A TOPOLOGICAL SORITES*

The paradigmatic cases of the sorites paradox—heaps of sand and bald heads—are cases where the changes in question are small but discrete. Trading on the vagueness of 'heap' and 'bald', we remove discrete units, grain-by-grain and hair-by-hair, to produce the paradox. Moreover, in the case both of the heap and of baldness, there is a natural ordering, in terms of the number of grains of sand and the number of hairs. Let us call such versions of the sorites paradox *discrete* and *numerical*. Most of the discussion in the literature concerns such cases, and why not? They are very difficult to solve and have lead to extremely fruitful work in philosophy of logic. But it is important to bear in mind that these are not the only versions of the sorites.

For a start, there are continuous versions of the sorites. Consider a sorites argument using the predicate *tall* and starting with someone who is 200 cm high and progressing *continuously* down to someone 125 cm high. Such versions of the paradox are usually shoe-horned into the above discrete format by considering a particular series of discrete transitions—1 mm steps, for example. Be that as it may, the underlying space here is continuous (putting aside, for the purposes of argument, debates about the discreteness or continuity of space-time). It seems that we ought to be able to formulate the sorites paradox in terms of continuous transitions and not merely discretize continuous cases. Indeed, it would seem that the smaller the increments, the more compelling the sorites argument, so the continuous version might well be thought to be the most compelling of all; see section II.

Next consider non-numerical cases of the sorites.[1] Here we have familiar examples of family resemblance concepts such as religion and sports. Consider an example of transitions from Hinduism (with its ritualistic dress and behavior, belief in supernatural beings with special powers, the passion-play of good versus evil, and a catalogue

* We are indebted to James Chase, Lloyd Humberstone, Dominic Hyde, and audiences at the Melbourne Logic Group and the AAPNZ 2009 for useful conversations on the topic of this paper. Research for this paper was funded by an Australian Research Council Discovery Grant to Mark Colyvan and Dominic Hyde (grant number DP0666020).

[1] Otávio Bueno and Mark Colyvan, "Just What Is Vagueness?" *Ratio*, XXV (2012), forthcoming.

0022-362X/10/0706/311–325

of hymns and chants); through the passionate following of an Australian Rules Football team (with slightly less ritualistic dress and behavior, belief in players blessed with extraordinary, if not superhuman, powers, the various heroes and villains, and the various chants and team songs); to a casual children's game of ball in a backyard. It is plausible that a sorites argument can be constructed here, but there is no natural ordering as in the numerical versions described at the start.

To take a more scientifically significant example, consider the concept of *endangered species*. Here we construct a transition from an abundance of some species, with ample connected habitat and with no population decline, through to a single individual member of some species, with little or only fragmented habitat and suffering rapid decline.[2] Both vagueness and its sorites paradox seem to be active in such a case, even if there is no salient textbook construction of a sorites argument ready to hand.

It is sometimes said that such family-resemblance cases are cases of higher-dimensional sorites, whereby each dimension (for example, the degree of ritualistic behavior) is well ordered, but there is no overall total ordering of the transition states.[3] Even this much numerical ordering strikes us as implausible, but be that as it may, cases such as this do not naturally lend themselves to the standard presentations of the sorites, and all the more so for other apparently vague notions, such as jokes, wisdom, or love. At the very least, we need to do violence to the case in order to get it to fit the standard discrete, numerical schema.

It is easy to set aside such cases or to insist that they conform to the discrete, numerical schema via suitable adjustments. In this paper, at least, we are not denying that such moves can be made. We are, however, questioning the wisdom of such moves. After all, on the face of it we have several quite different versions of the sorites. It may be that a narrow focus on the discrete, numerical versions such as the heap of sand obscures what really drives the paradox. Such a narrow focus may even lead to overconfidence in a solution that deals only with

[2] Helen M. Regan, Mark Colyvan, and Mark A. Burgman, "A Proposal for Fuzzy International Union for the Conservation of Nature (IUCN) Categories and Criteria," *Biological Conservation*, XCII (2000): 101–08; Regan, Colyvan, and Burgman, "A Taxonomy and Treatment of Uncertainty for Ecology and Conservation Biology," *Ecological Applications*, XII (2002): 618–28.

