

Continuity and the Logic of Perception

by

John L. Bell

In his *On What is Continuous* of 1914 ([2]), Franz Brentano makes the following observation:

If we imagine a chess-board with alternate blue and red squares, then this is something in which the individual red and blue areas allow themselves to be distinguished from each other in juxtaposition, and something similar holds also if we imagine each of the squares divided into four smaller squares also alternating between these two colours. If, however, we were to continue with such divisions until we had exceeded the boundary of noticeability for the individual small squares which result, then it would no longer be possible to apprehend the individual red and blue areas in their respective positions. But would we then see nothing at all? Not in the least; rather we would see the whole chessboard as violet, i.e. apprehend it as something that participates simultaneously in red and blue.

In this paper I will describe a simple and natural framework—a *logic of perception*—in which this “simultaneous participation” or *superposition* of perceived attributes is accorded a major role. (This framework was originally introduced in [1] for a different purpose.) The central concept of the framework is that of an attribute being *manifested* over a region or part of a *proximity space*—an abstract structure embodying key features of perceptual fields. An important property of the manifestation relation is *nonpersistence*, namely, the fact that a space may manifest an attribute not manifested by some part. This will be shown to be closely related to the idea of superposing attributes.

I will also show how this framework is tied up with the continuity of perceptual fields.

*

Let us think of *attributes* or *qualities* such as “blackness”, “hardness”, etc. as being *manifested over* or *supported by* parts of a (perceptual) *space*. For instance if the space is my total sensory field, part of it manifests blackness and part manifests hardness and, e.g., a blackboard manifests both attributes. Each attribute α is correlated with a *proposition* (more

precisely, a propositional function) of the form “— manifests the attribute α .”

I assume given a supply of *atomic* or *primitive* attributes, i.e., attributes not decomposable into simpler ones: these will be denoted by A, B, C . For each primitive attribute A and each space \mathbf{S} we may consider the total part of \mathbf{S} which manifests A ; this will be called the A -part of \mathbf{S} and denoted by $[[A]]_{\mathbf{S}}$. Thus, for instance, if \mathbf{S} is my visual field and A is the attribute “redness”, then $[[A]]_{\mathbf{S}}$ is the total part of my sensory field where I see redness: the red part of my visual field.

Attributes may be *combined* by means of the logical operators \wedge (and), \vee (and/or), \neg (not) to form *compound* or *molecular* attributes. The term “attribute” will accordingly be extended to include compound attributes. It follows that (symbols for) attributes may be regarded as the statements of a propositional language \mathbf{L} —the *language of attributes*.

In order to be able to correlate parts of any given space \mathbf{S} with compound attributes, i.e., to be able to define the A -part of \mathbf{S} for arbitrary compound A , we need to assume the presence of operations \sqcap, \sqcup, \sim corresponding respectively to \wedge, \vee, \neg , on the parts of \mathbf{S} . For then we will be able to define the α -part $[[\alpha]]_{\mathbf{S}}$ for arbitrary attributes α according to the following scheme:

$$\begin{aligned} [[\alpha \wedge \beta]]_{\mathbf{S}} &= [[\alpha]]_{\mathbf{S}} \sqcap [[\beta]]_{\mathbf{S}} \\ [[\alpha \vee \beta]]_{\mathbf{S}} &= [[\alpha]]_{\mathbf{S}} \sqcup [[\beta]]_{\mathbf{S}} \\ [[\neg\alpha]]_{\mathbf{S}} &= \sim[[\alpha]]_{\mathbf{S}} \end{aligned} \quad (*)$$

Once this is done, we can then define the basic relation $\sqsubseteq_{\mathbf{S}}$ of *inclusion* between attributes over \mathbf{S} :

$$\alpha \sqsubseteq_{\mathbf{S}} \beta \Leftrightarrow [[\alpha]]_{\mathbf{S}} \subseteq [[\beta]]_{\mathbf{S}}$$

where, as usual, “ \subseteq ” denotes the relation of set-theoretic inclusion.

