

Vagueness: A Conceptual Spaces Approach

Igor Douven · Lieven Decock · Richard Dietz ·
Paul Égré

Received: 14 April 2011 / Accepted: 18 October 2011 / Published online: 25 November 2011
© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract The conceptual spaces approach has recently emerged as a novel account of concepts. Its guiding idea is that concepts can be represented geometrically, by means of metrical spaces. While it is generally recognized that many of our concepts are vague, the question of how to model vagueness in the conceptual spaces approach has not been addressed so far, even though the answer is far from straightforward. The present paper aims to fill this lacuna.

Keywords Vagueness · Conceptual spaces · Voronoi diagrams · Gardenfors

1 Introduction

The conceptual spaces approach is enjoying growing popularity in the cognitive psychology and cognitive science communities. In this approach, concepts

I. Douven (✉)
Faculty of Philosophy, University of Groningen, Oude Boteringestraat 52, 9712 Groningen,
The Netherlands
e-mail: i.e.j.douven@rug.nl

L. Decock
Faculty of Philosophy, VU University Amsterdam, Amsterdam, The Netherlands
e-mail: lb.decock@ph.vu.nl

R. Dietz
Department of Philosophy, University of Tokyo, Tokyo, Japan
e-mail: rdietz@l.u-tokyo.ac.jp

P. Égré
Institut Jean–Nicod (CNRS), Paris, France
e-mail: paulegre@gmail.com

are represented geometrically, by means of so-called metrical spaces, an approach which allows for a formally precise treatment of how the human mind conceptualizes the world. The approach is relatively young and still developing, and there are some foundational issues that are yet to be addressed. One key issue concerns the fact that, while many of our concepts are vague, it is not evident how to model vagueness in the framework of the conceptual spaces approach.

The vagueness of concepts manifests itself most prominently in the occurrence of borderline cases, cases that we do not quite regard as falling under a given concept nor quite regard as *not* falling under that concept. Any theory of concepts will, on pain of material inadequacy, have to account for the existence of borderline cases, and it will have to do so in a way that explains our typical responses to such cases.

Theorists employing conceptual spaces in their work have to this date ignored the issue of vagueness. Yet their work does suggest an answer to the question of how there can be borderline cases, and, importantly, one that might explain why we respond to such cases the way we do. As we shall see, however, this answer cannot quite be correct, for it fails to do justice to some incontrovertible data concerning vagueness. To arrive at a satisfactory account of vagueness within the conceptual spaces approach, one that does do justice to the said data, this paper proposes two extensions of the formal framework of the approach.

Naturally, the account of vagueness to be offered can lay no greater claim to generality than can the conceptual spaces approach. And it must be admitted that, while this approach works very well for perceptual concepts, it does not obviously apply to more complex and abstract concepts. It may be a matter of time before the theory can satisfactorily deal with the latter kinds of concepts as well. Or, given that it is not a priori that all conceptualization must eventually be accounted for in the same terms, non-perceptual concepts may require some altogether different approach. In the latter case, non-perceptual concepts will require their own account of vagueness as well. We do not regard this as a downside of our proposal. After all, it is not a priori that all cases of vagueness must, in the end, be accounted for in the same way. It may even be prudent to assume that this is not so. In any event, we regard as the scope of our proposal concerning vagueness whatever the scope of the conceptual spaces approach will turn out to be.

As a further disclaimer, we note that our aim is to give an account of *borderline* vagueness within the conceptual spaces approach. We do not aim to give a full-fledged account of any of the other properties that may typically come with vagueness. Most notably, the treatment of the issue of soriticality in our framework must await another occasion.

We start by outlining the conceptual spaces approach and relating it to the issue of vagueness, particularly the question concerning borderline cases (Section 2). We then present two extensions of the conceptual spaces framework (Sections 3 and 4), and we argue that thereby we obtain a framework that enables us to deal with vague properties and concepts more satisfactorily than

the original one (Section 5). Finally, we point at some empirical consequences of the resulting account of vagueness (Section 6).

2 Conceptual Spaces and Vagueness

Many researchers in cognitive psychology and cognitive science have contributed to the development of the conceptual spaces approach, but the following exposition is mostly based on Gärdenfors's [6] exemplarily clear presentation of the approach and its main results. As intimated, the basic idea of this approach is that concepts can be represented geometrically, specifically as regions in certain well-defined spaces. These spaces are one-dimensional or multidimensional structures—sets of points—with a distance function (a metric) defined on them. Objects are mapped onto points in these spaces, and the dimensions of these spaces correspond to “qualities” objects may have. Formally, where S is a space, $\delta_S: S \times S \rightarrow \mathbb{R}_0^+$ is a distance function on S iff, for all $a, b, c \in S$: (i) $\delta_S(a, b) \geq 0$, and $\delta_S(a, b) = 0$ iff $a = b$; (ii) $\delta_S(a, b) = \delta_S(b, a)$; and (iii) $\delta_S(a, b) + \delta_S(b, c) \geq \delta_S(a, c)$. A common distance function is the Euclidean metric. Following Gärdenfors, we will mostly assume a generalized version of this metric, which, for an m -dimensional space S , defines the distance between points $p = \langle x_1, \dots, x_m \rangle$ and $p' = \langle x'_1, \dots, x'_m \rangle$ as

$$\delta_S(p, p') := \sqrt{\sum_{i \leq m} w_i \cdot (x_i - x'_i)^2}. \quad (1)$$

With $w_i = 1$ for all i , Eq. 1 gives the standard Euclidean metric. By varying the weights relative to one another, one can give more “salience” to some dimensions than to others.¹ Other metrics—most notably, the city-block metric—have also been discussed in the conceptual spaces literature, but for our later purposes the above is the most relevant one.

As defined on a conceptual space, a distance function measures the similarity between objects. More exactly, the similarity between two objects is assumed to be a monotonically decreasing function of their distance in the space: the greater the distance, the less similar the objects are. It is worth noting that the notion of similarity at stake here is not similarity *tout court*, but similarity in a given respect, where the respect corresponds to the given space.

One of the simplest examples of a conceptual space is a three-dimensional Euclidean space with a Euclidean metric defined on it which serves to represent proximity relations between objects in the world. Other well-known examples are temporal space, auditory space, olfactory and taste spaces, and spaces corresponding to various physical parameters such as density and hardness. In

¹Whenever we use “ δ_S ” in the following, we refer to *instances* of Eq. 1, with values for the w_i assumed to be in place.

his writings, Gärdenfors also discusses multidimensional shape spaces [6: 94–98] and action spaces [7], but these are among the more controversial examples hinted at in the introduction. A further well-known example, which will be invoked repeatedly in the following, is color space. This is standardly taken to be three-dimensional, and to be equipped with a Euclidean distance function. One dimension represents hue; this dimension can be pictured as a color circle with yellow, green, blue, violet, red, and orange (neighboring yellow again) lying in that order on the circle, with one color gradually merging into the next. The second dimension is saturation, or chromaticity, which represents the intensity of the color. And the third dimension is brightness, or luminosity, which ranges from white to black, through all shades of gray.

As already mentioned, in this approach concepts are identified with regions of conceptual spaces. For instance, the concept *red* corresponds to a particular region of the aforementioned color space. Of course, *red* is not just any concept; it is a *natural* concept. While, technically speaking, a region of a space is any subset of points of that space, not all regions in this sense correspond to natural concepts. Indeed, the regions of a space that do correspond to such concepts will typically be sparse. Gärdenfors does not give necessary and sufficient conditions for when a region represents a natural concept.² However, he does at least give a necessary condition for naturalness, to wit, convexity, meaning that for any pair of points in the region, the line segment connecting the points lies, in its entirety, in the region as well.³

In order to arrive at a plausible account of categorization, that is, of how people carve up conceptual spaces into regions corresponding to natural concepts, Gärdenfors [6: 87–91] invokes *prototype theory* together with the mathematical technique of *Voronoi diagrams*.