[3] Arthur W. Burks, "Empiricism and Vagueness," this JOURNAL, XLIII, 18 (Aug. 29, 1946): 477–86; Dominic Hyde, *Vagueness, Logic, and Ontology* (Burlington, VT: Ashgate, 2008), p. 17.

the special cases under consideration. It is at least plausible that the underlying phenomenon has little to do with discreteness or numerical ordering. Clearly, we would like a unified solution to the sorites; in order to achieve this, we first need a characterization of the sorites paradox in its full generality. Only then can we be confident that we are in a position to see what makes it tick.[4]

In this paper, we propose to provide such a general characterization. We will start with the canonical presentation of the sorites, then outline a continuous version, and then move to an even more general topological version of the sorites. The topological formulation is interesting in its own right, but it also leads very naturally to a new, more general definition of the problematic concept of *vagueness.*

The logic is classical throughout the paper, and the theorems are text-book. Accordingly, $\vdash$ represents classical consequence, and $\supset$ is the material conditional. All the proofs are well known and are given in thumbnail or omitted altogether. The message is that, just as classical logic and number theory make unbelievable predictions in canonical forms of the sorites, so classical topology makes exactly the same kind of paradoxical predictions in the more general case.

I. DISCRETE SORITES

Hyde offers a useful classification of the sorites paradoxes.[5] The first and most familiar is a long series of (material) conditional statements, with a true first sentence (0 grains is not a heap), seemingly true subsequent sentences (either 250 grains is a heap, or 251 is not), and a false conclusion (10,000 grains is not a heap). The second form of the sorites paradox is a generalization, called the *inductive* form. Let Φ be a predicate and $n \in \mathbb{N}$.

$$\begin{array}{ll} & \Phi 0, \\ & \forall n(\Phi n \supset \Phi(n+1)). \\ \vdash & \forall n \Phi n. \end{array}$$

This is just the mathematical induction schema. When Φ is a vague predicate the premises seem true, and this leads to trouble because a vague predicate is tolerant to small changes but does not apply to every object.

[4] Colyvan, "Vagueness and Truth," in Heather Dyke, ed., *From Truth to Reality: New Essays in Logic and Metaphysics* (New York: Routledge, 2008), pp. 29–40.

[5] Hyde, "Sorites Paradox," *Stanford Encyclopedia of Philosophy*, ed. Edward Zalta (2008). URL: http://plato.stanford.edu/entries/sorites-paradox/.

Since in the case of vague Φ the conclusion is false, we must reject the induction step (also called the sorites premise), and we thus arrive at the *line-drawing* form:

$$\begin{array}{ll} & \Phi 0, \\ & \neg \forall n \Phi n. \\ \vdash & \exists n(\Phi n \wedge \neg \Phi(n+1)). \end{array}$$

This is a valid argument with true premises, but it is still taken to be a paradox because it seems implicit in the notion of vagueness that a vague predicate cannot be sensitive to very small changes. And yet the line-drawing form concludes that there is a single second, a single grain of sand or hair on the head, that leads from being Φ to not; some straw breaks the camel's back.

In what follows, we will look for arguments analogous to these inductive and line-drawing forms which do not trade on the discrete ordering of $\mathbb{N}$. Let us emphasize that we are not challenging the legitimacy of the canonical sorites paradox qua paradox. Rather, we are looking for more abstract renderings that reveal the canonical sorites to be special cases of a more sweeping phenomenon.

II. CONTINUOUS SORITES

James Chase has generalized the sorites to the continuous case. His argument draws out consequences of the distinctive axiom for continuity, which is as follows.[6]

> *Axiom 1 (Dedekind)* Let $A \cup B = \mathbb{R}$ be nonempty and disjoint sets, with $a < b$ for every $a \in A$ and $b \in B$. There is a unique $k \in \mathbb{R}$ such that $a \leq k \leq b$ for every $a \in A$ and $b \in B$.

From Dedekind's axiom we have the (equivalent) proposition that any set of reals bounded from above has a least upper bound. Then, by a standard series of lemmas, beginning with the Archimedean property (that for all $x \in \mathbb{R}$ there exists an $n \in \mathbb{N}$ such that $x < n$), it follows that the reals are dense, in the sense that if $x < y$ then there is a real z such that $x < z < y$. So much for how Dedekind's axiom constitutes the reals.

