Consciousness and Continuity

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

WORD COUNT: 8401
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 experiences have continuous (rather than discrete) structures. This paper first explains what exactly it means to say that conscious experiences have continuous structures. Then I argue that the aforementioned experiences might instead be *contiguous*, where this means roughly that adjacent values of one experiential domain map to adjacent values of another experiential domain. I explain why every experience that is contiguous must be discrete, and I argue that our introspective evidence better aligns with the discrete theory than with the continuous theory. Throughout the paper, I use formal tools to develop precise analyses of continuity, discreteness, contiguity, and other relevant structural properties.

Introduction

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

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

In fact, nothing as exotic as a pink ice cube is needed to illustrate the idea. Consider 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 over 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 doesn't reveal any discontinuities in the phenomenal characters of our experiences. Since what it is for something to be continuous is for it to lack any discontinuities, the continuous theory seems to do better than the discrete theory in accounting for the phenomenology. Moreover, 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 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's hard to dismiss the idea that the experience is in fact continuous.

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, though, we will quickly encounter a number of complexities. A methodological aim of this paper is to show how formal tools can

[1] In this passage, Sellars' claim is about the contents of such experiences. But it's clear from context that Sellars is also making a claim about phenomenal character.

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 contest the idea that the continuous theory does better than the discrete theory in capturing the phenomenology of our experiences. I argue that such experiences might instead be *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. Though it may still turn out in the end that conscious experiences are in fact continuous, I'll argue that introspection provides no reason for thinking so.

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. These other issues will be mentioned only in passing, though it will be straightforward to draw out the implications of my arguments for other debates. Similarly, although my focus is on the structure of conscious experiences, my discussions of the notions of continuity, discreteness, and other structural properties will be applicable to analogous debates about the structure of space, time, matter, and representational format.[2]

Here's the plan for the paper. §**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** argue that all contiguous experiences are smooth; **§8** discusses whether the issue is merely verbal.

A preliminary remark: whenever I talk about the structure of an experience, I'll have in mind its phenomenal character. By contrast, a good deal of the

[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.

philosophical literature on consciousness and continuity focuses mainly on whether the physical correlates (especially the temporal correlates) of experience are continuous.[3] Though I think it's plausible that the phenomenal character of an experience is continuous just in case the physical correlate of that experience is continuous, I won't take a stance on questions about physical correlates. Instead, my focus will be on the sorts of issues raised by the introspective argument sketched earlier.

§1 Smooth vs. Gappy

To start, let's 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 room temperature gradually rises, or your temporal experience of an interval of time. It will be useful to invent a term that denotes this class of experiences but 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 half your body is submerged in hot water, or a temporal experience of a syncopated series of momentary 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 the phenomenal characters of smooth experiences are continuous, and the *discrete theory* be the view that the phenomenal characters of smooth experiences are discrete. The continuous theorist and the discrete theorist will agree on which experiences are smooth and which are gappy, since 'smooth' and 'gappy' are invented terms whose extensions are

[3] 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]. See Rashbrook [2013] for some useful disentanglement of that question from the question of this paper.

determined by ostensions. The disagreement is about the structure of smooth experiences: in particular, about whether smooth experiences are continuous. I'll eventually offer a precise account of smoothness and gappiness, but for now we'll work with the basic characterization from above.

I suspect the continuous theory will strike many as compelling. It's natural to say 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.

A clarification, before moving forward. The arguments of this paper are intended to be neutral on the metaphysical nature of phenomenal character. For representationalists, the relevant question is whether the representational contents of smooth experiences attribute continuous structures. For relationalists, the relevant question is whether smooth experiences involve perceptual awareness of continuous structures of external objects. For qualia theorists, the relevant question is whether the qualitative characters of smooth experiences have continuous structures. I'll frame my discussion simply in terms of whether smooth experiences are continuous, but it will be straightforward to translate the discussion into any of these frameworks.