Prototype theory capitalizes on the observation that among the instances of a concept, some are more representative of that concept than others. The one that is most representative is the prototype of the concept. While the original advocates of this theory had clearly universalist aspirations in that the classification of certain instances as prototypes was believed to transcend cultural boundaries (e.g., Rosch [15]), for the purposes of the conceptual spaces approach one can stay entirely noncommittal on this issue. For instance,

²Gärdenfors [6: 67] does say that natural concepts are “those that are natural for the *purposes* of problem-solving, planning, memorizing, communicating, and so forth” (original emphasis). This is compatible with a pragmatic conception of natural kinds, and does not commit him to the metaphysically more loaded conception of natural kinds as ultimate building blocks of the universe.

³A related topological property of regions that will be referred to below is connectedness. A region r is said to be *connected* iff, for all points p_i and p_j in r , there is a curve connecting p_i and p_j such that the curve lies entirely in r (or, to be more precise, such that all points lying on the curve lie in r). Clearly, a region that is convex is connected, but not necessarily the other way round.

for the purposes of the conceptual spaces approach, a prototype may just be a *typical case*, where this is understood as a case that competent members of a given language community in normal circumstances classify, or would classify, unhesitatingly as being typical. Whether such a classification is determined by something hard-wired in people's heads, or whether it merely reflects cultural bias, or whether it depends both on physiology and cultural background, is immaterial for the combined use of prototypes and Voronoi diagrams that Gärdenfors proposes.⁴

Voronoi diagrams have been applied in diverse fields, including physics, astronomy, geology, geography, meteorology, and recently also the cognitive sciences. A Voronoi diagram of a given space is a division of that space into cells such that each cell has a center and consists of all and only those points that lie no closer to the center of any other cell than to its own center. More formally put, given an m -dimensional space S and a set $P = \{p_1, \dots, p_n\} \subset \mathbb{R}^m$ of pairwise distinct points in S , we call the region

$$v(p_i) := \{p \mid \delta_S(p, p_i) \leq \delta_S(p, p_j), \text{ for all } j \in \{1, \dots, n\} \text{ with } j \neq i\}$$

the *Voronoi polygon/polyhedron associated with p_i* . Then the *Voronoi diagram generated by P* is the set

$$V(P) := \{v(p_i) \mid p_i \in P\}.$$


The points $p_i \in P$ are called the *generator points* of $V(P)$, and P is called the *generator set*. The vertices of the polygons/polyhedrons $v(p_i)$ are called *Voronoi points* of $V(P)$, the edges/faces of the polygons/polyhedrons the *Voronoi edges/faces*. Generator points that share a Voronoi edge/face are said to be *adjacent*.⁵ An example of a Voronoi diagram of a bounded two-dimensional space is given in Fig. 1.

According to Gärdenfors, “the Voronoi tessellation provides a constructive answer to how a similarity measure together with a set of prototypes determine a set of natural properties” [6: 88], emphasis omitted). That his criterion for naturalness—convexity—is satisfied follows from

⁴In the introduction, we already noted that the conceptual spaces approach is presently limited in that not all properties and concepts are handled equally well by it. We cannot exclude the possibility that if the idea of prototypes is taken aboard, this places a further limitation on the approach. After all, it is far from obvious that for every concept the idea of a prototypical instance makes sense. For instance, what is a prototypical instance of the concept *tall*, or of *poor*, or of *beautiful*? Even when relativized, such gradable adjectives may pose problems, given that it would be hard to say what a prototypical instance of a tall (or poor, or beautiful) person is. Perhaps instead of prototypes, the conceptual spaces account (as well as the approach to vagueness to be developed) can make do with *exemplars*, that is, clear examples which play a central role in learning the use of a predicate (this is rough; see Gärdenfors [6: 123 f] for a precise characterization). Whether or not the resulting account (or rather accounts) would be empirically adequate is not something we will try to determine here.

⁵See Okabe et al. [11, Ch. 2] for these definitions.

Fig. 1 A two-dimensional Voronoi diagram



Fact 2.1 All Voronoi polygons/polyhedrons of Voronoi diagrams based on a Euclidean metric are convex.⁶

But what reasons do we have for thinking that the designated “constructive answer” accords with how people actually categorize reality? Some support for Gärdenfors’s proposal may come from theoretical considerations, in particular considerations of cognitive economy:

A Voronoi tessellation based on a set of prototypes is a simple way of classifying a continuous space of stimuli. The partitioning results in a *discretization* of the space. The prime cognitive effect is that the discretization speeds up *learning*. The reason for this is that remembering the finite prototypes, which is sufficient to *compute* the tessellation once the metric is given, puts considerably less burden on *memory* than remembering the categorization of each single point in the space. In other words, a Voronoi tessellation is a cognitively *economical* way of representing information about concepts. Furthermore, having a space partitioned into a finite number of classes means that it is possible to give *names* to the classes. (Gärdenfors [6: 89])

Still, it might be objected that nothing much can be said about how economical the said way of representing information is as long as we do not know how much effort it takes to mentally generate a Voronoi diagram from a set of prototypical points. It would also be an exaggeration to say that the empirical studies Gärdenfors musters in support of his proposal are conclusive or even nearly so. Nevertheless, we adopt as our working hypothesis that, to a first approximation at least, the conceptual spaces approach, coupled with the ideas of prototypes and Voronoi diagrams as a way of generating categorizations, captures an important part of the truth about human cognition.

To accommodate the fact that concepts can be vague, one must be able to make sense of the notion of borderline case, we said in the introduction. The above theory would seem to have no difficulty with that. Just have another look at Fig. 1. In the present approach, concepts clearly have borderlines. A

⁶Gärdenfors [6: 88]. See Okabe et al. [11: 58] for a proof.

borderline case could thus be said to be a case that falls on the borderline between different natural concepts, a case for which there is no *unique* prototype to which it is closest, a case, again differently put, that is at equal distance from adjacent prototypes. The answer would seem to be of exactly the type we are after, in that it provides an explanation of why we respond to borderline cases the way we do: we do not quite regard a particular color shade as falling under, say, the concept *red*, nor quite regard it as *not* falling under that concept, because it is as far removed from the prototypical shade of red as it is from the prototypical shade of, for example, orange.

While a characterization of borderline cases in terms of their positions relative to prototypical instances in conceptual space is along the right lines, it is not complete enough to be empirically adequate. Specifically, the problem is that, on this characterization, all X/Y borderline cases (that is, borderline cases between properties X and Y) are at an equal distance from X as they are from Y . So, given the intended interpretation of the metric, all X/Y borderline cases should be as much X as they are Y . Surely, however, there exist, for instance, red/orange borderline cases which are more reddish than “orangish” as well as ones that are more orangish than reddish. The problem—which, for reasons that will become apparent, we term the “thickness problem”—is that, in the present picture, such cases cannot occur. In the following, we will propose a solution to this problem that fits broadly with Gärdenfors’s approach. It adds to the approach a more realistic conception of prototypes and a sort of Voronoi diagram that is slightly different from the standard one invoked by Gärdenfors.