Consider a vague predicate Φ mapped onto a real-number interval [0, 1], exhaustively partitioned into two nonempty sets,

$$\begin{aligned} A &= \{x \in [0,1] : \Phi(x)\}, \\ B &= \{x \in [0,1] : \neg\Phi(x)\}, \end{aligned}$$

[6] We are already assuming the other usual definitions and field properties of the real numbers $\mathbb{R}$, as can be studied in any text, for example, Michael Spivak, *Calculus*, 3rd ed. (New York: Cambridge, 2006).

with $a < b$ for all $a \in A$, $b \in B$. We assume that $\Phi(0)$ and $\neg\Phi(1)$, and that if some number is not Φ, then no numbers after it are Φ either. Thus A is the left side of the interval and B is the right. The left set has a least upper bound; call it *supA*. Now, Φ is vague, and in discrete cases we are prepared to admit that objects differing by whole number amounts (a hair, a grain of sand) are too similar for one to be Φ but not the other. Here the objects in question are much closer together. Therefore, since points vanishingly close to *supA* are Φ, and Φ is vague, also $\Phi(supA)$. By a symmetrical argument, $\neg\Phi(inf\,B)$. Knowing this we have a paradox.

By the linear order on $\mathbb{R}$, one of the following must be true:

$$\begin{array}{ll} & supA < inf\,B \\ \text{or} & inf\,B < supA \\ \text{or} & supA = inf\,B. \end{array}$$

Since the reals are dense, we have the following contradiction. If *supA* and *inf B* are different numbers, then there is some z between them, $supA < z < inf\,B$ or $inf\,B < z < supA$. But then Φz and $\neg\Phi z$, since by definition anything less than *inf B* is Φ but anything greater than *supA* is not. On the other hand, if $supA = inf\,B$ then again $\Phi supA$ and $\neg\Phi supA$. This exhausts all the cases. Therefore there is a point both Φ and $\neg\Phi$, a contradiction.

The Dedekind axiom can be used to derive the intermediate value theorem, and here we just have a special case of this. A continuous path must cross over from A to B at some distinct point. The transition is problematic if the sets are supposed to be partitioned by a vague property.

The argument used by Chase can be represented in analogy to the discrete inductive form. A sequence $\mathcal{X} = \{x_0, x_1, \ldots\}$ is *Cauchy* iff for all real ε there is some $n \in \mathbb{N}$ such that $|\, x_i - x_j \,| < \varepsilon$ as long as i, $j > n$. Let $\mathcal{X}$ range over Cauchy sequences in the interval [0, 1]. The soritical argument now runs:

$$\begin{array}{ll} & \Phi 0, \\ & \forall \mathcal{X}(\forall x(x \in \mathcal{X} \supset \Phi x) \supset \Phi(sup\mathcal{X})), \\ \vdash & \Phi 1. \end{array}$$

The second premise is the sorites premise. This is not entirely analogous to the discrete case, since this is not a generally valid mathematical schema. Priest calls it the Leibniz continuity condition: whatever is going on arbitrarily close to some limiting point is also going on at the limiting point; *natura non facit saltus.* Were it generally valid, we could prove all sorts of nonsense.[7] In the case of a vague predicate, though, the condition seems ineluctable. Since it causes trouble, similarly to the discrete case, we negate the sorites premise and get a line-drawing form:

$$\begin{array}{ll} & \Phi 0, \\ & \neg\Phi 1 \\ \vdash & \exists \mathcal{X}(\forall x(x \in \mathcal{X} \supset \Phi x) \wedge \neg\Phi(sup\mathcal{X})), \end{array}$$

again where $\mathcal{X}$ is Cauchy.

What can we learn from this version of the paradox? For a start, we see how the sorites can be constructed so that it relies upon a property of the real line—the property of being *connected*. This property can be expressed with the notion of *metric adherence* (where topological adherence is defined in section IV below): A point x is *adherent* to a set X iff for any ε, no matter how small, the ε-sized interval around x includes points in X. With this in hand, we see, just as a consequence of Dedekind's axiom, that the interval [0, 1], and the reals in general, cannot be broken into two isolated parts:

Theorem 1 If $\mathbb{R}$ is partitioned into two nonempty, disjoint sets, some number is adherent to both sets.

Proof. Let $A \cup B = \mathbb{R}$, with $a \in A$ and $b \in B$. Without loss of generality suppose $a < b$. Then $inf\{x \in B : a < x\}$ is adherent to both A and B.□

A very common response to the discrete forms of the sorites paradox is to see a problem with exclusively and exhaustively separating objects into two categories, Φ and not. We now see that this problem is well expressed in terms of connecteness. Connectedness as exemplified in Theorem 1 is an emergent property of Dedekind's axiom, and the key in generalizing from the discrete to the continuous. We can now use this property to generalize again.

