Is Consciousness Continuous?

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

WORD COUNT: 10,785
KEYWORDS: continuity of consciousness, formal phenomenology, continuous, discrete, structure of experience, phenomenal sorites, stream of consciousness

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. A number of philosophers have contended that these kinds of conscious experiences have continuous (rather than discrete) structures. This paper first explains what it means to say that conscious experiences have continuous structures. Then I argue that a better diagnosis of the aforementioned kinds of experiences is that they are *contiguous*, where this means roughly that adjacent values from one experiential domain map to adjacent values in another experiential domain. I explain why every experience that is contiguous must be discrete, and I argue that our introspective evidence favors the discrete theory rather than the continuous theory. Along the way, I use formal tools to precisify the notions of continuity, discreteness, and other relevant structural properties.

Introduction

A number of philosophers have contended that 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.
>
> Wilfred Sellars, "Philosophy and the Scientific Image of Man" [1963]

In fact, nothing as exotic as a pink ice cube is needed to illustrate the idea. Consider your visual experience of a blue sky on a cloudless day, your tactile experience when your whole hand is pressed against a surface, your auditory experience of a rising pitch rises, your thermal experience as the room temperature gradually increases, and your temporal experience of an interval of time. These experiences likewise strike many philosophers as continuous, rather than discrete.

The principal motivation for the continuous theory comes from the fact that introspection does not reveal any discontinuities in the phenomenal characters of such experiences. Since what it is to be continuous is to not have any discontinuities, it seems to follow that those experiences are continuous. And notice how this argument from introspection is markedly stronger than a structurally analogous argument from perception: though the sky may appear continuous, that gives us little reason to believe that the sky is in fact continuous, since obviously the way the sky appears can deviate from the way the sky actually is. On the other hand, if an experience of the sky seems continuous, then it is hard to dismiss the idea that the experience is in fact continuous. Even if we acknowledge the limits of introspection, it still seems that the continuous theory better accounts for the phenomenology.

This paper has two main goals. The first is to explain what exactly it means to say that conscious experiences have continuous structures. Prior discussions of the issue usually provide only cursory glosses of continuity. Once we start to examine the question more carefully, we will quickly encounter a number of complexities. A central theme in the paper is the methodology of *formal phenomenology*, or the application of formal tools to the study of conscious experiences. We will see how formal tools can elucidate philosophical questions about continuity and discreteness, as well as other philosophically relevant structural properties that I call 'smoothness', 'gappiness', 'adjacency', and 'contiguity'.

The second aim of this paper is to argue against the idea that the continuous theory does better than the discrete theory in capturing the phenomenology of the aforementioned experiences. I develop as an alternative hypothesis the idea that such 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, for any structure that is contiguous must also be discrete. Towards the end, I argue for the somewhat surprising conclusion that the discrete theory better aligns with our introspective evidence.

Whether conscious experiences are continuous or discrete has implications for a number of other philosophical issues, including issues about the limits of introspection, the contents of experience, the phenomenal sorites, and the physical correlates of consciousness.[1] Although limitations of space preclude directly addressing these other issues, it will be straightforward to draw out the implications of my arguments. Similarly, although the focus of this paper is on the structure of conscious experiences, most of my analyses will generalize to analogous debates about the structure of space, time, matter, and representation.[2]

§**1** demarcates two classes of experiences, which I call 'smooth' and 'gappy'; **§2** and **§3** explain what it means to say that conscious experiences are continuous and present the main argument in favor of the continuous theory; **§4** defines contiguity; **§5** develops an analysis of what makes an experience smooth versus gappy; **§6** and **§7** defend the claim that all contiguous experiences are smooth; **§8**

---

[1] For connections, see Sellars [1963], Maxwell [1975], Clark [1989], Dennett [1993], Lockwood [1993], VanRullen & Koch [2003], and Builes [2020] on the physical correlates of consciousness, Wright [1975] and Fara [2001] on the phenomenal sorites, and Blakemore [2002], Tye [2003], Dainton [2006, 2014], Rashbrook [2013], and White [2018] on temporal experience.

[2] 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 representational format. 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.

discusses whether other authors have used the term 'continuous' in a different sense; and **§9** argues that the discrete theory better aligns with our introspective evidence.

§1 | Smooth vs. Gappy

To start, let us draw a contrast between two classes of experiences. The first class consists of the kinds of experiences mentioned at the beginning of this paper—your visual experience of the clear blue sky, your auditory experience of a rising pitch, your tactile experience when your whole hand is pressed against a surface, your thermal experience as the temperature gradually rises, and your temporal experience of an interval of time. It will be useful to invent a term to denote this class of experiences that is neutral on whether they are continuous or discrete. Call these *smooth experiences*.

As a contrast class, consider the visual experience you have when looking at a pixelated low-resolution image, your color experience of a patch that is red on the left side and green on the right side, your auditory experience of C, then F#, then A♭ (and no notes in between), your tactile experience when only your palm and fingertips are pressed against a surface, your thermal experience when jumping into a pool on a hot day, and a temporal experience of a syncopated series of temporal snapshots. These are all experiences that involve abrupt changes in their phenomenal character. Call these *gappy experiences*.