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

A *state-space* is a structured set of the possible states that a system can be in. Consider the state-space for color experience, which is a bounded three-dimensional space with hue, saturation, and brightness as dimensions, where color experiences that are more similar are located closer within the space. Alongside color experience, there are state-spaces for auditory experience, spatial experience, and any other feature of experience.[4] At an abstract level, a state-space consists of a set of elements and a distance metric capturing the similarity relations between those elements. 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.[5] 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 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. It's sometimes pointed out that infinite divisibility isn't sufficient for continuity: the set $\mathbb{Q}$ of rational numbers is infinitely divisible (any interval of $\mathbb{Q}$ contains multiple rationals) but not continuous (every interval of $\mathbb{Q}$ is missing all the irrational numbers). For our purposes, though, it's unnecessary to detail the differences between continuity and infinite divisibility.

[4] See Young, Keller, & Rosenthal [2014] on the state-spaces for olfactory experiences, and Clark [2000] for a more general discussion of state-spaces for features of experiences.

[5] I'll 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.

This is because questions about infinite divisibility lie at the heart of philosophical debates about continuity. To my knowledge, no contemporary philosopher has seriously argued that space or time or matter or consciousness is infinitely divisible yet not continuous. Given this, I'll follow convention and assume that the state-space for a domain of experiences is infinitely divisible just in case it's continuous.

Here's a more important technical distinction. Infinite divisibility is often characterized as the thesis that for any elements x and y such that $x > y$, there is an element z such that $x > z > y$.[7] However, some features of experience have a cyclic, rather than linear, structure. For example, one experience cannot be "greater in hue" than another. If we wish to develop a characterization of continuity that is applicable to all experiential state-spaces, then we need a more general analysis of infinite divisibility. We can develop that by appealing to distances. Let d be the metric telling us the distance between any two elements of any given state-space.[8] Then 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 that $d(x, y) < \delta$.[10]

To simplify the discussion, I'll assume that all state-spaces under consideration are either wholly continuous or wholly discrete. This will enable us to 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 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 species.

---

[7] For a relevant example, see Clark [1989].

[8] 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.

[9] FORMAL DEFINITION: Let's suppose that a state-space is continuous just in case it's infinitely divisible. Then state-space X is **continuous** $\leftrightarrow \forall x \in X, \forall \delta \in \mathbb{R}, \exists y \in X (x \neq y$ and $d(x, y) < \delta)$.

[10] FORMAL DEFINITION: state-space X is **discrete** $=_{def} \forall x \in X, \exists \epsilon > 0, \forall y \in X \setminus \{x\}, d(x, y) > \epsilon$.

§3 The Core Question

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

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

Here's a hypothesis that is initially attractive but that 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 α instantiates exactly the values within a continuous region of the state-space for F-experiences. In this circumstance, let's say that experience α satisfies the *continuum condition* with respect to F-experience.[11] In the example from above, if the state-space for color experience is continuous and if red_1–red_{100} is a continuous region of that state-space, then your experience satisfies the continuum condition with respect to color. The hypothesis says that what it is for an individual 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. Suppose 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 somewhat analogous to the kind of visual experience you have when looking at a noisy static image, such as a television screen with no signal. Even though α satisfies the continuum condition for gray experience, α isn't continuous in gray experience. This

---

[11] 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 fn. 9).

example 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 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."[12] At first, it's not obvious whether this function-theoretic notion of continuity relates in any interesting way to the notion of continuity we are interested in. 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 experiences) by philosophers. However, the counterexample from the previous paragraph enables to see how these seemingly distinct senses of 'continuity' come together.