III. A TOPOLOGICAL SORITES

For millennia, geometers attempted to prove Euclid's parallel postulate. In the late eighteenth century came awareness that there are models of the first four Euclidean axioms that do not respect the parallel postulate. By the nineteenth century, in his landmark paper on the foundations of geometry, Riemann was able to diagnose *why* there are such models: The first four postulates, he saw, codify topological properties of the space, while the fifth is a specifically metric

[7] Graham Priest, *In Contradiction: A Study of the Transconsistent* (New York: Oxford, 2006), chapter 11.

property.[8] The lesson from Euclid is that there is a distinct science of space that does not deal in metric, quantitative notions, but only in qualitative notions like closeness.

It will be useful to describe the standard concepts of point-set topology.[9] The basic primitive (though intuitively familiar) notion is that of *open set*. Let X be a set. A *topology* is a collection of open subsets of X, closed under union and finite intersection, and including X and the empty set $\emptyset$. Let A be a member of the topology on X. A point x is interior to A, and A is a *neighborhood*[10] of x, iff there is an open set U where $x \in U \subseteq A$. A set A is open iff all its points are interior, that is, A is a neighborhood of all $x \in A$.

The interior of A is its largest open subset, the union of its open subsets, A°. The closure of A is its smallest closed superset, the intersection of closed supersets, A^-. The interior, the set, and the closure sit like this:

$$A^\circ \subseteq A \subseteq A^-.$$

A set A is open if A is contained in its interior, $A \subseteq A^\circ$, and A is closed if A contains its closure, $A^- \subseteq A$. Therefore a set is *both open and closed* if $A^\circ = A^-$.

Definition 1 A space X is *connected* iff the only sets in the topology of X that are both open and closed are X and $\emptyset$.

The following consequence could serve equally well as the definition of connectedness.

Theorem 2 A space is connected iff it cannot be partitioned into non-empty, disjoint, open sets.

At Theorem 1, for example, we saw that the reals $\mathbb{R}$ are connected.[11] We are now in a position to say why connected spaces are so useful for our present purposes.

[8] This and other insights are explored in Michael Spivak's *A Comprehensive Introduction to Differential Geometry*, vol. 2 (Berkeley: Publish or Perish, Inc., 1979), chapter 4.

[9] A standard reference is John L. Kelley's *General Topology* (New York: Springer-Verlag, 1955).

[10] The notion of a neighborhood is due to Hausdorff. He used the word *die Umgebung*, hence the common use of the symbol 'U'.

[11] There is a stronger notion, of a *path-connected space*, in which every two points a, $b \in A$ are connected by a path, a continuous function f: $[0, 1] \rightarrow A$ with $f(0) = a$ and $f(1) = b$. Every path-connected space is connected, but a connected space can still be impassible between two points (for example, the "topologist's sine wave"). See Lynn Arthur Steen and J. Arthur Seebach, Jr., *Counterexamples in Topology* (New York: Springer-Verlag, 1978). In multi-dimensional cases of vagueness, path connectedness seems to be the property that generates the paradox: We follow an arbitrary path through the space which takes us monotonically from one point in the space another. In pursuit of full generality, however, we will stick with the more general notion of connectedness.

Definition 2 A function f is *locally constant* iff for each $x \in X$ there is a neighborhood U_x such that the restriction of f to U_x is constant. A *globally constant* function always takes the same value, without restriction.

This is the key lemma.

Lemma 1 Let X be a connected space, Y a set, and f a function from X to Y. Suppose that f is locally constant. Then f is globally constant. A fortiori, if y is in the range of f, then $X = \{x : f(x) = y\}$.

Proof. Suppose f is not globally constant. Then there are objects x, $y \in X$ such that $f(x) \neq f(y)$. Then there is a $z \in X$ such that for any of its neighborhoods U_z, there are objects x, $y \in U_z$ and $f(x) \neq f(y)$.□

Heuristically, the set Y in Lemma 1 can be thought of as the pair $\{0, 1\}$, in which case the *characteristic function* σ of the set A is defined thus:

$$\sigma_A(x) = \begin{cases} 1 & \textit{if} \quad x \in A, \\ 0 & \textit{if} \quad x \notin A. \end{cases}$$

Consider a predicate Φ mapped onto a set A, and say that A is the extension of Φ. We have the analogous

$$\sigma_\Phi(x) = \begin{cases} 1 & \textit{if} \quad \Phi(x), \\ 0 & \textit{if} \quad \neg\Phi(x). \end{cases}$$

Since this set-up will lead to a paradox, there could be some objection to the language just employed: In using sets to represent predicates, we are assuming (with classical model theory) that predicates have extensions and these extensions are sets. We are eliding between predicates and sets, and it could be pointed out that, ever since Russell told Frege, we have known this is not always a harmless elision. There is, however, good reason to neglect such distinctions in this paper. The reason is that our goal is merely to formulate a problem using what looks like, in other cases, unproblematic language—to state, without jumping to solve, a paradox. We help ourselves to talk about extensions, only flagging that this language is not entirely innocent; see section v.