Let the *continuous theory* be the view that smooth experiences are continuous, and the *discrete theory* be the view that smooth experiences are discrete. Both the continuous theorist and the discrete theorist agree on which experiences are smooth and which are gappy, since 'smooth' and 'gappy' are invented terms whose extensions are determined by ostensions. The disagreement is about the structure of smooth experiences: in particular, about whether smooth experiences are continuous. I will eventually offer a precise analysis of smoothness and gappiness, but for now we will work with the intuitive characterization sketched above.[3]

[3] I will remain neutral on questions about the physical correlates and metaphysical nature of conscious experiences. At times, I will talk about experiences in terms of what they represent, but this should be taken as merely a convenient way of talking about phenomenal character. For most of this paper, I will treat questions about whether the phenomenal characters of

I suspect the continuous theory will strike many as compelling. It is natural to assert that your color experience of the color gradient varies continuously, that your spatial experience of a spatial region extends continuously, and that your temporal experience of an interval of time flows continuously. In fact, the discrete theorist may even be thought to be 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 is no analogous phenomenon of introspecting an experience "close up" versus "far away". This means that the standard method for explaining away the appearance of continuity is unavailable for the case of conscious experiences.

§2 | State-Spaces

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

Q1: Are the state-spaces for conscious experiences continuous or discrete?
Q2: Are individual conscious experiences continuous or discrete?

Since the main motivation for the continuous theory appeals to introspection, and since it is 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. However, to evaluate Q2, we will need to first understand Q1.[4]

experiences are continuous as equivalent to questions about whether the representational contents of experiences are continuous (though see my remark on p. 26).

[4] There are other interpretations of the sentence 'consciousness is continuous' that are less relevant here. Consider Q3: is consciousness itself continuous? Suppose consciousness is a quantity, meaning that creatures can be more or less conscious. Then we can ask whether that quantity is continuous. Or consider Q4: is the stream of consciousness continuous? This concerns the temporal structure of experience, rather than the structure of temporal experience (which arguably falls under Q2). For more on these questions, see Bayne, Hohwy, & Owen [2016] (on Q3) and Rashbrook [2013] (on Q4).

A *state-space* is a structured set of the possible states that a system can be in. Consider, for example, the state-space for color experience, which is canonically characterized as a bounded three-dimensional space with hue, saturation, and brightness as dimensions, where color experiences that are more similar to each other are closer to each other in the space. Each state-space may be thought of as corresponding to a different feature of conscious experiences: for example, there are state-spaces for color experience, auditory experience, spatial experience, and so forth. More abstractly, any state-space can be thought of a set of elements and a distance metric capturing the similarity relations between those elements: for example, the state-space for color experiences has as elements red_1, red_2, and so forth, and has a metric that entails 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.[5] In $\mathbb{R}$, every element is perfectly connected every other, 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 structure; $\mathbb{Z}$ is a paradigm of a discrete structure.[6]

In philosophical contexts, questions about continuity are usually treated as equivalent to questions about infinite divisibility. 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. In the example above, any bounded region of $\mathbb{R}$ is infinitely divisible into smaller regions containing distinct real numbers, while any bounded region of $\mathbb{Z}$ contains only a finite number of integers. Analogously, the question of whether the state-spaces for conscious experiences are continuous can be taken to turn on whether any region of a given state-space is infinitely divisible into smaller regions each containing distinct elements from that domain.

---

[5] I assume the standard metrics for $\mathbb{R}$, $\mathbb{Z}$, and every other set I discuss.

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

Strictly speaking, infinite divisibility is not sufficient for continuity. Consider the set $\mathbb{Q}$ of rational numbers, which is infinitely divisible (any interval of $\mathbb{Q}$ contains multiple rationals) but not continuous (every interval of $\mathbb{Q}$ is missing all the irrational numbers). For our purposes, it is not particularly important to discuss in detail the differences between continuity and infinite divisibility.[7] This is because even though the properties are not mathematically equivalent, it is really questions about infinite divisibility that 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 not continuous. Given this, I will follow convention and assume that an experiential state-space is infinitely divisible just in case it is continuous.

Now we can specify what it is for a state-space to be continuous versus discrete. Let d be the distance metric telling us the distance between any two elements of any given state-space.[8] A state-space is infinitely divisible (and, by assumption, *continuous*) just in case it contains no isolated elements, meaning that for any element x and any distance δ, there is a distinct element y such that $d(x, y) < \delta$.[9] By contrast, a state-space is *discrete* just in case every element is isolated, meaning that for any element x there is some distance δ where no distinct element y is such

[7] I suspect the best way of capturing what is meant by the claim that a space is continuous is by appeal to the notion of a connected space. A space X is *connected* just in case it is not the union of two disjoint non-empty open subsets, where a subset A of X is *open* just in case every element of A lies within an open ball of X, where an *open ball* $B_\delta(a)$ of X is the set of elements in X that are less than distance δ from element a. FORMAL DEFINITION: a space X is connected $=_{def} \neg\exists A,B \subset X\ (A \cup B = X$ and $\forall a \in A, \exists\delta > 0 : B_\delta(a) \in X$ and $(\forall b \in B, \exists\epsilon > 0 : B_\epsilon(b) \in X)$.

[8] For the purposes of this paper, I will assume that all of the spaces under consideration have metrics. This assumption is plausible, since a metric is needed to capture the similarity relations between different experiences within any experiential domain.