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

Can any feature of experience play either the F-role or the G-role? Well, it's natural to take *qualitative features* (such as color and auditory experience) to play the F-role and *locative features* (such as spatial and temporal experience) to play the G-role. It's easy to grasp what it means for an experience to be continuous in color with respect to space; it's hard to grasp what it means for an experience to be continuous in space with respect to color. I suspect this intuitive difference arises from structural differences between qualitative features and locative features.[13] Given this, I'll

---

[12] Spivak [2008: 115]. FORMAL DEFINITION: a function f with domain X is continuous $=_{def}$ for any point $a \in X$, $\forall \epsilon > 0$, $\exists \delta > 0$ such that $\forall x \in X$, if $0 < |x - a| < \delta$ then $|f(x) - f(a)| < \epsilon$. For a generalized definition of continuity as a property of individuals, see fn. 16.

[13] See Clark [2000: Ch.3] for discussion of these structural differences.

sometimes talk of F-*values* being instantiated at G-*locations*. However, my arguments will work whether or not one thinks the F-role and the G-role are restricted to particular kinds of experiences.[14]

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

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

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

[14] A natural thought is that so long as α satisfies the continuum condition for any locative feature G, then α is continuous with respect to G. I'm sympathetic to this idea, but I'll restrict my focus to the cases where we must ask whether α is continuous in F with respect to G. If there turn out to be simpler cases where we need not relativize to a second feature, then it will be straightforward to generalize my arguments.

respect to G, the state-space for G-experiences must itself be continuous.[15] And with that condition, we arrive at the following definition:

experience α is **continuous** in feature F with respect to feature G $=_{def}$[16]

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

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

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

[16] 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[α]. Then α is **continuous** in F with respect to G $=_{def}$ **(1)** α satisfies the continuum condition with respect to G-experience (see fn. 11), and **(2)** $\forall\epsilon > 0$ and $\forall g_a \in$ G[α], $\exists\delta > 0$ such that $\forall g_n \in$ G[α], if $d(g_a, g_n) < \delta$ then $d(f_a, f_n) < \epsilon$.

experience α is **discrete** in feature F with respect to feature G $=_{def}$[17]

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

Now, what happens if the mapping from α's G-values to F-values isn't even a function? Well, this situation will arise only if multiple F-values are associated with the same G-value.[18] Suppose, for example, that F is spatial experience, that G is color experience, and that red_1 maps to both l_1 and l_2. Then, according to the analysis, α isn't continuous in F with respect to G (though it may still be true that α is continuous in color experience with respect to spatial experience). This is intuitively correct.

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 some work to appreciate how the argument goes awry:

⊥ **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's gappy.

—

C: Some experiences are continuous.

The argument is valid. Both P1 and P2 are uncontestable, since the terms 'smooth' and 'gappy' were defined ostensively: one can 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

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

[18] If only locative experiences can play the G-role (and if each location value can be instantiated only once per experience), then every mapping from G-values to F-values is a function.

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 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—contiguity—where the basic idea is that adjacent values of one domain map to adjacent values of another domain. To illustrate contiguity, it's useful to start with a contrast between two kinds of sequences of integers. The A-sequences below are contiguous, while the B-sequences are not:

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

Let's say two integers a and b are *adjacent* just in case either $a=b$ or $a=b\pm1$. The A-sequences are contiguous because every subsequent integer is adjacent to its predecessor in the sequence. The B-sequences are discontiguous because some intermediate integers are missing: for example, sequence B_1 jumps from 3 to 5. The notion of a sequence may initially seem unrelated to 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.[19] 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

[19] FORMAL DEFINITION: $f_1 \in F$ is **adjacent** to $f_2 \in F =_{def} \neg\exists f_3 \in F\ (d(f_1, f_2) > d(f_1, f_3)\ \&\ d(f_1, f_2) > d(f_2, f_3))$.

second element of each preceding pair.[20] 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 α *instantiates a contiguous set* of G-values just in case α instantiates all and only the elements within a contiguous region of the state-space for G-experiences. With these definitions in place, we can define contiguity in a way that parallels the prior definition of continuity:

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

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

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

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

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

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

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

Contiguity and continuity are mutually exclusive.[22] To be contiguous, an experience must instantiate a contiguous region of the state-space for G-experiences.

