vs. topological notions doesn't matter much for the issues in this paper.

<sup>15</sup> See Clark [2000: Ch. 3] on locative experiences.

<sup>16</sup> Suppose  $\alpha$  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  $\alpha$  satisfies the continuum condition for spatial experience. But given that  $\alpha$  is the conjunction of a spatially disconnected visual experience and a spatially disconnected auditory experience, is it correct to say that  $\alpha$  is continuous in spatial experience?

---

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.<sup>17</sup> 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  $\alpha$  is continuous in F with respect to G is to say that the mapping from  $\alpha$ 's G-values to  $\alpha$ 's F-values is a continuous function, where this means that sufficiently small changes in  $\alpha$ 's G-values map to arbitrarily small changes in  $\alpha$ '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 count as continuous given our present analysis. Suppose that the state-space for spatial experience is discrete and that  $\alpha$  represents red at spatial location  $l_1$ , green at  $l_2$ , and blue at  $l_3$ . It would be bizarre to say that  $\alpha$  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  $\alpha$ 's G-values to  $\alpha$ 's F-values is continuous.

Fortunately, we've already encountered the tool that's needed to solve this problem. We need simply require that  $\alpha$  satisfies the continuum condition for *feature G*, meaning that  $\alpha$ '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  $\alpha$  to be continuous in F with

---

<sup>17</sup> See Clark [2000: Ch.3] for discussion.

respect to G, the state-space for G-experiences must itself be continuous.<sup>18</sup> And with that condition, we arrive at the following definition:

experience  $\alpha$  is **continuous** in feature F with respect to feature G  $\stackrel{\text{def}}{=}$ <sup>19</sup>

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

The remaining task is to define what it is for an individual experience  $\alpha$  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  $\alpha$  to trivially count as discrete in gustatory experience with respect to emotional experience simply because there is no mapping from  $\alpha$ 's emotional experience values to  $\alpha$ 's gustatory experience values. With this condition, we can define discreteness of individual experiences as follows:<sup>20</sup>

---

<sup>18</sup> 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  $\alpha$  is an experience as of looking at a uniformly red wall, where red<sub>i</sub> is represented at spatial locations  $l_1$ – $l_{100}$ . Then  $\alpha$  is continuous in color experience with respect to spatial experience, even though the former state-space is discrete.

<sup>19</sup> FORMAL DEFINITION: Let  $G[\alpha]$  be the set of G-values instantiated by  $\alpha$  that map onto F-values, and let  $f_i$  be the F-value mapped by any  $g_i \in G[\alpha]$ . Then  $\alpha$  is **continuous** in F with respect to G  $\stackrel{\text{def}}{=}$  **(1)**  $\alpha$  satisfies the continuum condition with respect to G-experience (see fn. 11), 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  $d(f_a, f_n) < \epsilon$ .

<sup>20</sup> If the mapping from  $\alpha$ 's G-values to F-values isn't a function, then the analysis entails that  $\alpha$  isn't continuous in F with respect to G.

experience  $\alpha$  is **discrete** in feature F with respect to feature  $G \stackrel{\text{def}}{=} G^{21}$

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

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:

A1:	(1, 2, 3, 4, ...)	B1:	(1, 3, 5, 7, ...)
A2:	(1, 1, 1, 1)	B2:	(1, 1, 1, 3)
A3:	(3, 2, 1, 0, 1, 2, 3)	B3:	(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 B<sub>1</sub> 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. Let's say two elements  $x$  and  $y$  of a state-space are *adjacent* just in case there are no other values between them.<sup>22</sup>

<sup>21</sup> FORMAL DEFINITION: Let  $G[\alpha]$  be the set of  $\alpha$ 's G-values that map to F-values.  $\alpha$  is **discrete** in F with respect to  $G \stackrel{\text{def}}{=} G[\alpha]$ ,  $\exists \epsilon > 0 : \forall y \in G[\alpha] \setminus \{x\}, d(x, y) > \epsilon$ , and **(2)**  $G[\alpha] \neq \emptyset$ .