Now the conventional meaning of “ \wedge ” dictates that, for any attributes α and β , we should have $\alpha \wedge \beta \sqsubseteq_{\mathbf{S}} A$ and $\alpha \wedge \beta \sqsubseteq_{\mathbf{S}} B$ and, for any γ , if $\gamma \sqsubseteq_{\mathbf{S}} \alpha$

and $\gamma \sqsubseteq_{\mathbf{S}} \beta$ then $\gamma \sqsubseteq_{\mathbf{S}} \alpha \wedge \beta$. In other words, $\llbracket \alpha \wedge \beta \rrbracket_{\mathbf{S}}$ should be taken to be the largest part (w.r.t. \sqsubseteq) of \mathbf{S} included in both $\llbracket \alpha \rrbracket_{\mathbf{S}}$ and $\llbracket \beta \rrbracket_{\mathbf{S}}$. By the first equation in (*) above, the same must be true of $\llbracket \alpha \rrbracket_{\mathbf{S}} \sqcap \llbracket \beta \rrbracket_{\mathbf{S}}$. Consequently, for any parts U, V of \mathbf{S} , $U \sqcap V$ should be the largest part of \mathbf{S} included in both U and V .

Similarly, now using the conventional meaning of “ \vee ”, we find that, for any parts U, V of \mathbf{S} , $U \sqcup V$ should be the smallest part of \mathbf{S} which includes both U and V .

We shall suppose that there is a *vacuous* attribute \perp for which $\llbracket \perp \rrbracket_{\mathbf{S}} = \emptyset$, the empty part of \mathbf{S} . In that case, for any attribute α , we have

$$\llbracket \alpha \rrbracket_{\mathbf{S}} \sqcap \sim \llbracket \alpha \rrbracket_{\mathbf{S}} = \llbracket \alpha \rrbracket_{\mathbf{S}} \sqcap \llbracket \neg \alpha \rrbracket_{\mathbf{S}} = \llbracket \alpha \wedge \neg \alpha \rrbracket_{\mathbf{S}} = \llbracket \perp \rrbracket_{\mathbf{S}} = \emptyset.$$

Consequently, for any part U of \mathbf{S} we should require that $U \sqcap \sim U = \emptyset$, i.e. that U and $\sim U$ be *mutually exclusive*.

It follows from these considerations that we should take the parts of a perceptual space \mathbf{S} to constitute a *lattice of subsets* of (the underlying set of) \mathbf{S} , on which is defined an operation \sim (‘complementation’) corresponding to negation or exclusion satisfying the condition of mutual exclusiveness mentioned above. Formally, a *lattice of subsets* of a set S is a family \mathbb{L} of subsets of S containing \emptyset and S such that for any $U, V \in \mathbb{L}$ there are elements $U \sqcap V, U \sqcup V$ of \mathbb{L} such that $U \sqcap V$ is the largest (w.r.t. \sqsubseteq) element of \mathbb{L} included in both U and V and $U \sqcup V$ is the smallest (w.r.t. \sqsubseteq) element of \mathbb{L} which includes both U and V . $U \sqcap V, U \sqcup V$ are called the *meet* and *join*, respectively, of U and V . A lattice \mathbb{L} of subsets of S equipped with an operation $\sim: \mathbb{L} \rightarrow \mathbb{L}$ satisfying $U \sqcap \sim U = \emptyset$ for all $U \in \mathbb{L}$ is called a \sim -*lattice* of subsets of S .

We can now formally define a *perceptual space*, or simply a *space*, to be a pair $\mathbf{S} = (S, \mathbb{L})$ consisting of a set S and a \sim -lattice \mathbb{L} of subsets of S . Elements of \mathbb{L} are called *parts* of \mathbf{S} , and \mathbb{L} is called the *lattice of parts* of S .

The perceptual spaces that most closely resemble actual perceptual fields are called *proximity spaces*. These in turn are derived from *proximity structures*. A *proximity structure* is a set S equipped with a *proximity relation*, that is, a symmetric reflexive binary relation \approx . Here we think of S as a field of perception, its points as *locations* in it, and the relation \approx as representing *indiscernibility of locations*, so that $x \approx y$ means that x and y are “too close” to one another to be perceptually distinguished. (Caution: \approx is *not* generally transitive!) For each $x \in S$ we define the *sensum* at x , Q_x , by

$$Q_x = \{y \in S: x \approx y\}.$$

We may think of the sensum Q_x as representing the *minimum perceptibilium* at the location x . Unions of families of sensa are called *parts* of S . Parts of S correspond to *perceptibly identifiable subregions* of S . It can be shown that the family $\text{Part}(S)$ of parts of S forms a \sim -lattice of subsets of S (actually, a complete ortholattice) in which the join operation is set-theoretic union, the meet of two parts of S is the union of all sensa included in their set-theoretical intersection, and, for $U \in \text{Part}(S)$,

$$\sim U = \{y \in S: \exists x \notin U. x \approx y\}.$$

The pair $\mathbf{S} = (S, \text{Part}(S))$ is called a *proximity space*.