[9] FORMAL DEFINITION: state-space X is **continuous** $\Leftrightarrow$ X is infinitely divisible $=_{def} \forall x \in X, \forall\delta \in \mathbb{R}$, $\exists y \in X (x \neq y$ and $d(x, y) < \delta)$. Note that this formulation is expressed as a biconditional, since the assumption that a state-space is continuous just in case it is infinitely divisible is substantive. See footnote 6 for a more direct definition of 'continuous space'.

that $d(x, y) < \delta$.[10] A noteworthy virtue of these analyses is that they apply even to features of experience that lack linear orders, such as hue.[11]

To simplify the discussion, I will assume that all the state-spaces under consideration are one-dimensional and either wholly continuous or wholly discrete. These simplifications will enable us to maintain focus on the philosophical questions without getting too entangled in formal definitions. For those interested in the more complex cases, it will be straightforward to see how my arguments and analyses generalize to state-spaces of arbitrary dimensionality and structure.[12] For similar reasons, I will 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 species. This ensures that questions about continuity and discreteness (and smoothness and gappiness) do not become relativized to particular creatures or species.

§3 | The Core Question

We can now turn to our core question of whether individual conscious experiences are continuous or discrete. This question requires a bit more analysis to understand. Let us begin with a simple observation: for any individual experience α and any feature of experience F, there will be a set of elements of the state-space for F-experiences that are instantiated in α and another set of elements of that state-space that are not instantiated in α. Suppose, for example, that you are looking at a red wall: your color experience might instantiate red_1–red_{100} from the state-space for

---

[10] FORMAL DEFINITION: state-space X is **discrete** $=_{def}$ $\forall x \in X, \exists \epsilon > 0, \forall y \in X \setminus \{x\}, d(x, y) > \epsilon$. Note that while discrete structures can contain infinitely many elements (consider $\mathbb{Z}$ or $\{1/n : n > 0\}$), no discrete structure is infinitely divisible (see footnote 9).

[11] By contrast, many prior discussions of continuity (such as Clark [1989]) appeal to *density*: for any x and y such that $x < y$, there is a z such that $x < z < y$. However, not only does density not generalize to non-linearly ordered structures, but it also does not even entail infinite divisibility: for example, $\{1/n : n > 0\}$ is dense but not infinitely divisible (see footnote 9).

[12] Strictly speaking, the dimensionality assumption means that instead of talking about the state-space for (say) color experience, we should talk about the state-space for (say) hue experience. But I think friendlier prose is worth the cost of some slight technical transgressions.

color experiences but not $blue_1$–$blue_{100}$ from that same state-space. Call the elements of the state-space for F-experiences that are instantiated by α the *F-values of α*.

Here is a hypothesis that is initially attractive but turns out to be false: an individual experience α is continuous in feature F just in case α instantiates a continuous set of F-values, meaning that α instantiates exactly the values within a continuous region of the state-space for F-experiences. Call this the *continuum condition*.[13] The following example illustrates why satisfying the continuum condition is not sufficient for continuity. Suppose that the state-space for gray experience is continuous (ranging from pure white to pure black), and suppose α instantiates every gray quality within a continuous region of that state-space. However, suppose there is no systematic correspondence in α between the darkness of a gray quality and the spatial location associated with that quality. You might imagine α as similar 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 gray experience, α is not continuous in gray experience. This tells us something important about what it means to ascribe continuity to individuals.

In mathematics, the term 'continuity' is standardly used to denote a property of functions, where *continuous functions* are those where 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."[14] At first, it is not obvious whether this function-theoretic notion of continuity relates in any interesting way to the notion of continuity deployed by philosophers. After all, it seems that continuity is being ascribed to fundamentally different kinds of things: functions by mathematicians and worldly things (such as space, time, matter, or conscious experiences) by philosophers. However, the observation about the continuum

---

[13] FORMAL DEFINITION: Let F_α be the set of F-values instantiated by α and d be the metric for the state-space for F-experiences. Then α satisfies the **continuum condition** with respect to F-experience $=_{def}$ F_α is continuous under d (see footnote 8).

[14] Spivak [2008, p.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 $|f(x) - f(a)| < \epsilon$. For a generalized definition of continuity as a property of individuals, see footnote 18.

condition puts us in position to see how these seemingly distinct senses of continuity come together.

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

Are there any cases where we can simply ask whether an experience α is continuous in feature F, without relativizing the question to a feature G? The best candidates for this are "locational" experiences, such as spatial and temporal experience. A natural thought is that so long as α satisfies the continuum condition for spatial experience, α is continuous in spatial experience. Perhaps that is right, but it seems to me the answer depends on some unresolved questions, such as whether spatial and temporal experience are modality-specific.[15] Since addressing such questions would take us away from the main goals of this paper, I will 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, it will be straightforward to generalize my analyses.

We are now in position to see why the relevant formal tool for thinking about whether individual experiences are continuous is the function-theoretic notion of continuity used by mathematicians. To say that α is continuous in F with respect to G is to say that the mapping from α's G-values to α's F-values is a continuous function, where this means that sufficiently small changes in α's G-values map to arbitrarily small changes in α's F-values. Now, there is a question of what exactly it

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

means for a G-value to map to an F-value. But I take the notion to be intuitive enough, and it is easy to interpret particular cases. If feature F is color experience and feature G is spatial experience, then continuity means that so long as we consider a small enough interval of spatial experience, we will be able to find changes in color experience that are as little as we like.