3 Prototypical Areas

It was probably only as a simplification that some have supposed every concept to have a *unique* prototype. For most concepts, the thought of there being exactly one most representative instance appears to make little sense. This is particularly obvious if the notion of prototype is understood in a metaphysically lightweight, non-universalist sense, according to which the prototypical instances of a concept are its typical instances. Few people will consider exactly one shade of red to be typically red. Rather, we count many—slightly different—shades of red as being typical instances of red, though we count many more as being nontypical instances. The same is true for the other colors and for a great many other concepts. It is worth noting that this is not just intuitively plausible. At least insofar as colors are concerned, there is also ample empirical evidence for the non-uniqueness of typical cases. In their [1], Berlin and Kay report by now famous research in which they asked subjects to mark, on a board with all the Munsell color chips ordered on it, “the best, most typical examples of x ” (p. 7). They found that, in response to this question, the subjects frequently marked more than one chip (p. 10).

As Gärdenfors [6: 139] shows, supplanting the assumption of prototypical points by that of prototypical areas hardly necessitates a large-scale revision

of the conceptual spaces approach, for one can as easily carve up conceptual spaces on the one assumption as one can do on the other. Specifically, one can also have prototypical areas generate a Voronoi diagram. To show how to do this, Gärdenfors assumes that prototypical areas are circles.⁷ Then for a set of prototypical circles with centers $\{p_1 = \langle p_{1_1}, \dots, p_{1_m} \rangle, \dots, p_n = \langle p_{1_n}, \dots, p_{m_n} \rangle\}$ and radii $\{c_{p_1}, \dots, c_{p_n}\}$, a “generalized Voronoi diagram” is defined in the same way as an ordinary Voronoi diagram is, except that now for each circle with center p_i the associated Voronoi polygon/polyhedron is defined to be the set

$$\left\{ p \mid \sqrt{\sum_{k \leq m} w_k \cdot (x_k - p_{k_i})^2} - c_{p_i} \leq \sqrt{\sum_{k \leq m} w_k \cdot (x_k - p_{k_j})^2} - c_{p_j}, \right. \\ \left. \text{for all } j \in \{1, \dots, n\} \text{ with } j \neq i \right\},$$

where $p = \langle x_1, \dots, x_m \rangle$. Gärdenfors (*ibid.*) notes that, given the metric assumed here, all points lying on a prototypical circle come out as having distance zero from the center of the area, and points lying within the circle—including the center itself—as having imaginary distances from the center. These mathematical facts are not easy to interpret if distances are supposed to reflect degrees of similarity.

But even apart from this oddity, the above way of combining the non-uniqueness of prototypes with Voronoi diagrams is not the only way, and not the best, at least not for the purposes of representing vagueness. In particular, appealing to generalized Voronoi diagrams is not going to help with solving the thickness problem, in as much as the boundaries of such diagrams still consist of lines of points all of which are equidistant (given the above metric) from the relevant prototypical areas.

4 Collated Voronoi Diagrams

Though the main textbook on Voronoi diagrams (Okabe et al. [11]) lists many more variants of such diagrams than the just-described generalized version, the type of Voronoi diagram we want to make use of is, to our knowledge, new. Instead of calculating a single diagram on the basis of a set of prototypical areas—as the generalized Voronoi diagram does—we consider a construction built out of a set of ordinary Voronoi diagrams. The basic idea is that each choice function that picks from all prototypical areas exactly one point determines a single diagram, and that these diagrams can be collated, or projected onto each other. We call the result a *collated Voronoi diagram* (or *collated diagram*, for short).

⁷This assumption is not necessary; see Okabe et al. [11: 186–189], who show how to generate Voronoi diagrams on the basis of arbitrary areas. However, Gärdenfors’s assumption may have been motivated by considerations of cognitive economy; see below.

Before stating the definition, we first introduce the notion of a *restricted* Voronoi polygon/polyhedron:

Definition 4.1 Given an m -dimensional space S and a set $P = \{p_1, \dots, p_n\} \subset \mathbb{R}^m$ of pairwise distinct points, the *restricted Voronoi polygon/polyhedron associated with p_i* is the region

$$\underline{v}(p_i) := \{p \mid \delta_S(p, p_i) < \delta_S(p, p_j), \text{ with } j: 1 \leq j \leq n\}.$$

Thus, less formally, a restricted Voronoi polygon/polyhedron associated with a given generator point is the (unrestricted) Voronoi polygon/polyhedron associated with that point minus its edges/faces.

Let $R = \{r_1, \dots, r_n\}$ be a set of pairwise disjoint regions of a space S . Then

$$\Pi(R) := \prod_{i=1}^n r_i = \{\langle p_1, \dots, p_n \rangle \mid p_i \in r_i\}$$

is the set of all ordered sequences $\langle p_1, \dots, p_n \rangle$ such that $p_i \in r_i \in R$ for $1 \leq i \leq n$; that is, each sequence in $\Pi(R)$ contains exactly one point out of each region in R . We note that, where $|r|$ denotes the number of points in r and $R = \{r_1, \dots, r_n\}$, the cardinality of $\Pi(R)$, $|\Pi(R)|$, equals $\prod_{i \leq n} |r_i|$.

Given a set R of disjoint regions of a space S , we can consider the set of all Voronoi diagrams of S generated by elements of $\Pi(R)$:

$$\mathcal{V}(R) := \{V(P) \mid P \in \Pi(R)\}.$$

It is easy to see that not every Voronoi diagram corresponds to a unique set of generator points. Therefore, $|\mathcal{V}(R)| \leq |\Pi(R)|$.

Let us introduce the following abbreviations for the set of all Voronoi polygons/polyhedrons $v(p)$ associated with the various points p in region $r_i \in R$, and similarly for restricted polygons/polyhedrons:

$$\{v(p)\}_{r_i \in R} := \{v(p) \mid p \in r_i \wedge v(p) \in V(P) \in \mathcal{V}(R)\};$$

$$\{\underline{v}(p)\}_{r_i \in R} := \{\underline{v}(p) \mid p \in r_i \wedge v(p) \in V(P) \in \mathcal{V}(R)\}.$$

Given a set $\mathcal{V}(R)$, from its elements we can construct a collated Voronoi diagram, as follows:

Definition 4.2 Let $R = \{r_1, \dots, r_n\}$ be a set of pairwise disjoint regions of a space S . Then the region

$$\underline{u}(r_i) := \bigcap \{\underline{v}(p)\}_{r_i \in R}$$

is the *collated polygon/polyhedron associated with r_i* , and the set

$$\underline{U}(R) := \{\underline{u}(r_i) \mid 1 \leq i \leq n\}$$

is the *collated diagram generated by R* . The r_i are called the *generator regions* of $\underline{U}(R)$. The set $S \setminus \bigcup \underline{U}(R)$ is called the *boundary region* of $\underline{U}(R)$.

So, where an ordinary Voronoi polygon/polyhedron associated with a generator point consists of precisely those points that lie at least as close to that generator point as they lie to any other generator point, a collated polygon/polyhedron associated with a given generator region consists of precisely those points that are closer to *every* point in that region than they are to *any* point in *any* other generator region; see Fig. 2 for an illustration. Note that, unlike an ordinary Voronoi diagram, a collated diagram does not in general exhaust a space, though a collated diagram together with its boundary region does. Also, we stick to the terminology of polygons/polyhedrons, even though, depending on the shapes of the generator regions, the elements of a collated diagram need not literally be polygons/polyhedrons but may have different shapes as well (e.g., they may be circles/spheres).

It is natural to define a notion dual to that of a collated polygon/polyhedron, which we may call the *expanded* polygon/polyhedron associated with a region:

Definition 4.3 Let $R = \{r_1, \dots, r_n\}$ be a set of pairwise disjoint regions of a space S . Then the region

$$\bar{u}(r_i) := \bigcup_{r_i \in R} \{v(p)\}$$


is the *expanded polygon/polyhedron associated with r_i* .