<sup>22</sup> FORMAL DEFINITION:  $f_1 \in F$  is **adjacent** to  $f_2 \in F \stackrel{\text{def}}{=} \neg \exists f_3 \in F (d(f_1, f_2) > d(f_1, f_3) \ \& \ d(f_1, f_2) > d(f_2, f_3))$ .

Next, let's say a *contiguous region* of a state-space is a region where for any two elements  $x$  and  $y$ , there is a sequence of pairs of adjacent elements starting with  $x$  and ending with  $y$  such that the first element of each subsequent pair is the same as the second element of each preceding pair.<sup>23</sup> In other words, contiguous regions are those where we can move from any value to any other via a chain of adjacency pairs. Finally, let's say experience  $\alpha$  *instantiates a contiguous set* of G-values just in case  $\alpha$  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  $\alpha$  is **contiguous** in feature F with respect to feature G =<sub>def</sub><sup>24</sup>

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

In the examples below,  $\alpha_1$  and  $\alpha_2$  are contiguous and  $\beta_1$  and  $\beta_2$  are discontinuous in color experience with respect to spatial experience (assume the state-spaces for both color experience and spatial experience are discrete):<sup>25</sup>

- $\alpha_1$ : an experience of a spectrum of color, representing red<sub>1</sub>–red<sub>50</sub> at locations  $l_1$ – $l_{50}$  (respectively).
- $\alpha_2$ : an experience of a homogenous field of color, representing red<sub>1</sub> at locations  $l_1$ – $l_{50}$ .

<sup>23</sup> FORMAL DEFINITION: Let F be a state-space and  $F[\alpha] \subseteq F$ . Then region  $F[\alpha]$  is **contiguous** =<sub>def</sub>  $\forall f_1, f_2 \in F[\alpha], \exists f_a \dots f_z \in F[\alpha]$  ( $f_1$  is adjacent to  $f_a$ ,  $f_a$  is adjacent to  $\dots$   $f_z$ ,  $f_z$  is adjacent to  $f_2$ ).

<sup>24</sup> FORMAL DEFINITION: Let  $G[\alpha]$  be  $\alpha$ '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  $\alpha$  is **contiguous** in F with respect to G =<sub>def</sub> **(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  $f_2$ .

<sup>25</sup> Lee, A [2021] argues that experiences correspond to regions, rather than individual elements, in their respective state-spaces. This complicates the analysis of contiguity, since it's not obvious how to generalize contiguity to such a framework. But I suspect it's possible to construct a degree-theoretic analogue of contiguity for regions, where "degree of contiguity" between two regions is the proportion of elements in each region that are contiguous to some element within the other region.



- $\beta_1$ : an experience of a series of colors with every other color skipped over, representing red<sub>1</sub> at location  $l_1$ , red<sub>3</sub> at  $l_2$ , red<sub>5</sub> at  $l_3$ , and so forth.
- $\beta_2$ : an experience of a spectrum of colors with a gap, representing red<sub>1</sub>–red<sub>50</sub> (except for red<sub>37</sub>) at locations  $l_1$ – $l_{50}$  (except for  $l_{37}$ ).

Contiguity and continuity are mutually exclusive.<sup>26</sup> 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 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

---

<sup>26</sup> The sole exception is the degenerate case where  $\alpha$  instantiates only a single G-value and a single F-value: such an experience is both continuous and contiguous in F with respect to G.

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  $\alpha$  is **gappy** in feature F with respect to feature G  $\leftrightarrow$

- $\alpha$  instantiates non-adjacent values  $g_1$  and  $g_2$  mapping to  $f_1$  and  $f_2$ .
- $\alpha$  is missing intermediate F-values at intermediate G-locations.

I think this analysis of gappiness is fundamentally correct. But we will need to be careful in how we interpret the analysis. Let's say that  $z$  is *between*  $x$  and  $y$  just in case  $d(x, z) < d(x, y)$  and  $d(y, z) < d(x, y)$ . At first, it's tempting to interpret the analysis in the following way:  $\alpha$  is gappy in feature F with respect to feature G  $\leftrightarrow$  for any F-values  $f_1$  and  $f_2$  instantiated at G-locations  $g_1$  and  $g_2$  in  $\alpha$ , there is some value  $f_n$  between  $f_1$  and  $f_2$  in the state-space for F-experiences such that  $f_n$  is not instantiated in  $\alpha$  at any location between  $g_1$  and  $g_2$ . However, this interpretation only if we restrict our focus to one-dimensional linear state-spaces: it breaks down if we begin thinking about multidimensional or cyclical state-spaces.