If the mapping from α's G-values to F-values is not even a function, then α is not continuous in F with respect to G. Suppose, for example, that F is spatial (instead of color) experience, that G is color (instead of spatial) experience, and that red_1 maps to both l_1 and l_2. Then it is intuitively false that α is continuous in spatial experience with respect to color experience (though it may still be true that α is continuous in color experience with respect to spatial experience). A speculative thought: there may be structural constraints on which features can play the F-role or the G-role. Consider, for example, the idea that each value along the state-spaces for locational experiences (such as spatial and temporal experience) can be instantiated only once per total experience.[16] If we then accept the already intuitive claim that only locational experiences can play the G-role, we arrive at the propitious result that every mapping from G-values to F-values is a function. For limits of space, I will not speculate further on which features of experience can play the F-role or the G-role. And in any case, the analyses in this paper are fully general, meaning they do not rely upon any particular view about such constraints.

We are now close to a precise definition of what it means to say that an individual experience is continuous. However, there remains one last technical complication. Any function with a discrete domain trivially satisfies the mathematical definition of continuity. But this means that some features of individual experiences that intuitively are not continuous will count as continuous given our present analysis. Suppose, for example, that the state-space for spatial experience is discrete and that experience α represents red at spatial location l_1, green at l_2, and blue at l_3. It would be bizarre to say that α is continuous in color experience with respect to spatial experience. Yet it turns out, given the formal definition of a continuous function, that the function mapping α's G-values to α's F-values is continuous.

---

[16] See Clark [2000, Ch. 3] for a defense of this idea.

Fortunately, we have more or less already encountered the solution to this problem. We need simply require that α satisfies the continuum condition for *feature* G, meaning that α's G-values are exactly those within a continuous region of the state-space for G-experiences. In other words, the requirement is that the domain of the function from G-values to F-values be continuous. In fact, this requirement not only solves the technical problem described above, but also forges a neat connection between continuity of individual experiences and continuity of state-spaces: in order for α to be continuous in F with respect to G, the state-space for G-experiences must itself be continuous.[17] And with this condition, we finally have a definition of what it means to ascribe continuity to an individual conscious experience:

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

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

The remaining task is to define what it is for an individual experience α to be discrete in feature F with respect to feature G. There is no standard definition of what it means for a function to be discrete, but the most 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 understanding 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 does not overgeneralize: for example, we do not want α to trivially count as discrete in taste with respect to emotion simply because there is no

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

[18] FORMAL DEFINITION: Let G_α be the set of G-values instantiated by α that map onto F-values, and let f_i be the F-value mapped by any $g_i \in G_\alpha$. Then α is **continuous** in F with respect to G $=_{def}$ **(1)** α satisfies the continuum condition with respect to G-experience (see footnote 8), 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 $d(f_a, f_n) < \epsilon$.

mapping from α's emotional experience values to α's taste experience values. With this, we can define discreteness of individual experiences as follows:

experience α is **discrete** in feature F with respect to feature G $=_{def}$[19]

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

This brings us to a somewhat subtle point: 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, that there are infinitely many values between red_1 and red_{10} and between l_1 and l_{10}, and that α instantiates red_1 at l_1, red_2 at l_2, red_3 at l_3, and so forth. Then α is discrete in color experience with respect to spatial experience, even though the relevant state-spaces are continuous. As an analogy, consider how any function f from $\mathbb{Z}$ to $\mathbb{R}$ is discrete, even though $\mathbb{R}$ is continuous. In this case, f's domain (i.e. $\mathbb{Z}$) is analogous to α's G-values, f's codomain (i.e. $\mathbb{R}$) is analogous to the state-space for F-experiences, and f's image (i.e. the set of elements of $\mathbb{R}$ that are outputs of f) is analogous to α's F-values. Although this point may feel a bit technical, the underlying idea is intuitive: so long as α instantiates a discrete set of G-values, it is guaranteed that all of α's corresponding F-values are isolated from one another. That ensures that α is discrete in F with respect to G.

We are now ready for the argument in favor of the continuous theory. To my knowledge, the argument below has never been formulated explicitly in the philosophical literature. But I think it captures the phenomenological considerations that motivate the continuous theory, and it will take a good amount of work to appreciate where the argument goes awry:

---

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

⊥ **The Argument for Continuity**

P1: Some experiences are smooth.

P2: Smooth experiences are not gappy.

P3: An experience is either continuous or discrete.

P4: If an experience is discrete, then it is 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 and as contrasting properties: 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 us restrict the quantifier to experiences that are either wholly continuous or wholly discrete. The premise I wish to challenge instead is P4. Call this the *discrete-implies-gappy premise*. I will argue, contra this premise, that some discrete experiences are smooth.

§ 4 | Contiguity

To develop my argument against the discrete-implies-gappy premise, I need to define a new structural property, which I call 'contiguity'. To get a grip on contiguity, it is useful to contrast two classes 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

Call two integers a and b *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. By contrast, the B-sequences are discontiguous because some intermediate integers are missing: for example, sequence B1 jumps from 3 to 5. The notion of a sequence may initially seem unrelated to the formal tools we have invoked so far, but notice that sequences are basically 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 is easy to see that for contiguous sequences of integers, adjacent G-values (i.e. the indices of the sequence) map to adjacent F-values (i.e. the integer at a given index).