---

[20] FORMAL DEFINITION: Let F be a state-space and $F[\alpha] \subseteq F$. Then region $F[\alpha]$ 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).

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

[22] 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.

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.

§5 The Analysis of Smoothness

To argue that contiguous experiences are smooth, I need to first discuss 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 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 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 $\mathbb{R}\backslash\mathbb{Z}$, or the set of real numbers minus the set of integers, which is gappy (it's missing all the integers) but not discrete (its elements are not all isolated from one another).

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

experience α is **gappy** in feature F with respect to feature G ↔

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

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 $d(y, z) < d(x, y)$. Then it's natural to interpret the analysis in the following way: α is gappy in feature F with respect to feature G ↔ for any F-values f_1 and f_2 instantiated at G-locations g_1 and g_2 in α, 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 mapped to in α by any location between g_1 and g_2. This interpretation indeed works if we restrict our focus to one-dimensional linear state-spaces. But it breaks down if we begin thinking about multidimensional or cyclical state-spaces.

The following two challenge cases 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 α instantiates $\langle red_1–red_{10}, blue_1\rangle$ at locations l_1–l_{10} and $\langle red_{10}, blue_1–blue_{10}\rangle$ at locations l_{11}–l_{20}. In other words, as we move from l_1 to l_{20}, α's color values first shift in redness (while blueness remains fixed) and then shift in blueness (while redness remains fixed). Intuitively, α ought to count as smooth in color with respect to space. Yet α is missing some intermediate color values (such as $\langle red_5, blue_5\rangle$) 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 $hue_0 = hue_1$. Then suppose α instantiates $hue_{.05}$–$hue_{.95}$ at locations l_1–l_{90} (and doesn't instantiate any values between $hue_{.95}$ and $hue_{.05}$). Intuitively, α ought to count as smooth in hue with respect to space. Yet α is missing some intermediate hue values (such as hue_0) at intermediate locations.

To ensure that the analysis above applies to all cases, we need to be careful about what we mean by the claim that α is missing intermediate F-values at intermediate G-locations. The challenge 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 not mapped to 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 α's F-values are disconnected from each other relative to α's G-locations. Speaking loosely, if we were to track α's F-values as we move across α'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 α instantiates two F-values f_1 and f_2, then α also instantiates all the F-values between f_1 and f_2. But the challenge cases from above indicate that smoothness is compatible with α 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 α's F-values is connected relative to α'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 α is **smooth** in feature F with respect to feature G $\leftrightarrow$

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

These characterizations of gappiness and smoothness are intuitively plausible, and correctly categorize the smooth and gappy experiences mentioned in §1. And since both the continuous theorist and the discrete theorist can make sense of the idea of missing intermediate values, these characterizations of gappiness and smoothness are neutral between the continuous theory and the discrete theory.[23]

Here's 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

[23] 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.

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 this analysis of smoothness is structural, rather than epistemic. Those who express sympathy towards the discrete theory usually account for smooth experiences by appealing to limits in our introspective capacities.[24] 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 three objections to my arguments.

§6 The Feels-Discrete Objection

Here's the first objection:

⊥ **The Feels-Discrete Objection**

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

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

—

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

The argument is obviously valid, and each premise seems compelling on its own. But the argument equivocates. Although there is a reading of P1 that is true and a reading of P2 that is true, there is no univocal precisification of 'feels discrete' that renders both premises true. The equivocation is between two senses of 'feels φ':

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

(**1**) an experience *phenomenally* feels φ just in case it has phenomenal character φ, and (**2**) an experience *epistemically* feels φ just in case it strikes its subject as φ.[25]

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

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

---

[25] 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 one being disposed to believe that α is φ (see Werner [2014]), one having the intuition that α is φ (see Bengson [2015]), or α having presentational phenomenology that φ (see Chudnoff [2012]). Note that if α does *not* strike one as φ, then that doesn't necessarily mean that α strikes one as not φ, nor that one isn't in a position to know that α is φ. And striking one as φ might well be a matter of degree, in which case I'll assume that sentences of the form 'α (epistemically) feels φ' are true just in case α strikes one as φ to a sufficiently high degree.