Consequently, while the collated polygon/polyhedron is the intersection of the associated polygons/polyhedrons excluding their edges/faces, the expanded polygon/polyhedron corresponds to the union of the associated Voronoi polygons/polyhedrons including their edges/faces.

With these definitions in place, we can think of at least two plausible ways to define the notion of *boundary region associated with a collated polygon/polyhedron* $\underline{u}(r_i) \in \underline{U}(R)$.

One possibility is to view this region as the set of points p such that p lies on a Voronoi edge/face, or is a Voronoi point, of some $v(p') \in V(P) \in \mathcal{V}(R)$ with $p' \in r_i$. On this definition, to be a boundary point of a collated polygon/polyhedron means to lie on an edge/face of any of the polygons/polyhedrons in the collation. We shall use the notation $b_1(\underline{u}(r_i))$ for boundary regions in this first sense.

Fig. 2 A collated diagram with a gappy boundary region; the collated polygon associated with generator region $r = \{a, b\}$ is the smallest area containing r that is enclosed by boundary lines



Another possibility is to let the boundary region be the difference between the expanded polygon/polyhedron and the collated polygon/polyhedron associated with r_i , that is, $\overline{u}(r_i) \setminus \underline{u}(r_i)$. On this definition, the boundary region associated with a collated polygon/polyhedron $\underline{u}(r_i)$ consists of the points that lie in the union of all the inclusive Voronoi polygons/polyhedrons associated with the various points in r_i but not in the intersection of the corresponding restricted polygons/polyhedrons. We shall use the notation $b_2(\underline{u}(r_i))$ for boundary regions in this second sense.

The two definitions agree on the idea that if a point is in the boundary region of $\underline{u}(r_i)$, then it falls entirely out of $\underline{u}(r_i)$. The first definition, however, is more constraining than the second. Indeed, a point may belong to $b_2(\underline{u}(r_i))$ without belonging to $b_1(\underline{u}(r_i))$ if it does not lie on a Voronoi edge/face. On the other hand, if a point belongs to $b_1(\underline{u}(r_i))$, then it lies on a Voronoi edge/face of some polygon/polyhedron associated with a point in r_i , and so it belongs to the union of the polygons/polyhedrons, and therefore to $b_2(\underline{u}(r_i))$. Though boundary regions in the first sense are not in general boundary regions in the second sense, it will be seen that the two notions coincide provided further topological restrictions are assumed to hold on the regions r_i , namely, regarding the distribution of prototypical instances.

In the next section, we will want collated diagrams to serve much the same goal in the conceptual spaces approach that Gärdenfors had in mind for ordinary Voronoi diagrams. To that end, we give two important facts about collated diagrams. The first is this:

Fact 4.1 All collated polygons/polyhedrons are convex.

To see this, recall Fact 2.1, and note that by removing the edges/faces from a Voronoi polygon/polyhedron we again obtain a convex set.⁸ Fact 4.1 then follows via the observation that the intersection of a collection of convex sets is itself convex.⁹

Stating the second fact takes a bit more preparation. Note that if, for some $P \in \Pi(R)$, a point p in S lies on a boundary line of $V(P)$, then p is an element of $S \setminus \bigcup \underline{U}(R)$. To be precise, given a set $R = \{r_1, \dots, r_n\}$ of pairwise disjoint

⁸*Proof* Suppose, towards a contradiction, that, for some p , $v(p)$ is convex but $\underline{v}(p)$ is not. Then there must be points $a, b \in \underline{v}(p)$ such that \overline{ab} lies in its entirety in $v(p)$ but does not lie in its entirety in $\underline{v}(p)$. Then \overline{ab} must intersect some edge(s)/face(s) of $v(p)$. One easily verifies that there are only two types of line segments that intersect some edge(s)/face(s) of $v(p)$, to wit, first, line segments that connect points lying outside $\underline{v}(p)$, and second, line segments that connect a point in $\underline{v}(p)$ with one outside $\underline{v}(p)$. But, clearly, if \overline{ab} were of the first type, neither a nor b would lie inside $\underline{v}(p)$; if it were of the second, a and b would not both lie inside $\underline{v}(p)$. Either way it follows that there cannot be points $a, b \in \underline{v}(p)$ such that \overline{ab} lies in its entirety in $v(p)$ but not in its entirety in $\underline{v}(p)$. Hence, given that $v(p)$ is convex, $\underline{v}(p)$ must be convex, too. \square

⁹*Proof* Any pair of points in the intersection will be in each of the intersecting sets. By the convexity of these sets, they each also contain the whole line segment connecting the points. Because this line segment is in each of the sets, it is in their intersection. Hence, the intersection is convex. \square

regions in a space S , the set $B(R, S)$ containing the points of S that lie on a boundary line of $V(P) \in \mathcal{V}(R)$, for some $P \in \Pi(R)$, is equal to the set of points p such that there is a $P \in \Pi(R)$ for which it holds that, for at least two points $p_i, p_j \in P$, $\delta_S(p, p_i) = \delta_S(p, p_j)$, and for no $p_k \in P$, $\delta_S(p, p_k) < \delta_S(p, p_i)$.

Does it necessarily hold that $B(R, S) = S \setminus \bigcup \underline{U}(R)$? That is to say, can we claim that a point p is an element of $S \setminus \bigcup \underline{U}(R)$ not just *if*, but also *only if*, p lies on a boundary line of $V(P)$, for some $P \in \Pi(R)$? The answer is negative, just as we observed when comparing $b_1(\underline{u}(r_i))$ to $b_2(\underline{u}(r_i))$. If, apart from disjointness, no assumptions are made about the generator regions, then the associated collated diagram may well have a boundary region that contains *gaps*, in the following precise sense:

Definition 4.4 Given a space S and a collated diagram $\underline{U}(R)$ of S , a region g is a *gap* of $S \setminus \bigcup \underline{U}(R)$ iff (i) $g \subset S \setminus \bigcup \underline{U}(R)$; (ii) g is connected; (iii) $g \cap B(R, S) = \emptyset$; and (iv) for every set $S' \subseteq S$, if $g \subset S'$, then $S' \cap B(R, S) \neq \emptyset$.

To put this differently, a gap in a boundary region $S \setminus \bigcup \underline{U}(R)$ is any connected region of points all of which lie in that boundary region, none of which lie on a boundary line of any elements of $\mathcal{V}(R)$, such that if all points are properly included in another set, then that other set does contain points lying on a boundary line of some element of $\mathcal{V}(R)$.

We say that a boundary region of a collated diagram $\underline{U}(R)$ of a space S is *full* iff $S \setminus \bigcup \underline{U}(R)$ contains no gaps, or equivalently, iff $B(R, S) = S \setminus \bigcup \underline{U}(R)$; otherwise it is said to be *gappy*. Figure 2 shows a collated diagram generated by four regions each consisting of a pair of points lying relatively close to each other. The boundary region of this diagram is obviously gappy.

For the kind of application of collated diagrams that we are envisioning, an important question is whether, under certain conditions for generator regions, boundary regions are full. The following theorem provides the requisite information:

Theorem 4.1 *Let $R = \{r_1, \dots, r_n\}$ be a set of pairwise disjoint regions of a given space each of which is connected. Then the collated diagram $\underline{U}(R)$ on that space has a full boundary region.*

The formal proof of Theorem 4.1 is given in the [Appendix](#). But to also give an intuitive feel for it, we briefly point to a result concerning so-called dynamic Voronoi diagrams. Consider a set of points in a space that move along continuous paths through space, as time passes. At each particular time, the set generates a Voronoi diagram. One can think of the family of diagrams as a “dynamic” Voronoi diagram, a diagram constantly in flux. As is shown in Roos and Noltemeier [14], dynamic Voronoi diagrams change in an entirely continuous fashion.¹⁰ This is so even though the movements of the various

¹⁰See also Roos [13] and Gavrilova and Rokne [8].

points may bring about what are called *topological events*, that is, changes in adjacency relations between generator points. Such changes can obviously only be discontinuous, adjacency being a categorical matter (two generator points either share a Voronoi edge/face or they do not). But even these discontinuous *topological* changes are accompanied by continuous *geometrical* changes in the dynamic Voronoi diagram, that is, continuous movements of the diagram's Voronoi points and edges/faces. For a clear explanation of this phenomenon, see Gavrilova and Rokne [8: 94–96].