The notion of contiguity can be generalized beyond sequences of integers. To do this, we need to first generalize the notion of adjacency. Call two elements x and y of a state-space *adjacent* just in case there are no other values in between them.[20] Next, we need the basis for an analogue to the continuum condition (a contiguum condition, so to speak). Let a *contiguous region* of a state-space be 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.[21] 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 us say that 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:[22]

---

[20] FORMAL DEFINITION: Let F be a state-space and $f_1, f_2 \in F$. Then f_1 is **adjacent** to $f_2 =_{def} \neg\exists f_3 \in F$ $[d(f_1, f_2) > d(f_1, f_3), d(f_2, f_3)]$. To generalize to multidimensional state-spaces, we would need to first reinterpret this definition as characterizing adjacency between values of dimensions, and then take two elements of a state-space to be adjacent just in case each of their values along its dimensions are adjacent.

[21] FORMAL DEFINITION: Let F be a state-space and $F_\alpha \subseteq F$. Then region F_α is **contiguous** $=_{def}$ $\forall f_1, f_2 \in F_\alpha, \exists f_a \ldots f_z \in F_\alpha$ (f_1 is adjacent to f_a, f_a is adjacent to… f_z, f_z is adjacent to f_2).

[22] Lee [forthcoming] argues that experiences should be modeled using regions, rather than individual elements, in the respective state-spaces. This complicates the analysis of contiguity, since it is not obvious how to generalize contiguity to regions. However, I suspect it is possible to construct a degree-theoretic analogue of contiguity for regions, where the "degree of contiguity" between two regions is the proportion of elements in each region that are contiguous to some element within the other region.

experience α is **contiguous** in feature F with respect to feature G $=_{def}$[23]

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

Here are some toy examples to illustrate how contiguity applies to experiences. Suppose that the state-spaces for both color experience and spatial experience are discrete. Consider four experiences α, β, γ, and ζ. The first two, α and β, are contiguous in color with respect to space: α is as of a spectrum of color, and represents red_1–red_{50} at locations l_1–l_{50} (respectively), while β is as of a homogenous field of color, and represents red_1 at locations l_1–l_{50}. The second two, γ and ζ, are discontiguous in color with respect to space: γ is as of a series of colors with every other color skipped over, representing red_1 at location l_1, red_3 at l_2, red_5 at l_3, and so forth, while ζ is as of a spectrum of colors with a gap, representing red_1–red_{50} (except for red_{37}) at locations l_1–l_{50} (except for l_{37}).

Contiguity and continuity are mutually exclusive.[24] In order for an experience to be contiguous, it 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. Likewise, in order for an experience to be continuous, it 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. Therefore, these definitions yield a satisfying connection between contiguity of individuals and contiguity of state-spaces, and between continuity of individuals and continuity of state-spaces.

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.

---

[23] FORMAL DEFINITION: Let G_α be α's set of G-values that map onto F-values, and let f_i be the F-value mapped by any $g_i \in G_\alpha$. Then α is **contiguous** in F with respect to G $=_{def}$ **(1)** G_α is contiguous (footnote 21), 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.

[24] The sole exception is the degenerate case where α 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.

§ 5 | The Analysis of Smoothness

To make the case that contiguous experiences are smooth, I need to develop a more precise analysis of smoothness and gappiness. Let us 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♭, 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 $\mathbb{R}\setminus\mathbb{Z}$, or the real numbers minus all the integers, which is gappy (it is missing all the integers) but not discrete (its elements are not all isolated from one another).

A better diagnosis is that gappy experiences are those where some intermediate values are missing. Your gappy visual experience is missing intermediate color values between adjacent pixels, your gappy auditory experience is missing C#, D, and D#, and your gappy tactile experience is missing tactile sensations for the spatial locations between your fingertips. More generally, gappy experiences are those where some intermediate F-values are missing at the relevant G-locations. To make this precise, let z be *between* x and y just in case $d(x,y) > d(x,z), d(y,z)$. Then:

experience α is **gappy** in feature F with respect to feature G ⇔[25]

- α instantiates non-adjacent values g_1 and g_2 mapping to f_1 and f_2.
- for some f_n between f_1 and f_2, no g_n between g_1 and g_2 maps to f_n.

---

[25] FORMAL DEFINITION: Let G_α be α's set of G-values, let F be the state-space for F-experiences, and let f_i be the F-value mapped by any $g_i \in G_\alpha$. Then α is **gappy** in F with respect to G ⇔ **(1)** $\exists g_1, g_2 \in G_\alpha$ (g_1 is not adjacent to g_2), and **(2)** $\exists f_n \in F$ (f_n is between f_1 and f_2) and $\neg\exists g_n \in G_\alpha$ (g_n is between g_1 and g_2)).

The first condition precludes degenerate cases (where an experience instantiates only a single G-value) from counting as gappy. The second condition specifies that some intermediate F-values are missing at the relevant G-locations. A virtue of this analysis is that it yields the right predictions whether the continuous theory or the discrete theory is true, since both the continuous theorist and the discrete theorist can make sense of the idea of missing intermediate values.

The natural corollary hypothesis is that smooth experiences are those where all intermediate F-values are instantiated at the relevant G-locations. This characterization is intuitively plausible, and correctly categorizes all the smooth and gappy experiences mentioned at the beginning of the paper. To construct a precise analysis, we again need a first condition precluding degenerate cases and a second condition specifying that all intermediate F-values are instantiated at the relevant G-locations:

experience α is **smooth** in feature F with respect to feature G $\Leftrightarrow$[26]

- α instantiates non-adjacent values g_1 and g_2 mapping to f_1 and f_2.
- every f_n between f_1 and f_2 is mapped to by some g_n between g_1 and g_2.