I suspect that part of the appeal of the feels-discrete objection comes from a temptation 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 I'm 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 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.

§7 The Contains-Gaps Objection

Here's the second objection:

⊥ **The Discrete-Implies-Gappy 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 is 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 α 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 are not 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 α is gappy in F with respect to G just in case α is discontinuous in F with respect to G. These two definitions yield different diagnoses of contiguous structures: only the discontinuity definition says that contiguous structures are gappy. It's obvious that we are now in the vicinity of a verbal dispute. But we can avoid that trap by focusing on the substantive question: which definition best captures the class of experiences I originally labeled 'gappy'? I'll argue that the discontinuity definition yields plausible results only if we already presume that the relevant structures are continuous.

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

The advantage of the missing-values definition becomes more obvious when we consider non-experiential gaps. Suppose the physical world turns out to be fundamentally discrete. Then, according to the discontinuity definition, every physical structure contains gaps (since no physical structures are continuous). But there is clearly still a sense in which we can talk about some physical things containing gaps and other physical things lacking gaps. A wall that is half black and half white is gappy (in color with respect to space), while a wall that is uniformly black is not; a blinking message in Morse code 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 at least some credence that the discrete theory is true, yet suppose you 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 simply classify every experience as gappy, then we 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.

§8 The Verbal Objection

Here's the last objection:

> **The Verbal Objection:** The way this paper uses the term 'continuous' differs from the way other authors use the term 'continuous. It's implausible that every author who has claimed that consciousness is continuous intended to say that sufficiently small changes in some feature G map to arbitrarily small changes in some feature F. Perhaps the term 'continuous' ought to be interpreted more loosely, to mean either 'continuous or contiguous' or 'strikes one as continuous'.

I'll make three points in response.

First, there is textual evidence of authors using the term 'continuous' in the sense discussed in this paper. Alongside this paper's opening quote from Sellars [1963], consider Clark [1989: 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," Fara [2001: 924]'s explicit appeal to the standard epsilon-delta definition of 'continuity' used in mathematics, Dainton [2014: 130]'s characterization of temporal experiences as "phenomenal continua" that are infinitely divisible, or Prentner [2019: 29]'s claim that the "phenomenology of consciousness is such that it seems composed of an indefinite number of (phenomenal) parts."[26] Moreover, if philosophers were using 'continuity' in a non-standard way, then it would be

[26] See Rashbrook [2013] for further argument that various philosophers discussing consciousness and continuity have had in mind the sense of continuity deployed in this paper.

puzzling why the "apparent continuity" of consciousness has often been thought to be in tension with the apparent discreteness of its physical correlates.

Second, we should distinguish speaker meaning from 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 doesn't 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 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, for example, that we were to use 'continuous' to mean 'continuous or contiguous'. Then even many paradigmatically discrete structures (such as $\mathbb{Z}$) would count as continuous, questions about continuity would no longer have direct implications for questions about infinite divisibility, and there would be a confusing disparity between the sense of 'continuity' invoked in discussions of consciousness versus the sense of 'continuity' invoked by mathematicians, scientists, metaphysicians, and philosophers of science. Instead of proliferating senses of 'continuity', it's better to respect the established terminological standards.

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 our earlier discussions. First, 'continuous' doesn't merely mean 'smooth'—instead, it denotes the structural property that was defined earlier and that is 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. Speaking for myself, once I recognize these facts, I lose any feeling 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.

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.

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.

Young, Benjamin D, Keller, Andreas & Rosenthal, David (2014). Quality-space theory in olfaction. *Frontiers in Psychology* 5.