This geometric continuity ensures that, when all the Voronoi edges/faces and Voronoi points of the separate Voronoi diagrams generated at each particular point in time are projected onto each other, the result is a diagram with a full boundary region. Theorem 4.1 should then be more or less obvious once we realize that, first, because the generator regions are connected, from any selection of points from these regions we can reach any other selection of points from those regions by having all points move along continuous trajectories that lie entirely within the generator regions; and second, we can exhaust the possible selections of points from the generator regions by, so to speak, chaining together infinitely many such operations.

We further state, without proof, that Lemma A.2, which is used for proving Theorem 4.1 in the [Appendix](#), has the following corollary:

Corollary 4.1 *Let $\underline{U}(R)$ be a collated diagram on a given space. Then, if R satisfies the conditions of Theorem 4.1, it holds that $b_1(\underline{u}(r)) = b_2(\underline{u}(r))$, for all $r \in R$.*

Because in the following we will be concerned only with collated diagrams $\underline{U}(R)$ for which R satisfies the conditions of Theorem 4.1, we will henceforth simply use the notation $b(\underline{u}(r))$ for the boundary region associated with collated polygon/polyhedron $\underline{u}(r) \in \underline{U}(R)$.

5 Vagueness Reconsidered

With the new Voronoi construction in place, and assuming prototypical *areas* rather than points, we obtain a picture of how human beings form categories that is not so different from the one put forth in the current conceptual spaces literature. In this picture, the set of all clear instances of a concept is identified with the collated polygon/polyhedron associated with the region representing the concept's typical instances in the relevant conceptual space. That the concepts thus defined satisfy Gärdenfors's criterion for naturalness follows from Fact 4.1.

The picture also offers a straightforward explanation of the phenomenon that something can be borderline X/Y while being a bit more X -ish than Y -ish. To make the explanation entirely precise, we first note that, while the metrics so far assumed are defined for pairs of points only, they can be swiftly generalized in a way that allows us to measure distances between points and

regions, or between regions. Let Φ and Ψ be nonempty regions of a space S . Then the *Hausdorff distance* between Φ and Ψ is defined to be

$$\max\{\sup_{x \in \Phi} \inf_{y \in \Psi} \delta_S(x, y), \sup_{y \in \Psi} \inf_{x \in \Phi} \delta_S(x, y)\}.$$

Figure 3 illustrates this definition. Note that because Φ or Ψ may consist of only one point, we can meaningfully speak of the Hausdorff distance between single points and regions.

Given our assumption that, at least for the kind of properties we focus on, prototypes are in general non-unique, one natural way to explicate expressions like “ a is more X than Y ” would be in terms of the above measure. More precisely, the idea would be that “ a is more X than Y ” is true iff a is closer to the area of prototypical X ’s than to the area of prototypical Y ’s, where closeness is measured by the Hausdorff metric. Then—to state our solution to the thickness problem—we can say that, because in the present picture, borderlines need not be lines in the strict sense, but may be broader areas, there can be points p lying in $b(\underline{u}(r)) \cap b(\underline{u}(r'))$, where r is the region of prototypical instances of X and r' the region of prototypical instances of Y , such that p is closer (on the Hausdorff metric) to r than to r' , which, given the intended interpretation of the metric, means that p is more X than it is Y .

We note at once that if boundary regions of the relevant conceptual spaces could be gappy, the explanation might not be so straightforward after all. In particular, it might not be immediately obvious how to interpret the existence of gaps, given that points lying in a gap could easily appear to have some special status, or at least to be not quite on a par with other points in the boundary region. For—it might be said—they correspond to cases that on each particular choice of points from the prototypical areas come out as being non-borderline, while still ending up as borderline cases.

However, Theorem 4.1 establishes, together with some highly reasonable assumptions about prototypical areas, that the possibility of gappy boundary regions cannot obtain. For consider that, first, connectedness of prototypical areas is an exceedingly weak assumption. We saw above that, for the purposes of constructing a generalized Voronoi diagram, Gärdenfors even assumed prototypical areas to be circles. While he does not provide an argument for that assumption, one way to justify it would be by pointing to the considerations of cognitive economy cited in Section 2: storing pairs of a prototypical point and a radius presumably does not exact a much greater toll on the brain than




Fig. 3 Shortest distances from the +’s to some element of the set of •’s (a); shortest distances from the •’s to some element of the set of +’s (b); the longest of those shortest distances, which is the Hausdorff distance between the two sets (c)

does storing prototypical points alone. We also note that the connectedness assumption is supported by the findings of Berlin and Kay [1, Section 1.5], given that whenever their test subjects pointed at more than one chip as being a typical instance of a color, these chips were always adjacent. Second, we can be assured that the areas of typical instances will be pairwise disjoint: it would be absurd to think that we might find typical instances of, say, red that are also typical instances of orange. This, too, is supported by Berlin and Kay's experimental work. Thus, where prototypical areas serve as generator regions of the collated diagram of the relevant space, Theorem 4.1 entails that this collated diagram has a full boundary region.

To summarize so far, we have arrived at a picture of categorization that permits a characterization of borderline cases which very naturally explains the phenomenological data about vagueness. In particular, we obtain a unified set of answers to what are unarguably among the most central questions about vagueness. What is a borderline case relative to a set of concepts? Answer: a case represented by a point in the boundary region of the conceptual space representing the concepts. What is, more specifically, a borderline case of X -ness (or, if one likes, an X /not- X borderline case)? Answer: a case represented by a point in $b(\underline{u}(r))$, where r is the region of typical instances of X . What is a borderline case between X and Y (an X/Y borderline case)? Answer: a case represented by a point in $b(\underline{u}(r)) \cap b(\underline{u}(r'))$, where r is the region of typical instances of X and r' the region of typical instances of Y . How can there be X/Y borderline cases that are more X than they are Y ? Answer: because borderlines between concepts need not be lines consisting of points at equal distance from two prototypes; they may have a certain thickness, so that some X/Y borderline cases can be closer to the prototypical X than to the prototypical Y area, some can be closer to the prototypical Y than to the prototypical X area, and some can be equidistant from the two areas.

We end this section with two comments on the foregoing. The first concerns the worry that our account of how to model vagueness in the conceptual spaces framework is seriously incomplete. Specifically, it might be thought that already the prototypical areas in which the typical cases of the concepts are to be found cannot be sharply delineated, so that we would have a residual problem of vagueness that could not be explained along the above lines.

We believe that this worry is misplaced. In particular, it would be wrong to think that if x is typically red and y has a color that is, so to speak, "directly adjacent" to the color of x , then y must be typically red, too. If typical instances are taken to have a metaphysically special status, then it may simply be a brute fact of nature that x is typically red and y is not, however similar they look. Even if typicality is taken in a metaphysically more neutral sense, it is far from obvious that very similar shades of a color must have the same status *qua* typical instance. For example, when going from somewhere in the prototypical area of red to the (non-prototypical) clear cases of red, at some point a little hesitancy may set in when one is asked, or otherwise contemplates, whether *this* specific shade of red is typically red. And, on the presumed understanding of typicality, it would seem reasonable to take any hesitancy on our part as

being evidence that this particular shade of red is *not* typically red, at least given normal circumstances.¹¹ Thus, as far as we can see, there is no reason to suppose that the area of typical instances of a vague concept must itself have vague boundaries.