We use the notions of local constancy and characteristic function to propose a definition of vagueness.[12]

Definition 3 (*Vagueness*) A predicate is *vague* iff its characteristic function is locally constant but not globally constant.

The definition says that a vague predicate is tolerant of small changes but does run out somewhere. The principle of tolerance is found in the

[12] Thanks to Lloyd Humberstone for his contribution to formulating this definition.

standard literature on the sorites.[13] All the same, this definition is quite unlike any of the usual definitions in the literature.[14] But this is to be expected for two reasons.

First, it is well recognized that it is extremely difficult to provide a definition of vagueness that does not beg questions about its proper treatment.[15] For example, a common definition of vagueness in terms of permitting borderline cases, which in turn are defined as gaps, begs the question against glutty approaches.[16] While it is not the purpose of the present discussion to defend the above definition of vagueness against all charges of being question begging, its generality does suggest that it will do better on this front than some of the others—at least it does not presuppose that vagueness is a gappy rather than a glutty, or even nonclassical, phenomenon.[17] In any case, it is a very natural definition in the context of a more general conception of vagueness and is worth laying on the table.

This brings us to the second reason it is not surprising that this new definition is different from the standard ones: the standard definitions have a much narrower phenomenon as their targets—typically, vagueness associated with discrete, numerical sorites. Our aim is to provide a more general account of the sorites, and this must be accompanied with a more general definition of vagueness. Often, generalizations lead to new and more fecund definitions of the target concepts.[18] Still, we need to show that this definition does capture the intuitive notion. We do this

[13] Crispin Wright, "On the Coherence of Vague Predicates," *Synthese*, XXX, 3/4 (April–May 1975): 325–65.

[14] See for examples: Kit Fine, "Vagueness, Truth and Logic," *Synthese*, XXX, 3/4 (April–May 1975): 265–300; Hyde, "Sorites Paradox," *op. cit.*; Rosanna Keefe, *Theories of Vagueness* (New York: Cambridge, 2000); Roy Sorensen, *Vagueness and Contradiction* (New York: Oxford, 2001); Stewart Shapiro, *Vagueness in Context* (New York: Oxford, 2006); Nicholas J. J. Smith, "Vagueness as Closeness," *Australasian Journal of Philosophy*, LXXXIII (2005): 157–83; Timothy Williamson, *Vagueness* (New York: Routledge, 1994); Crispin Wright, "On the Characterisation of Borderline Cases," forthcoming.

[15] Shapiro, *op. cit.*; Bueno and Colyvan, *op. cit.*

[16] Hyde and Colyvan, "Paraconsistent Vagueness: Why Not?" *Australasian Journal of Logic*, VI (2008): 107–21; Zach Weber, "A Paraconsistent Model of Vagueness," *Mind*, to appear.

[17] Substituting 'continuous' for 'constant' in the definition, which would make no great difference in what follows, a fuzzy account can also be allowed for. Smith briefly entertains a definition of vagueness that does just this: A predicate is vague if its characteristic function is continuous (*Vagueness and Degrees of Truth* (New York: Oxford, 2008), p. 182). Smith works with degrees of truth (cf. his "A Plea for Things That Are Not Quite All There: Or, Is There a Problem about Vague Composition and Vague Existence?" this JOURNAL, CII, 8 (August 2005): 381–421); to make his arguments relevant to our restriction to truth values of only 0 or 1, we would say the characteristic function is constant. Smith points out some difficulties with a topological theory of vagueness, and this proposed definition in particular (*ibid.*). Since his objections are tied to continuity per se they do not impact the proposal here. We briefly take this up in section V.

[18] For example, to generalize the concept of a *straight line* to that of a *geodesic*, one does not use the seemingly obvious idea of the shortest distance between two points,

by way of a couple of formal results and show that it can be used to formulate different versions of topological sorites arguments.

Theorem 3 Let X be connected and $A \subseteq X$, with A the extension of a vague Φ. Then either $A = X$ or $A = \emptyset$.