I will soon say more to defend these analyses. But first, observe that both continuity and contiguity satisfy the analysis of smoothness. If the relevant state-spaces are continuous, then lacking gaps is a matter of continuity; if the relevant state-spaces are discrete, then lacking gaps is a matter of contiguity. Since contiguous experiences are discrete, this means that discrete experiences can be smooth. By consequence, the discrete-implies-gappy premise is false, and the Argument for Continuity is unsound.

It is worth highlighting that this analysis of smoothness is structural rather than epistemic. From my own experience, those who express sympathy towards the discrete theory nearly always appeal to limits in our introspective capacities.[27] In

---

[26] FORMAL DEFINITION: Let G_α be α's set of G-values, let F be the state-space for F-experiences, and let f_i be the F-value mapped by any $g_i \in G_\alpha$. Then α is **smooth** in F with respect to G $\Leftrightarrow$ **(1)** $\exists g_1, g_2 \in G_\alpha$ (g_1 is not adjacent to g_2), and **(2)** $\forall f_n \in F$ (f_n is between f_1 and $f_2 \rightarrow \exists g_n \in G_\alpha$ (g_n is between g_1 and g_2)).

[27] See Clark [1989] for the most explicit defense of the discrete theory that I have found.

particular, one might think that smooth experiences are those where the differences in phenomenal character between adjacent values are too small for subjects to introspectively discern. Although I am sympathetic to introspective limitations playing a role in explaining why a given experience strikes its subject as smooth, it seems to me that the best analysis of smoothness itself is structural rather than epistemic. An epistemic analysis like the one stated above yields the result that whether an experience is smooth or gappy depends on the introspective capacities of the subject. That strikes me as counterintuitive: at least, it seems to me that smoothness is more naturally understood in a way that is not subject-relative and that is applicable even for idealized subjects with perfect introspective capacities. Since my analysis of smoothness is structural, it avoids those sorts of issues.

§6 | The Feels-Discrete Objection

Let us now turn to two objections against my claim that contiguous experiences are smooth. Here is the first objection:

⊥ **The Feels-Discrete Objection**

P1: If an experience is discrete, then it feels discrete.

P2: If an experience feels discrete, then it is gappy.

—

C: If an experience is discrete, then it is gappy.

The argument is obviously valid, and each premise seems compelling on its own. But the problem is that 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. To see the equivocation, we need to distinguish 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.[28]

---

[28] The phrase 'strikes one as φ' is intended as an intuitive gloss rather than a philosophical analysis. Some candidates for what it is for an experience α to strike one as φ include being disposed to believe that α is φ (see Werner [2014]), having the intuition that α is φ (see Bengson [2015]), or α having presentational phenomenology with respect to φ (see Chudnoff

Here is 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 does not 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 does not mean that the visual experience strikes you as prime-colored: as an analogy, consider how you can know that π is a transcendental number if you think carefully enough, even though π does not 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 is 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. The missing premise is the following: 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 is gappy. Yet this is exactly the discrete-implies-gappy premise, which I have already argued against.

I suspect that part of the appeal of the feels-discrete objection comes from tacitly assuming that what it is like to have a discrete experience is structurally analogous to what it is like to imagine an experience as discrete. If I am asked to imagine an experience as discrete, then I imaginatively represent the discrete experience using a mental image of a pixelated image, where individual pixels

[2012]). Note that if α does *not* strike one as φ, then that does not necessarily mean that α strikes one as not φ, nor that one is not 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 assume that sentences of the form 'α (epistemically) feels φ' are true just in case α strikes one to a sufficiently high degree as φ.

correspond to the discrete units of the target experience. And that imaginative experience itself may well be gappy. But just because the experience of imagining an experience as discrete is gappy does not 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.

It may be tempting to respond by appealing to the following principle: in order to imagine an experience as φ, one's imaginative experience must itself be φ. I think this principle is dubious: for example, I can imagine an extremely painful experience without my imaginative experience itself being extremely painful. But even if the principle were true, it would be irrelevant in the present context. The principle entails that in order to successfully imagine an experience as discrete, one's imaginative experience must itself be discrete. But that does not mean that if the imagining of a discrete experience is gappy, then the discrete experience that is imagined must itself be gappy. The direction of entailment given by the principle is the opposite of the direction of entailment needed to secure the feels-discrete objection's missing premise. And the reverse principle—that if one's imaginative experience is φ, then the experience imagined must be φ—lacks even prima facie plausibility.

As an analogy, consider what happens when we imagine physical space as discrete. When I do so, I imagine a two-dimensional grid where each unit square represents a spatial atom. This sort of imaginative exercise can be useful: for example, my imaginative representation makes vivid the fact that the Pythagorean theorem does not hold if physical space is discrete.[29] However, it is obvious that not all structural properties of my imaginative representation of discrete physical space are structural properties of discrete physical space itself. Though spatial atoms are imaginatively represented via squares, spatial atoms are obviously not themselves squares since squares are extended whereas spatial atoms are not. Similarly, although it is natural to use an imaginative experience of a pixelated image as a representation of a discrete experience, that does not mean that what it is like to have a discrete experience is what it is like to see a pixelated image. The gappiness is a feature of the imaginative experience, rather than of the experience imagined.

---

[29] See Weyl [1949] on this argument.

§7 | The Contains-Gaps Objection

Here is the second objection against my claim that contiguous experiences are smooth:

⊥ **The Discrete-Implies-Gappy Objection**

P1: If an experience is discrete, then it contains gaps.