Our second comment concerns the worry that Gärdenfors' original model of categorization does not only make borderlines too thin to be empirically adequate but also makes them come out as too *sharp*. This is a problem if Sainsbury [17] (among others) is right that a concept's having borderline cases is consistent with its having sharp boundaries, and that vagueness is not merely a matter of having borderline cases, but also of having blurred boundaries. More importantly, this so-called problem of higher-order vagueness may also seem to affect our model, given that collated diagrams by themselves do nothing to "fuzzify" borderlines: even if borderlines are really areas, these areas themselves have sharp boundaries.

In response, we make three remarks. The first is that it is not uncontroversial that there is such a thing as higher-order vagueness.¹² Secondly, we note that even if we have to grant an intelligible notion of higher-order vagueness, and also grant that our model is idealizing to some extent and therefore does not accommodate higher-order vagueness, we could still claim that it is less idealizing than Gärdenfors' model, which even fails to accommodate first-order vagueness. So our model would be a step forward in any case.¹³ Finally, however, and less indirectly, in addressing the problem of higher-order vagueness, we may exploit the fact, pointed out by Gärdenfors [6: 89], that psychological measures, such as the metrics defined on conceptual spaces, tend to be imprecise and that, in particular, people may be uncertain about the exact metrical structure of a given space, for example, because they are uncertain about the values for the weights in the relevant instance of Eq. 1, that is to say, about the relative scaling of the axes.¹⁴ For, as Gärdenfors (*ibid.*) notes, as a result of this uncertainty "the borderlines [of the Voronoi diagram of a

¹¹Doubts about the lighting conditions, or about one's present judgmental capacities, may of course cause some hesitancy. But often we can be sure enough that the lighting is normal and our judgment unimpaired. In such circumstances, hesitancy about whether a given case is typical may be taken to rule out typicality.

¹²According to Wright [18], for instance, higher-order vagueness is illusory.

¹³Sainsbury [16] argues that any semantics that allows the collection of definite truths to be a set, with sharp boundaries, is inadequate, since it eliminates vagueness, that is, the very phenomenon it is supposed to describe. But this point seems to overgenerate insofar as it may be turned into a general objection to the application of mathematical models of whatsoever, which are always idealizing in some way and not perfectly accurate descriptions of reality. For this point, compare Cook [3] and Gaifman [5: 43 ff].

¹⁴On uncertainty about the exact structure specifically of color space, see Clark [2: 130 f], Fairchild [4: 303], Wyszecki and Stiles [19: 165, 825–830], and Kuehni [10: 357].

conceptual space] will not be exactly determined.” As collated diagrams are composed out of simple Voronoi diagrams, the indeterminacy affecting the borderlines of the latter will carry over to the boundary regions of the former in that it will not be exactly determined where a collated polygon/polyhedron ends and the boundary region begins. At the same time, given that the boundary region is an area and not a line, there can be perfectly determinate borderline cases between the indeterminate beginnings and endings of the boundary region. For example, assuming a simple one-dimensional space, we may be uncertain where between coordinates 0 and 1 the boundary region between two concepts starts and where between 2 and 3 it stops—which may account for the fact that we can experience cases as being vaguely borderline—but we may be certain that any point lying in the interval $[1, 2]$ represents a borderline case, which may account for the fact that some borderline cases are experienced as being clearly borderline.¹⁵ This may be all there is to higher-order vagueness.

6 Matters Empirical

In arguing for the proposed extension of the conceptual spaces framework, we have referred to data concerning vagueness, like data showing that people do not experience sharp transitions from clear cases to borderline cases, or that we can differentiate between clear borderline cases with regard to their distance from clear (non-borderline) cases. These data are introspectively accessible to anyone and do not need the backing of experiments. In the present section, we want to argue that the proposed account of vagueness suggests experiments that may generate further data relevant to vagueness, data which are not as easily accessible as the aforementioned ones, and that may support or undermine the account (as the case may be).

To introduce the idea for what is arguably the simplest and most direct experiment we have in mind, note that, for any one-dimensional conceptual space, it follows from our proposal that the length of the borderline area between two concepts is a linear function of the lengths of the prototypical areas, to wit, the sum of the lengths of the prototypical areas divided by

¹⁵It may at first be tempting to try and explain borderline vagueness purely in terms of the uncertainty of psychological measures. However, this approach would seem to supply insufficient means for doing justice to such intuitions as that there can be perfectly determinate borderline cases, and that of two determinate red–orange borderline cases (say), one is more reddish than the other.

2.¹⁶ Just for expository purposes, consider a one-dimensional color space for monochromatic color samples, having wavelength as its only dimension. Now suppose that, for a given person, the typical cases of green are to be found between 500 and 520 nanometers, the typical cases of yellow between 570 and 580 nm, and the typical cases of red between 650 and 660 nm.¹⁷ Then our account predicts that, for the given person, the clear green/yellow borderline cases are to be found between, roughly, 535 and 550 nm, and the clear yellow/red borderline cases between, roughly, 610 and 620 nm (the qualifier “roughly” is there on account of the presumed imprecision of the relevant metric). At a minimum, it predicts that there is more variation between the clear green/yellow borderline cases than between the clear yellow/red borderline cases.

To make this idea more general, recall that, given a collated diagram $\underline{U}(R)$, the region $b(\underline{u}(r))$ was defined to be the boundary region associated with a collated polygon/polyhedron $\underline{u}(r) \in \underline{U}(R)$. Further notice that, depending on the dimensionality of the relevant space, we can speak of the length, surface, or volume of $b(\underline{u}(r))$. Then we want to make the following conjecture:

Conjecture 6.1 For every collated diagram $\underline{U}(R)$ of a given space and every $r \in R$, the length/surface/volume of $b(\underline{u}(r))$ correlates with the length/surface/volume of r , all else being equal.

While, as we saw in note 16, for one-dimensional spaces this conjecture—even a claim more specific than this conjecture—can be proven, we do not currently see how to prove the conjecture for higher-dimensional spaces. Nevertheless, it is easy to convince oneself of the correctness of the conjecture

¹⁶*Proof* Given a one-dimensional space S , consider two prototypical areas (line segments) of this space, ab and cd , and let ef be the boundary region between them, as defined by a collated diagram of S . Assume, without loss of generality, that b lies between a and e and c lies between f and d , and note that $\delta_S(a, e) = \delta_S(a, c)/2$ and $\delta_S(a, f) = \delta_S(a, b) + \delta_S(b, d)/2$. Then we have

$$\begin{aligned} \delta_S(e, f) &= \delta_S(a, f) - \delta_S(a, e) \\ &= \delta_S(a, b) + \frac{\delta_S(b, d)}{2} - \frac{\delta_S(a, c)}{2} \\ &= \delta_S(a, b) + \frac{\delta_S(b, c) + \delta_S(c, d)}{2} - \frac{\delta_S(a, b) + \delta_S(b, c)}{2} \\ &= \delta_S(a, b) + \frac{\delta_S(c, d) - \delta_S(a, b)}{2} \\ &= \frac{2\delta_S(a, b) + \delta_S(c, d) - \delta_S(a, b)}{2} \\ &= \frac{\delta_S(a, b) + \delta_S(c, d)}{2}. \end{aligned}$$

□

¹⁷There is ample experimental evidence showing that different people will designate somewhat different areas as containing the typical instances of various colors; see Berlin and Kay [1, Section 1.5] and Hardin [9: 162].

at least for two-dimensional spaces by considering simulations such as the ones depicted in Fig. 4. Simulations that allow generator areas to have different surfaces—such as the ones depicted in the lower row of Fig. 4—also make it easy to appreciate that the volume (respectively, length/surface) of the boundary region associated with $\underline{u}(r)$ depends not only on the volume (length/surface) of r but also on the volumes (lengths/surfaces) of the generator regions of the collated polygons/polyhedrons neighboring $\underline{u}(r)$. This is one reason why the *ceteris paribus* clause is required. The other reason is that the volume (etc.) of the said boundary region also depends on the *locations* of the neighboring generator regions. This cannot be gleaned from Fig. 4, but it is not hard to imagine.