Proof. The characteristic function σ_A of a vague predicate A is locally constant, by Definition 3. So, by Lemma 1, $\sigma_A : X \longrightarrow \{0, 1\}$ is globally constant.□

What does this mean? In general, for a connected space X, to prove that every $x \in X$ has some property Φ it suffices to show

BASE: Some $x \in X$ is Φ, and
INDUCTION: x is Φ iff all the points in a sufficiently small neighborhood of x are Φ.

These establish that all $x \in X$ are Φ because the induction asserts that the characteristic function for Φ is locally constant, and the base step asserts that the global value is not 0. This gives us a *topological inductive* version of the sorites:

Sorites Paradox, topological inductive version:
Let Φ be a vague predicate and X be a connected space.
Then if any member of X is Φ, every member of X is Φ.

Alternatively, with A as the extension of vague Φ,

$$\begin{array}{ll} & \emptyset \neq A \subseteq X, \\ & X \text{ is connected.} \\ \vdash & A = X. \end{array}$$

The 'induction step' is that X is connected, because connected spaces support the local-global property of Lemma 1, now built into the definition of vagueness. Faced with the discrete inductive sorites, we reject the induction step. Similarly, here we reexamine the situation from a line-drawing perspective.

Theorem 4 For any X such that $A \subset X$, that is, any space containing A other than A itself, if A represents a vague predicate then X is disconnected.

Proof. The characteristic function on A is locally constant, by Definition 3. Therefore it is globally constant, by Lemma 1. Now, $A = \{x : \sigma_A(x) = 1\}$. If X is connected, then $X = \{x: \sigma_A(x) = 1\}$, too; so if $X \neq A$, then X must be disconnected.□

---

but rather that of a curve whose acceleration is identically zero. The latter turns out to be more flexible and informative. See John M. Lee, *Riemannian Manifolds: An Introduction to Curvature* (New York: Springer, 1997), p. 47.

The result is a *topological line-drawing* sorites:

Sorites Paradox, topological line-drawing version:
Some things in X are Φ and some are not, for some vague Φ.
Then X is disconnected.

Alternatively, again with A as the extension of Φ,

$$\begin{array}{ll} & \emptyset \neq A \subseteq X, \\ & A \neq X. \\ \vdash & X \text{ is not connected.} \end{array}$$

Recall that, for the moment, we are only finding natural formulations of the paradox, in analogy to the discrete cases, without insinuating anything about how to interpret or resolve the problem. (It is interesting to notice how difficult it is only to state a problem, without trying to solve it.) The aim of the exercise is to show how topology represents vagueness, under classical assumptions—for example, that predicates can be represented with extensions. Under this assumption, classical topology predicts that the host-space of a vague predicate is not connected, because otherwise vague predicates would apply to everything.

IV. A CLOSER LOOK AT BOUNDARIES

Connectedness is a global property; it cannot be determined locally. But disconnectedness is a very local property. If σ is not globally constant, then it is not locally constant, either (by contraposition). So the disconnection becomes a local property, and the familiar counterintuitive aspect of line-drawing emerges. There is an $x \in A$, some particular point, in the neighborhood of which σ_A changes value.

Owing to the extreme locality of disconnection, we can study vagueness by studying the behavior of the characteristic function at the boundary of A. What would the boundary have to be like to support a sorites? Let us have a closer look at the boundary of a space.

In the following, the complement of A is $\mathcal{C}A = \{x : x \notin A\}$ and $X - A = X \cap \mathcal{C}A$.

Definition 4 A point $x \in X$ is adherent to A iff every neighborhood of x in X intersects A. The *boundary* of A, $\partial(A) := A^{-} \cap (X - A)^{-}$, is the set of all points adherent to both A and $X - A$.

A set shares its boundary with its complement, $\partial(A) = \partial(X - A)$, and a boundary is 'stable' in the sense that $\partial(\partial(A)) \subseteq \partial(A)$. Moreover, unpacking definitions,

$$\begin{aligned} \partial(A) &= A^{-} - A^{\circ}, \\ A^{\circ} &= A - \partial(A), \\ A^{-} &= \partial(A) \cup A^{\circ}. \end{aligned}$$

If $A \subseteq X$, then X is the pairwise disjoint union of $\partial(A)$, A^{-}, and $\mathcal{C}A^{-}$.

Now we have an alternative characterization of fundamental notions, summarized in a theorem:

Theorem 5 A is open iff $\partial(A) \subseteq X - A$, and A is closed iff $\partial(A) \subseteq A$.