The most natural proximity structures (and proximity spaces) are derived from *metrics*. Any metric d on a set S and any nonnegative real number ε determines a proximity relation \approx given by $x \approx y \Leftrightarrow d(x, y) \leq \varepsilon$. When $\varepsilon = 0$ the associated proximity relation is the *identity relation* $=$: the corresponding proximity space is then called *discrete*. It can be shown that, if a proximity space \mathbf{S} has a *transitive* proximity relation, then it is *almost discrete* in the sense that its lattice of parts is isomorphic to the lattice of parts of a discrete space.

Given a perceptual space $\mathbf{S} = (S, \mathbb{L})$ we define an *interpretation* of the language \mathbf{L} of attributes to be an assignment, to each primitive attribute A , of a part $\llbracket A \rrbracket_{\mathbf{S}}$ of \mathbf{S} . Then we can extend the assignment of parts of \mathbf{S} to all attributes as in (*) above. Given an attribute α and a part U of \mathbf{S} , we think of the relation $U \subseteq \llbracket \alpha \rrbracket_{\mathbf{S}}$ as meaning that U is *covered* by the attribute A . Now there is another relation between parts and attributes the *manifestation relation* $\Vdash_{\mathbf{S}}$ —which reflects more closely the way compound attributes are built up from primitive ones. $U \Vdash_{\mathbf{S}} \alpha$, which is

read “ U manifests α ” or “ α is manifested over U ” is defined as follows:

$$\begin{aligned}
 U \Vdash_{\mathbf{S}} A &\Leftrightarrow U \subseteq \llbracket A \rrbracket_{\mathbf{S}} \text{ for primitive } A, \\
 U \Vdash_{\mathbf{S}} A \wedge B &\Leftrightarrow U \Vdash_{\mathbf{S}} A \text{ and } U \Vdash_{\mathbf{S}} B, \\
 U \Vdash_{\mathbf{S}} A \vee B &\Leftrightarrow V \Vdash_{\mathbf{S}} A \ \& \ W \Vdash_{\mathbf{S}} B \text{ for some parts } V, W \text{ of } \mathbf{S} \text{ such that} \\
 &U = V \cup W, \\
 U \Vdash_{\mathbf{S}} \neg A &\Leftrightarrow (\text{for all parts } V \text{ of } \mathbf{S}) V \Vdash_{\mathbf{S}} A \Rightarrow V \subseteq \sim U.
 \end{aligned}$$


Thus U manifests a disjunction $A \vee B$ provided there is a “covering” of U by two “subparts” manifesting A and B respectively, and U manifests a negation $\neg A$ provided any part of \mathbf{S} manifesting A is included in the “complement” of U .

In general, the manifestation and covering relations fail to coincide in proximity spaces. The reason for this is that, while the latter has a certain *persistence* property, the former, in general, fails to possess this property. By persistence of the covering relation is meant the evident fact that if a part U of a space is covered by an attribute, then this attribute continues to cover any subpart of U . However, as we shall see, this is not the case for the manifestation relation: there are attributes manifested over a part of a space which fail to be manifested over a subpart.

Let us call an attribute **S**-*persistent* (or persistent over **S**) if for all parts U, V of \mathbf{S} we have

$$V \subseteq U \ \& \ U \Vdash_{\mathbf{S}} A \Rightarrow V \Vdash_{\mathbf{S}} A.$$


(Note that a *primitive* attribute is always persistent. More generally, it is not hard to show that the same is true for any compound attribute not containing occurrences of the disjunction symbol \vee .) Let us call a space **S** *persistent* if every attribute is **S**-persistent (for any interpretation of **L** in **S**). We now give an example of a *nonpersistent* proximity space, a one-dimensional version of Brentano’s chessboard.