Conjecture 6.1 is a purely mathematical statement about purely mathematical constructs (collated Voronoi diagrams). However, if these constructs provide faithful models of human categorization, then this statement points to ways in which we may obtain evidence for our account of vagueness.

The simplest way to gain such evidence is probably to record the cases that persons classify as being typical instances of various concepts and those they classify as being borderline cases of those concepts, and then check for the relevant correlations. For example, a straightforward test would be to repeat the experiment reported in Berlin and Kay [1] that was mentioned at the beginning of Section 3, this time also paying close attention to subjects'




Fig. 4 An ordinary Voronoi diagram (*top left*); a collated diagram with generator regions uniformly of diameter 1 (*top middle*); a collated diagram with generator regions uniformly of diameter 4 (*top right*); collated diagrams with generator regions varying in diameter from 1 to 4 (*bottom row*)

responses to the boundary areas of colors (which Berlin and Kay [1: 13] found too unreliable to give much weight in their analysis). Supposing Conjecture 6.1 to be correct, the prediction would be that there will be a correlation between the surfaces of the areas in which the typical instances of the colors are located and the surfaces of the color boundaries.

There may be other, less direct ways in which Conjecture 6.1 can manifest itself experimentally. Raffman [12: 52] notes that when we go through a series of color patches, ranging from a clearly red one to a clearly orange one, with each new patch slightly different in color from the foregoing one, then, if we are forced to judge them as either being red or being orange, the point at which our judgments will switch from red to orange (or conversely) may depend on whether we start at the “red end” and then proceed, patch by patch, to the “orange end,” or whether we go through this process in the opposite direction. This so-called hysteresis effect is presumed to be due to “a kind of judgmental inertia” created by the starting point (*ibid.*). In the paper just cited, Raffman based her claim not on an experiment but on introspection and extrapolation from known data about seemingly analogous phenomena (such as the occurrence of order effects in experiments concerning assessments of sizes of objects). However, experimental results reported in her [Raffman, D., forthcoming, Ch. 5] seem to buttress the claim. These same results further suggest that subjects who are forced to choose in the way that was described tend to stick to their initial judgment as far as possible, only switching judgments when they encounter clear cases of the other color. Now, we will call the point at which one switches from red to orange when one starts at the red end the “red/orange switching point,” and the point at which one switches from orange to red when one goes in the other direction the “orange/red switching point.” If Raffman is right, then the distance between the red/orange and orange/red switching points (for instance, as measured in number of patches that lie between them) will not only be *constrained* by the size of the boundary region between red and orange, but may be expected to *correlate* with the size of that region. Given that, on our account, the size of this region in turn correlates with the sizes of the prototypical areas of red and orange, the distance between the red/orange and orange/red switching points should correlate with the sizes of those areas. We may hence be able to find further support for our account by checking for correlations between, on the one hand, distances between switching points in the kind of processes described by Raffman and, on the other hand, sizes of prototypical areas. For example, a supporting finding would be that the prototypical areas for red and orange are both small compared to the prototypical areas of yellow and green, and that the red/orange and orange/red switching points are closer to one another than the yellow/green and green/yellow switching points.

Here, we have only indicated some ways in which our proposal can be put to the test. An obvious avenue for future research is to carry out these and similar experiments. Needless to say, if these experiments turn out negatively—for instance, if the predicted correlations between boundary regions and prototypical regions are found not to exist—then that will deliver a serious blow

to our proposal. In fact, this would also have negative implications for the conceptual spaces approach more generally (even if it would not refute it), in as much as our account seems the natural way to extend that approach to vague concepts, and given that, as we said, any account of concepts must be able to accommodate the phenomena related to vagueness if it is to be adequate.

Acknowledgements A version of this paper was presented at the Similarity and Vagueness conference held in Paris in May 2010, at the Universities of Amsterdam, Groningen, Leuven, Lund, and Toronto, and at the London School of Economics. We thank the audiences for stimulating questions and comments. We are also grateful to Jake Chandler, Vincent de Gardelle, David Etlin, Peter Gärdenfors, Leon Horsten, Jennifer Nagel, Wlodek Rabinowicz, Dianna Raffman, David Ripley, Jos Uffink, Christopher von Bülow, Jonathan Weisberg, and three anonymous referees for this journal for helpful comments and discussions. Hans van Ditmarsch gave useful editorial advice, for which we thank him. For financial support, we would like to thank the Government of the Flemish Community (project “Formal Epistemology: Foundations and Applications”) and the Agence Nationale de la Recherche (project “Cognitive Origins of Vagueness,” ANR-07-JCJC-0070-01).

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

Appendix: Proof of Theorem 4.1

In this Appendix, we prove Theorem 4.1. The proof of the theorem, as well as the proofs of the lemmas we make use of, assume a standard Euclidean metric, which assigns equal weight to each dimension. This assumption allows us to speak of circles, spheres, and so forth, but otherwise does not affect the generality of the proofs.

Lemma A.1 *Given a Voronoi diagram $V(P) = \{v(p_1), \dots, v(p_n)\}$ on some space S , if some $p_{i_1}, \dots, p_{i_m} \in P$ lie on a circle (or, more generally, an m -dimensional sphere) that is otherwise empty of elements of P , then the center of that circle (sphere) lies on a Voronoi edge (face), or is a Voronoi point, of $V(P)$.*

Proof Suppose the condition holds. By the definition of a Voronoi diagram, each point in S either lies in $v(p)$ for precisely one $p \in P$ or lies on a Voronoi edge (face) or is a Voronoi point. Because the center of the circle (sphere) is closer to any of the p_{i_k} than it is to any other point in P , it is not an element of $v(p)$ for some $p \neq p_{i_k}$ ($1 \leq k \leq m$). Because the center is equally close to each of p_{i_1}, \dots, p_{i_m} , it is neither an element of $v(p_{i_k})$ for precisely one $k \in \{1, \dots, m\}$. So it lies on a Voronoi edge (face) or is a Voronoi point. \square


Lemma A.2 *Let $\underline{U}(R)$ be a collated diagram of a space S , with $R = \{r_1, \dots, r_n\}$ such that the r_i are pairwise disjoint and each r_i is connected. Then consider points $a, b \in r$ for some $r \in R$ (possibly with $a = b$) and two line segments l and m such that (i) for some $v(a)$ and P such that $v(a) \in V(P) \in \mathcal{V}(R)$, l is a Voronoi*

edge (in the higher-dimensional case, a Voronoi face) of $v(a)$; (ii) for some $v'(b)$ and P' such that $v'(b) \in V(P') \in \mathcal{V}(R)$, m is a Voronoi edge (face) of $v'(b)$; and (iii) a line segment with a as one endpoint intersects l in point s , and m in point s' , such that $\overline{as} < \overline{as'}$ (see Fig. 5). Then any point c lying on $\overline{ss'}$ is an element of $B(R, S)$.