P2: If an experience contains gaps, then it is gappy.

—

C: If an experience is discrete, then it is 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 is an obvious question here: what does it mean for an experience to contain gaps?

Here is my analysis from earlier: what it is for experience α to be gappy in feature F with respect to feature G is for there to be some intermediate F-values that are not mapped to by any intermediate G-values. Since contiguous experiences map adjacent values to adjacent values, and since there are no intermediate values between adjacent values, contiguous experiences are not gappy. Supposing we interpret P2 as a tautology, it follows that P1 false: just because an experience is discrete does not 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 render P1 true. Call the characterization of 'gap' defined above 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. Obviously, there is some danger here of a verbal dispute. But we can avoid that danger by focusing on the substantive question: which definition best captures the class of experiences I originally labeled 'gappy'? Once we consider that question, the core problem with the discontinuity definition will become evident: namely, it yields plausible results only if we already presume that the relevant structures are continuous.

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

The advantage of the missing-values definition is particularly obvious if we consider non-experiential gaps. Suppose the physical world turns out to be fundamentally discrete. Then, according to the discontinuous 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); a blinking message in Morse code is gappy (in light with respect to time); a slice of cheese full of holes is gappy (in density with respect to area). The missing-values definition works whether the target structures are continuous or discrete; the discontinuous definition does not.

Here is one more way of illustrating the point. Suppose you have at least some credence that the discrete theory is true yet 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 that we labeled 'smooth'—are in fact gappy. Yet one of the central goals of this paper has been to explain the difference between the class of experiences I called 'smooth' and the class of experiences I called 'gappy'. If we were to simply classify every experience as gappy, then we have lost sight of that initial explanandum. After all, 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.

§8 | The Verbal Objection

One might object that I have been using the term 'continuous' in a different way than how other philosophers use the term. After all, it seems unlikely that every philosopher who has claimed that consciousness is continuous really intended to say that sufficiently small changes in some feature G map to arbitrarily small changes in some feature F. Let me make three points in response to this objection.

First, there is textual evidence of philosophers using the term 'continuous' in the sense discussed in this paper. Consider this paper's opening quote from Sellars [1963] describing the visual experience of a pink ice cube as involving a "pink continuum," or Clark [1989, p.277]'s remark 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," or Fara [2001, p.924]'s explicit appeal to the standard function-theoretic definition of 'continuity' used in mathematics, or Dainton [2014, p.130]'s characterization of temporal experiences as "phenomenal continua" that are infinitely divisible.[30] Moreover, if philosophers of consciousness were using 'continuity' in a non-standard way, then it would be puzzling why the "apparent continuity" of consciousness has often been thought to be in tension with the apparent discreteness of its physical correlates.

Second, speaker meaning is not the same as semantic meaning. Most casual ascriptions of continuity are probably not sensitive to the sorts of distinctions that have been drawn in this paper. But that does not mean that we should adopt a non-standard interpretation of 'continuous' when evaluating 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 'consciousness is continuous' (and provides no explicit

[30] See Rashbrook [2013] for more thorough argument that various philosophers discussing the continuity of consciousness have had in mind the sense of continuity deployed in this paper (i.e. where continuous structures contain no gaps and are infinitely divisible).

specification of what they mean by 'continuous'), then we should interpret 'continuous' in the standard sense when evaluating their claim.

Third, using the term 'continuous' in a non-standard way would be a step backwards in conceptual clarity. Suppose we were to use 'continuous' to mean something looser, such as 'continuous or contiguous', or 'strikes one as continuous', or 'not gappy'. 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 versus the sense of 'continuity' invoked by mathematicians, scientists, metaphysicians, and philosophers of science. Instead of proliferating senses of 'continuity', it is better to respect the established terminological standards.

§9 | Introspection Favors the Discrete Theory

This paper began by remarking that our introspective evidence seems to favor the continuous theory. My goal so far has been to explain why the discrete theory can account for the phenomenological differences between smooth and gappy experiences. In this final section, I will explain why I think the introspective evidence actually favors the discrete theory.

We need to start by defining a few more terms. Let an *introspective capacity* be a capacity to make introspective discriminations between experiences within an experiential domain, where two experiences are *introspectively discriminable* by a subject just in case that subject can know solely on the basis of introspection that those two experiences are distinct. Let an introspective capacity be *infinite* just in case it enables infinitely many introspective discriminations, and *finite* otherwise. Analogously, let an experience be *infinite* just in case it instantiates infinitely many values of a state-space, and *finite* otherwise. Now here is the argument:

— **An Argument for Discreteness**[31]

P1: If our introspective capacities are finite, then our experiences are finite.

P2: Our introspective capacities are finite.

P3: If our experiences are finite, then they are discrete.

—

C: Our experiences are discrete.

Let us take the premises in reverse order. P3 is justified by the following observation: if the continuous theory is true, then any experience that is continuous in F with respect to G and instantiates at least two distinct G-values is infinite. This is because in continuous structures there are infinitely many elements between any two distinct elements. The only exceptions are degenerate cases where an experience instantiates only a single F-value and a single G-value: for example, consider an experience that represents only a single color at a single spatial location. Since few or none of our experiences are like this, I think we can set aside such cases and presume that continuity implies infinitude.