Consider the real line with the proximity relation \approx defined by $x \approx y \Leftrightarrow |x - y| \leq \frac{1}{2}$, and let \mathbf{R} be the associated proximity space. The sensum at a point x is then the closed interval of length 1 centred on x . Suppose now we are given two primitive attributes B ('blue') and R ('red'). Let the B -part of \mathbf{R} be the union of all closed intervals of the form $[2n, 2n + 1]$ and let the R -part of \mathbf{R} be the union of all closed intervals of the form $[2n - 1, 2n]$. To put it more vividly, we "colour" successive unit segments alternately blue and red. Clearly, then, \mathbf{R} manifests the disjunction $R \vee B$. But if U is the sensum $Q_1 = [\frac{1}{2}, 1\frac{1}{2}]$, then $R \vee B$ is *not* manifested over U , since U is evidently not covered by two subparts over which R and B are manifested, respectively—indeed U has no proper subparts.

Thus arises the curious phenomenon that, although we can tell, by surveying a (sufficiently large part of) the whole space \mathbf{R} , that the part U is covered by redness and blueness, nevertheless U —unlike \mathbf{R} —does not split into a red part and a blue part. In some sense redness and blueness are conjoined or *superposed* in U : it seems natural then to say that U manifests a *superposition* of these attributes rather than a disjunction. If we take the unit of length on the real line sufficiently small (or equivalently, redefine $x \approx y$ to mean $|x - y| \leq \varepsilon$ for sufficiently small ε) so that each interval of unit length represents the minimum length discernible to human visual perception, we have (essentially) Brentano's chessboard in one dimension. In that case, the "superposition" of the two attributes blue and red turns out to be *violet*, which is what we actually see.

Actually, the covering of our proximity space by parts like U looks like this:



while Brentano's chessboard looks like this:



But the two arrangements are obviously isomorphic.

The concept of superposition of attributes admits a very simple rigorous formulation. In the example we have just considered, the part U manifests a superposition of the attributes R and B just when there is a part V of the space which includes U and manifests $R \vee B$ (in this case, V may be taken to be the whole real line). This prompts the following

definition. Given a proximity space \mathbf{S} , an interpretation of \mathbf{L} in \mathbf{S} and attributes A, B , we say that a part U of \mathbf{S} *manifests a superposition of A and B* if there is a part V of \mathbf{S} such that $U \subseteq V$ and $V \Vdash_{\mathbf{S}} A \vee B$. Now for any attribute C , it is readily shown that

$$U \Vdash_{\mathbf{S}} \neg\neg C. \Leftrightarrow V \Vdash_{\mathbf{S}} C \text{ for some part } V \text{ such that } U \subseteq V.$$

So the condition that U manifest a superposition of A and B is just

$$U \Vdash_{\mathbf{S}} \neg\neg(A \vee B).$$

It follows that *a superposition is a double negation of a disjunction*. In the human visual field, then, the attribute “violet” is the double negation of the attribute “blue or red”. Similarly, the attribute “grey” is the double negation of the attribute “black or white”, etc.

Finally, we discuss the relationship between these ideas and continuity. Let us call a proximity structure (S, \approx) *continuous* if for any $x, y \in S$ there exist z_1, \dots, z_n such that $x \approx z_1, z_1 \approx z_2, \dots, z_{n-1} \approx z_n, z_n \approx y$. Continuity in this sense means that any two points can be joined by a finite sequence of points, each of which is indistinguishable from its immediate predecessor. (Thus, in the case of our nonpersistent proximity space above, continuity means that a red segment and a blue segment can always be joined by a violet line provided that the coloured segments are taken to be sufficiently small.) If d is a metric on S such that the metric space (S, d) is connected, then every proximity structure determined by d is continuous. When S is a perceptual field such as that of vision, the fact that it does not fall into separate parts means that it is connected as a metric space with the inherent metric. Accordingly every proximity structure on S determined by that metric is continuous. Note that this continuity emerges even when S is itself an assemblage of discrete “points”. This would seem to be the way in which continuity of perception is engendered by an essentially discrete system of receptors.

References

- [1] Bell, J.L., *A New Approach to Quantum Logic*. British Journal for the Philosophy of Science, **37**, 1986.

[2] Brentano, Franz, *Philosophical Investigations of Space, Time and the Continuum*. Barry Smith, translator. London: Croom Helm, 1988.