A closed set includes all its adherent points; the closure of A is the set of all points adherent to A; so A is closed iff every point adherent to A is in A. It should now be clear that *set-theoretic extensions of vague predicates* are closed. For example, in the continuous case in section II, both the left set and the right set were closed by virtue of vagueness. Our further assumption that the sets be disjoint induced a contradiction.

We have now collected enough tools to make the last more precise, and to draw some morals. From the definition of adherence, it follows that a point k is adherent to A iff k is not interior to $X - A$, and k is interior to A iff k is not adherent to $X - A$. A fortiori, $\partial(A)$ is all the points interior to neither A nor $X - A$. In this terminology, A is open iff $A \cap \partial A$ is empty. Finally, A is both open and closed if $\partial(A) \subseteq A \cap (X - A)$. In sum,

Theorem 6 The following are equivalent:
The characteristic function on A is constant;
A is both open and closed, $A^{\circ} = A^{-} = A$;
$\partial(A) = \emptyset$.

This tells us that $x \in \partial(A)$ implies $x \in A$ and $x \notin A$, just as we saw in section II. Classically, this means that the boundary of A is empty, on pain of contradiction. From the definition of connectedness (Def. 1), the only sets with empty boundaries in a connected space X are X itself and $\emptyset$. Similarly, a connected space X is both open and closed, $X^{\circ} = X = X^{-}$; so if X is the set-theoretic extension of a vague predicate, then by the last theorem its boundary is overloaded. In section II, Chase derived a contradiction owing to the simple fact that if X is connected, $A \subseteq X$, and $B = X - A$, then A is open iff B is closed. The sorites premise made both sides of the interval both open and closed, overloading the boundary.

The sorites is a paradox. Classical theory encounters difficulties in the face of paradoxes, and such is the case here. The possibilities for a nonclassical treatment may be more flexible—for example, we might employ characteristic *relations* rather than functions, as is done in relational semantics.[19] On this approach, we could also organize things like $\vdash$ so that a glutty boundary is not disastrous. But to a fair extent these strategies await the development of more nonclassical mathematics. In the meantime, we have reached our destination, and step back to examine it.

[19] Priest, *An Introduction to Non-Classical Logic* (New York: Cambridge, 2008).

V. OPEN QUESTIONS

A number of questions arise. For now we consider an interconnected few concerning extensions and the representation of vagueness.

Let Φ be our vague predicate, and let A be the set-theoretic extension of Φ. One might wonder, following Smith,[20] as to whether there is a topology on the domain of the characteristic function of Φ. Simply stipulating that there is a topology on the domain, he argues, may be an "onerous" assumption not grounded in empirical experience. It is true that generating a topology is not always successful. For example, working with his own definition of vagueness in terms of a three-place 'closeness' relation, Smith is only able to generate the *discrete* topology,[21] which trivializes the exercise. And it is true that it is not always obvious that the space of a predicate like 'is an endangered species' has a natural topological structure. On the other hand, it does not seem to us very presumptuous to assume a nontrivial topology on the domain of the characteristic function of Φ. This depends, to an extent, on what the domain is. But whether or not we provide a recipe for generating a topology, it is entirely plausible that there is one in many interesting cases. Nontrivial topologies are not so hard to come by. In many cases, topologies can be inherited from Euclidean space. Or, returning to the definitions of boundary, for example, $\{\mathcal{C}(A \cup \partial(A)) : A \subseteq X\}$ is a topology.

More fundamentally, we can ask whether or not the extension of a predicate can always be interpreted with sets. If we are taking extensions to represent predicates, then this is asking whether or not vagueness can be represented by standard model theory. Perhaps, as has been the solution to Frege's naive comprehension woes, the answer here is to deny that every extension is a set. Perhaps extensions of vague predicates are not sets. However, this is a most unappealing thought, once we notice the large number of vague predicates in both our conversational language and more rigorous scientific language. The predicates in question here are not unusual like 'is not a member of itself' or 'is an ordinal', as in the Russell and Burali-Forti paradoxes, but banal predicates about sports and loud noises. We have been assuming since Descartes that the world is uniquely quantifiable by maps from the world to the real numbers; we have been assuming since Einstein that differential geometry and tensor calculus provide the tools to understand macroscopic space and time. These assumptions rest on basic representations of the world via set theory. If it turns out that vague properties cannot be treated in this way, then that would be very, very surprising.