Proof Without loss of generality, let the geometric situation be as depicted in Fig. 5. So l is the Voronoi edge a shares with a' , and m is the Voronoi edge b shares with b' . Then consider a point c lying between l and m on $\overline{ss'}$, and let a' be an element of generator region $r' \in R$. Suppose all points in r' lie within the circle C with center c and such that a lies on C , and suppose, again without loss of generality, and as is the case in Fig. 5, that a is closer to c than b is to c . Then c is always closer to some point in r' than it is to b , so that for no $v^*(b)$ would $c \in v^*(b)$. But then m cannot be a Voronoi edge of b , contrary to our assumption. (This also holds if a and b are equally close to c ; if b is closer to a , one can rerun the argument with the tags interchanged.) Thus, some points in r' must lie outside C . Given that r' is connected, there must be points in r' lying on C . Now either there are regions $r^* \in R$ alongside r' and the region r in which a lies such that r^* has points lying in C , or there are no such further regions.

- (i) Consider the latter case first. Say that $p \in r'$ lies on C . Then for some $P \in \Pi(R)$, a and p lie on a circle that otherwise does not contain points of P . It then follows from Lemma A.1 that c lies on a Voronoi edge, or is a Voronoi point, of $V(P)$. Because $V(P) \in \mathcal{V}(R)$, c will be an element of $B(R, S)$.
- (ii) Next consider the case in which there are further regions r_1^*, \dots, r_k^* that have points lying inside C . Again, none of these regions can lie entirely within the circle, for then c would never be closer to b than to any other point, which, we saw, excludes that m is a Voronoi edge that b shares with b' . So each r_i^* will have points lying outside C and therefore, by connectedness, also points lying on C . There will thus be a $P \in \Pi(R)$ such

Fig. 5 Illustration of Lemma A.2




that P contains of each of r', r_1^*, \dots, r_k^* exactly one point which lies on C and no other points picked out by P lie inside C . Again by Lemma A.1, it follows that c lies on a Voronoi edge, or is Voronoi point, of $V(P)$ (it is in effect easy to prove that c is a Voronoi point of $V(P)$). So, c will again belong to $B(R, S)$.

Hence, c will be an element of $B(R, S)$ in either case. As c was arbitrarily chosen among the points lying on $\overline{ss'}$, it follows that all points lying on that line segment belong to $B(R, S)$. \square

Theorem A.1 *Let $R = \{r_1, \dots, r_n\}$ be a set of pairwise disjoint regions of a space S such that each region is connected. Then the collated diagram $\underline{U}(R)$ of S has a full boundary region.*

Proof Suppose, towards a contradiction, that $\underline{U}(R)$ satisfies the conditions of the theorem, yet that $S \setminus \bigcup \underline{U}(R)$ contains gaps. Let g be one such gap, and let a be a point lying in g (see Fig. 6 for an illustration). Further, consider a point $b \in S \setminus \bigcup \underline{U}(R)$ such that b is a limit point of g (where “limit point” is understood in the topological sense: each open set that contains b has a nonempty intersection with g). Then for some $P \in \Pi(R)$ there will be a $p \in P$ such that b lies on the boundary line of $v(p) \in V(P) \in \mathcal{V}(R)$. Furthermore, let b be so chosen that $\delta_S(b, p) > \delta_S(a, p)$ (such points must exist if a lies in g). Then $a \in v(p)$. Let $r \in R$ be the generator region containing p . Then note that there cannot be a $p' \in r$ and a $P' \in \Pi(R)$ such that $a \notin v'(p') \in V(P') \in \mathcal{V}(R)$, for else the conditions of Lemma A.2 would be satisfied, and thus we would have that $a \in B(R, S)$. Moreover, a cannot lie on a boundary line, given that otherwise, again, a would be in $B(R, S)$. So, for all $v(p) \in \{v(p)\}_{r \in R}$, $a \in v(p)$, and so $a \in \underline{u}(r)$. But, by Definition 4.4, $g \subset S \setminus \bigcup \underline{U}(R)$. Hence, because $a \in g$, also $a \in S \setminus \bigcup \underline{U}(R)$. However, together with the just-established fact that $a \in \underline{u}(r)$, it follows that $\underline{u}(r) \cap S \setminus \bigcup \underline{U}(R) \neq \emptyset$. This is inconsistent with the

Fig. 6 Point a lying in gap g of a boundary region (partly represented, in gray); if there were a $v'(p')$ that did not contain a (e.g., the area partly represented by the thin lines), the conditions of Lemma A.2 would be satisfied



assumption that $\underline{U}(R)$ is a collated diagram, for it is a direct consequence of Definition 4.2 that, for all $r_i \in R$, $\underline{u}(r_i) \cap S \setminus \bigcup \underline{U}(R) = \emptyset$. So, since our assumption that there is a gap in the boundary region of $\underline{U}(R)$ leads to a contradiction, and since $\underline{U}(R)$ was an arbitrary collated diagram generated by a set of pairwise disjoint regions each of which is connected, it follows that there can be no such diagrams with gappy boundary regions. \square

References

1. Berlin, B., & Kay, P. (1969/1999). *Basic color terms*. Stanford CA: CSLI Publications.
2. Clark, A. (1993). *Sensory qualities*. Oxford: Clarendon.
3. Cook, R. (2002). Vagueness and mathematical precision. *Mind*, *111*, 225–247.
4. Fairchild, M. D. (1998). *Color appearance models*. Reading, MA: Addison–Wesley.
5. Gaifman, H. (2010). Vagueness, tolerance and contextual logic. *Synthese*, *174*, 5–46.
6. Gärdenfors, P. (2000). *Conceptual spaces*. Cambridge, MA: MIT Press.
7. Gärdenfors, P. (2007). Representing actions and functional properties in conceptual spaces. In T. Ziemke, J. Zlatev, & R. M. Frank (Eds.), *Body, language and mind* (Vol. 1, pp. 167–195). Berlin: De Gruyter.
8. Gavrilova, M., & Rokne, J. (1999). Swap conditions for dynamic Voronoi diagrams for circles and line segments. *Computer Aided Geometric Design*, *16*, 89–106.
9. Hardin, C. L. (1988). *Color for philosophers*. Indianapolis, IN: Hackett.
10. Kuehni, R. (2003). *Color space and its divisions*. New York: Wiley.
11. Okabe, A., Boots, B., Sugihara, K., & Chiu, S. N. (2000). *Spatial tessellations* (2nd ed.). New York: Wiley.
12. Raffman, D. (1994). Vagueness without paradox. *Philosophical Review*, *103*, 41–74.
13. Roos, T. (1993). Voronoi diagrams over dynamic scenes. *Discrete Applied Mathematics*, *43*, 243–259.
14. Roos, T., & Noltemeier, H. (1991). Dynamic Voronoi diagrams in motion planning. In H. Bieri, & H. Noltemeier (Eds.), *Lecture notes in computer science* (Vol. 553, pp. 227–236). New York: Springer.
15. Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, *104*, 192–232.
16. Sainsbury, M. (1990). Concepts without boundaries. Inaugural Lecture, Department of Philosophy, King's College, London; Reprinted In R. Keefe, & P. Smith (Eds.), *Vagueness. A reader* (pp. 251–264). Cambridge, MA: MIT Press, 1996.
17. Sainsbury, M. (1991). Is there higher-order vagueness? *Philosophical Quarterly*, *41*, 167–182.
18. Wright, C. (2010). The illusion of higher-order vagueness. In R. Dietz, & S. Moruzzi (Eds.), *Cuts and clouds: Vagueness, its nature and its logic* (pp. 523–549). Oxford: Oxford University Press.
19. Wyszecki, G., & Stiles, W. S. (2000). *Color science*. New York: Wiley.