Here are two cases that exhibit the limits of that interpretation. **Case 1:** Suppose feature F is a two-dimensional color space, with the dimensions corresponding to redness and blueness. Suppose  $\alpha$  instantiates  $\langle \text{red}_1\text{--red}_{10}, \text{blue}_1 \rangle$  at locations  $l_1\text{--}l_{10}$  and  $\langle \text{red}_{10}, \text{blue}_1\text{--blue}_{10} \rangle$  at locations  $l_{11}\text{--}l_{20}$ . In other words, as we move from  $l_1$  to  $l_{20}$ ,  $\alpha$ 's color values first shift in redness (while blueness remains fixed) and then shift in blueness (while redness remains fixed). Intuitively,  $\alpha$  ought to count as smooth in color with respect to space. Yet  $\alpha$  is missing some intermediate color values (such as  $\langle \text{red}_5, \text{blue}_5 \rangle$ ) at intermediate locations. **Case 2:** Suppose feature F is hue, which has a cyclical (rather than linear) structure. Let all hue values be mapped to points on the circumference of a circle, where each point corresponds to a real number

between 0 and 1, and where  $\text{hue}_0 = \text{hue}_1$ . Then suppose  $\alpha$  instantiates  $\text{hue}_{0.01}$ – $\text{hue}_{0.90}$  at locations  $l_1$ – $l_{90}$  (and doesn't instantiate any values between  $\text{hue}_{0.90}$  and  $\text{hue}_{0.01}$ ). Intuitively,  $\alpha$  ought to count as smooth in hue with respect to space. Yet  $\alpha$  is missing some intermediate hue values (such as  $\text{hue}_0$ ) at intermediate locations.

The cases above show that gappiness requires more than merely that there exists *some* value  $f_n$  between  $f_1$  and  $f_2$  in the state-space for F-experiences that is uninstantiated by any location between  $g_1$  and  $g_2$ . Rather, gappiness requires also that the values that are missing are values that would “fill in” the gap. In other words, gappiness occurs whenever some pair of  $\alpha$ 's F-values are disconnected from each other relative to  $\alpha$ 's G-locations. If we were to track  $\alpha$ 's F-values as we move across  $\alpha$ 's G-locations, we would sometimes skip over intermediate F-values from the state-space for F-experiences.

The natural corollary hypothesis is that smooth experiences are those where every intermediate G-value maps to an intermediate F-value. In one-dimensional linear state-spaces, this is equivalent to requiring that if  $\alpha$  instantiates two F-values  $f_1$  and  $f_2$ , then  $\alpha$  also instantiates all the F-values between  $f_1$  and  $f_2$ . But the challenge cases from above indicate that smoothness is compatible with  $\alpha$  missing some intermediate F-values at intermediate G-locations, so long as we never skip over F-values as we move across G-locations. In other words, smoothness occurs whenever every pair of  $\alpha$ 's F-values is connected relative to  $\alpha$ 's G-values. 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  $\alpha$  is **smooth** in feature F with respect to feature G  $\leftrightarrow$

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

These characterizations of gappiness and smoothness are intuitively plausible, and 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 characterizations of gappiness and smoothness are neutral between the continuous theory and the discrete theory.<sup>27</sup>

Now for an important result: both continuity and contiguity satisfy the analysis of smoothness. If the relevant state-spaces are continuous, then lacking gaps 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 lacking gaps 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.

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.<sup>28</sup> 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:

---

<sup>27</sup> 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 claims about what the experiences we labeled as 'smooth' and 'gappy' have in common.

<sup>28</sup> Clark [1989] is the most explicit defense of the discrete theory I have found. Fara [2001], like me, contends that introspection is neutral as to whether smooth experiences are continuous or discrete, but offers an epistemic account of smoothness.