[20] Smith, *Vagueness and Degrees of Truth*, p. 152.
[21] See Kelley, *op. cit.*, p. 37.

With set-theoretic extensions looked at in this way, then, there is a further open question. Is the topological characterization of sorites paradoxical at all? Candidates for multi-dimensional vagueness are commonplace but are more abstract than predicates like redness or baldness. The space of jokes, wisdom, or love may well be disconnected, being abstract fragments of logical space to begin with. Was there reason to have suspected otherwise? Rather than a paradox, perhaps in these cases we have learned something structural—that, in a sense, the transition from religion to sport involves traversing a disconnection.

Nevertheless, it is also clear that there will be at least some cases where multi-dimensional sorites traverse abstract connected spaces. Some of these spaces (for example, Minkowski space) will be of scientific interest, and there will be sorites arguments in these spaces which will be hard to represent as examples of canonical discrete, numerical sorites. The space of threatened species, for example, has several degrees of freedom: species numbers, area and quality of habitat, rate of decline, and others. While each of these is arguably numerical, it is not at all clear how to combine them in a meaningful fashion. Any proposed metric in this space, it seems, will be problematic. Yet the space in question is very plausibly connected. Importantly, whether or not this and other spaces are unexpectedly disconnected, we have made progress, because the characterization we have given here will help in formalizing the alleged sorites arguments in all such cases. Our topological characterization leaves open the question of whether the space really is connected; it leaves the question sharpened.

Perhaps the concern can be put differently—as a concern about maintaining a neutral dialectic. We have proven that whenever we have a sorites series, the underlying space must be disconnected. This should be music to the ears of epistemicists about vagueness, for they take the lesson learned from the sorites paradox to be that all such spaces are (surprisingly) disconnected. It seems we have just vindicated the epistemic approach and thus trivialized the debate by ruling out other serious contenders such as supervaluational and fuzzy approaches. Any definition that begs all the important questions is no definition at all.

This, however, is to misunderstand where things stand. As we just indicated, some of the spaces in question are connected—some are provably connected (for example, Minkowski space and $\mathbb{R}^n$), while others have strong cases to be made for their connectedness (for example, the space of threatened species). This is the heart of the sorites and why it is a paradox. Just as vague predicates in the company of classical logic lead to genuine paradox, so too does classical topology with vague set descriptions. All along, our purpose has been to provide a generalization of the sorites; we have not been trying to solve the paradox. Of course if one rigidly sticks with classical topology, then epistemicism appears to be the only viable approach; this is no different from the canonical sorites, where if one refuses to depart from classical logic, epistemicism appears to be the only viable option. But all the standard moves are here to be made; no proposed solution is ruled out. For example, a very natural move to make here is to entertain a nonclassical theory of topology, one underwritten by a nonclassical set theory—fuzzy, intuitionisitic, or paraconsistent—or through new directions in topological computation.[22] Our discussion of a generalized sorites via an appeal to classical point-set topology was not intended to suggest that classical topology should be used in solving the problem.

We do not expect that we have said enough here to convince everyone that the sorites is essentially topological. An important concern is whether the generalization is too general. A topological characterization of vagueness is a generalization on the canonical discrete, numerical sorites in the same way that topology itself is a generalization on certain properties of the real numbers. One important consequence of the topological approach is that, since a topological sorites does not require assumptions about order, any proposed solutions to sorites that trade on the details of order are not going to be general solutions. Conversely, this could mean that the generalization has gone too far and has lost grip on the essence of the sorites. If so, topology is the wrong tool for the job, and what we are calling a sorites is, in fact, not.

If, on the other hand, the argument form we display with only topological tools is still a recognizable sorites, then the 'lost' information is inessential. With some doubts now aired, we do think that a topological sorites is recognizably a generalization of the canonical sorites and that the topological characterization captures the essential ingredients—namely, connectedness and local and global constancy. It is not hard to see that the core notion of local constancy is a generalization of the principle of tolerance, and that the topological sorites is a generalization in the sense that the canonical cases can be recovered as special cases. If we are right about this, then progress has been made in exposing what is invariant about vagueness in a variety of cases, and we are closer to understanding a very resilient puzzle.

ZACH WEBER

University of Sydney and University of Melbourne

MARK COLYVAN

University of Sydney

[22] Viggo Stoltenberg-Hansen and John V. Tucker, "Computability on Topological Spaces via Domain Representations," in *New Computational Paradigms: Changing Conceptions of What Is Computable* (New York: Springer, 2008.)