P2, it seems to me, can be accepted with something approaching certainty. We are finite creatures, and given that all of our other mental capacities (perception, cognition, memory, imagination, and so forth) are finite, there seems little reason to think that our introspective capacities are by contrast infinite. If our introspective capacities were infinite, then we should expect to be able to make introspective discriminations at arbitrarily fine levels of precision. Consider, for example, your color experience of a color gradient: given infinitude, you should be able to make introspective discriminations between arbitrarily fine slices of your color experience. But if we can partition your color experience as fine as we like, then we are guaranteed to eventually reach a point where your introspective capacities cannot discriminate between color values across sufficiently close segments.

Finally, I think P1 is best justified by appeal to an evidential connection between the grain of introspection and the grain of experience. If a subject can make x number of introspective discriminations over their color experiences, then that is

[31] Clark [1989] also argues for the discrete theory by appeal to limits in our introspective capacities. The principal differences between Clark's view and mine consist in the arguments against the continuous theory and in the diagnoses of smoothness (see p.17).

evidence that the subject can undergo x distinct kinds of color experiences. This inferential scheme is based upon the following principle: ascribe no more structure to experience than what is needed to best explain the introspective data (at least unless we have compelling reason to think otherwise). The stronger one takes that imperative to be, the stronger the case for the discrete theory.

Nevertheless, P1 strikes me as defeasible. There are reasons to think that the grain of experience outstrips the grain of introspection, meaning that our introspective capacities are not sensitive to every difference in phenomenology.[32] The extent to which this is true turns on some unresolved theoretical issues, but the possibility opens the door for ascribing more structure to experience than what is given by our introspective observations. However, notice that in order for the Argument for Discreteness to be defeated, the grain of experience would have to be *infinitely* finer than the grain of introspection. That is a much stronger hypothesis. While I do not think we can rule out such a hypothesis a priori, I also do not think the support for that hypothesis surpasses the evidential weight in favor of P1. Therefore, I think the current balance of considerations favors the discrete theory, even acknowledging that the continuous theory may nevertheless turn out true.

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 is worth briefly reiterating some of the upshots from our earlier discussions. First, remember that 'continuous' does not merely mean 'smooth'—instead, it denotes the kind of structural property defined earlier and that is referenced in mathematics, science, and other parts of philosophy. Second, remember that the fact that an experience is discretely structured does not entail that the experience will strike its subject as discretely structured. Third, remember that what it is like for one to imagine an experience as discrete need not be what it is like to actually undergo the discrete experience that is imagined. Speaking for myself, once I fully appreciate these facts, I lose any feeling that the discrete theory cannot be phenomenologically adequate.

These results have implications for a puzzle concerning the physical correlates of experience. Many philosophers have found the following triad of claims to be compelling: (1) phenomenal properties supervene on physical

[32] See Clark [1989], Fara [2001], Schwitzgebel [2006], and Lee [2019] for a few discussions.

properties, (2) empirical investigation reveals the physical correlates of consciousness to be discrete, and (3) introspection reveals consciousness to be continuous. Although these claims are not strictly speaking contradictory, it is easy to see the internal tension: if the physical correlates of consciousness are discrete and supervenience is true, then we ought to expect consciousness to be discrete as well. This paper provides a solution to the puzzle by explaining why we can justifiably reject (3). Consciousness might turn out to be continuous, but introspection provides no reason for thinking so.

One final thought: instead of asking whether your visual experience of the pink ice cube is continuous, we could have instead asked whether the pink ice cube looks continuous to you. The former question concerns phenomenal character; the latter concerns representational content. Although the questions are not quite equivalent, many of the arguments in this paper generalize to both issues. However, I wish to make one brief remark specifically about the latter issue. If the pink ice cube looks continuous to you, then your experience of the ice cube is veridical only if the ice cube is in fact continuous. But this would mean that the veridicality of your experience is beholden to the fundamental structure of reality. Yet on the contrary, it seems that your experience of the pink ice cube is veridical just in case you see a pink ice cube, irrespective of whether the ice cube is continuous or discrete. Just as the fact that some experiences are smooth does not entail that those experiences are continuous, the fact that some physical objects look "smooth" does not entail that those physical objects look continuous. Given this, I think it is false to say that the pink ice cube looks continuous to you; instead, your experience is simply agnostic as to whether the ice cube is continuous or discrete.

Conclusion

This paper began with a contrast between smooth and gappy experiences. The continuous theorist says that smooth experiences are continuous; the discrete theorist says that they 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 is 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 is not 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.

The bulk of this paper focused on arguing for the negative claim that our introspective evidence does not support the continuous theory. The last section advanced the positive claim that our introspective evidence actually favors the discrete theory. Supposing that we ought not ascribe more structure to experience than what is needed to best explain the introspective data, the fact that our introspective capacities are finite is evidence that our experiences are discrete. Given this, I think the current balance of considerations favors the discrete theory.

A general goal of this paper has been to exhibit the prospects for formal phenomenology, or the application of formal tools to the study of conscious experiences. I have aimed to use such tools to sharpen our sight of the theoretical space and to more carefully assess the dialectical situation. These kinds of endeavors strike me as ways of making philosophical progress. Even when we do not yet know the answers, we can still move forward by advancing our understanding of the questions.

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.

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.

Lee, Andrew Y. (2019). The Microstructure of Experience. Journal of the American Philosophical Association 5 (3):286-305.

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

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.

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.

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.

Weyl, Hermann (1949). Philosophy of Mathematics and Natural Sciences. Princeton University Press.

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.