---

⊥     **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  $\varphi$ ': (1) an experience *phenomenally* feels  $\varphi$  just in case it has phenomenal character  $\varphi$ , and (2) an experience *epistemically* feels  $\varphi$  just in case it strikes its subject as  $\varphi$ .<sup>29</sup>

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  $\pi$  is a transcendental number if you think carefully enough, even though  $\pi$  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.

---

<sup>29</sup> The phrase 'strikes one as  $\varphi$ ' is intended as an intuitive gloss rather than a philosophical analysis. Some candidates for what it is for an experience  $\alpha$  to strike one as  $\varphi$  include one being disposed to believe that  $\alpha$  is  $\varphi$  (see Werner [2014]), one having the intuition that  $\alpha$  is  $\varphi$  (see Bengson [2015]), or  $\alpha$  having presentational phenomenology that  $\varphi$  (see Chudnoff [2012]). Note that if  $\alpha$  does *not* strike one as  $\varphi$ , then that doesn't necessarily mean that  $\alpha$  strikes one as not  $\varphi$ , nor that one isn't in a position to know that  $\alpha$  is  $\varphi$ . And striking one as  $\varphi$  might well be a matter of degree, in which case I'll assume that sentences of the form ' $\alpha$  (epistemically) feels  $\varphi$ ' are true just in case  $\alpha$  strikes one as  $\varphi$  to a sufficiently high degree.

---

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 noted 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 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<sub>17</sub> and the adjacent pixel red<sub>54</sub>. 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<sub>17</sub> and red<sub>54</sub>. 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 Contains-Gaps Objection

Here's the second objection:

### ⊥ The Contains-Gaps Objection

**P1:** If an experience is discrete, 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 seems to follow from the nature of discreteness and P2 sounds tautological. But there's an obvious question here: what does it mean for an experience to contain gaps?

Here's my analysis from earlier: what it is for experience to be gappy in feature F with respect to feature G is for some intermediate F-values to be 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 'gap' that would render P1 true. Let's call the characterization of 'gap' defined above the *missing-values definition*. The natural alternative is the *discontinuity definition*, which says that  $\alpha$  is gappy in F with respect to G just in case  $\alpha$  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  $\alpha$  is contiguous in color experience with respect to spatial experience and that  $\text{red}_1$  and  $\text{red}_2$  are adjacent color values instantiated by  $\alpha$ . Is

---

$\alpha$  gappy? In other words, does  $\alpha$  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  $\alpha$  doesn't have this feature since there are no color experiences between  $\text{red}_1$  and  $\text{red}_2$ . This is reason to think that  $\alpha$  isn't the kind of experience that would strike its subject as gappy, which is evidence that  $\alpha$  isn't gappy. To hold otherwise, one would have to say that  $\alpha$  is gappy even though it's impossible to "fill in" that gap. Since only the missing values definition classifies  $\alpha$  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 that 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

My arguments have some implications for introspection. Let's start with an obvious consequence. Since I've argued that introspection leaves open whether smooth experiences are continuous or discrete, my arguments entail that the following principle is false:

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.

Now, an argument against LUMINOSITY is nothing new—plenty of philosophers have argued that the principle is false. But I think the case of continuity and discreteness provides an especially compelling case against LUMINOSITY.

Here's a more interesting consequence: 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 is needed to account for the introspective data, while the continuous theory entails that there is infinitely more structure to experience than what introspection reveals.

To see why, note first that our introspective capacities are finite: for any given experience, we 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 is an infinite number of 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 is only a finite number of distinct F-values.

The point above is dialectically significant. Those concerned with quantitative parsimony may take this as a strike against the continuous theory. But more importantly, the point undercuts a potential criticism of the discrete theory. Most philosophers 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 positing more structure to our experience than what introspection reveals. But it turns out that the continuous theory faces the same cost. In fact, it's instead the discrete theorist who is able to deny that smooth experiences involve non-introspectable changes. 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.

This is a striking turnaround. At the beginning of the paper, it seemed that introspection is on the side of the continuous theorist. The primary task of the paper was to develop a discrete model of smooth experiences that's phenomenologically adequate. But now it's the 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.

## §11 Conclusion

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