Forthcoming in D. Stokes, S. Biggs and M. Matthen (eds.) *Perception and Its Modalities*, Oxford University Press.

The Space of Sensory Modalities

Fiona Macpherson

University of Glasgow

Abstract: Is there a space of the sensory modalities? Such a space would be one in which we can represent all the actual, and at least some of the possible, sensory modalities. The position of the senses in this space would indicate how similar and how different the senses are from each other. I argue that we can construct such a space in a non-arbitrary manner. The space is one with a high number of dimensions. I explain that one can use principal component analysis (PCA) to reduce the number of dimensions, and that doing so might (a) reveal factors important for individuating the senses that we had not previously considered, (b) help us to define what a sense is, and (c) provide grounds for deciding what is one token sense rather than multiple token senses. Philosophers should consider whether PCA might be useful in taxonomising other entities.

Keywords: senses, sensory modalities, taxonomy, principal component analysis, perception, space, representation, phenomenal character, sensory organ, proximal stimuli

1. Introduction

Is there a space of the sensory modalities? Such a space would be one in which we can represent all the actual, and at least some of the possible, sensory modalities. The relative position of the senses in this space would indicate how similar and how different the senses were from each other. The construction of such a space might reveal unconsidered features of the actual and possible senses, help us to define what a sense is, and provide grounds that we might use to decide what is one token sense rather than multiple token senses.

In this chapter, I explore, refine, and defend the idea that we can construct such a space—an idea that I briefly proposed in earlier work (Macpherson 2011a and 2011b). In doing so, I defend the idea from an objection that Richard Gray (2012) has voiced. Gray claims that the dimensions of the space of sensory modalities must generate a non-arbitrary ordering of the senses. He argues that we cannot find dimensions that do so. I disagree. I identify different ways in which a space could be arbitrary but I argue that we can define a space of the sensory modalities in a non-arbitrary manner. I give examples of what the dimensions might be.

In the next section of this chapter, I outline the background issues. I describe the debate about how we should individuate the senses. I summarise the position that I have argued for elsewhere (Macpherson 2011a and 2011b), namely, that the senses do not form natural kinds. They differ in degree and we should use all of the four standard criteria, which are typically taken to be competing criteria, to individuate the senses in a fine-grained manner. In section

3, I introduce the idea of the space of modalities and describe the role of the criteria for individuating the senses in building the space of sensory modalities. In section 4, I show that we can find dimensions that non-arbitrarily order the senses, focusing particularly on the sensory organ criterion. In section 5, I give examples of dimensions for the other criteria that meet the conditions of non-arbitrariness laid out in section 4. In section 6, I show that if one produces a high-dimensional space, then it is often possible to use Principal Component Analysis (PCA) to produce a second space that has fewer dimensions but preserves all the similarity and difference relationships between the items plotted in the first space. One can also use the analysis to reduce the number of dimensions to any number that one wishes, if one is prepared to lose some similarity and difference data. However, the analysis ensures that as much data as possible is retained, compatible with the reduction. Such a reduction would allow us to reduce the space of the sensory modalities to three dimensions, allowing visualisation of the space of the senses. Finally, in that section, I discuss how this PCA may lead us to identify factors important for individuating the senses that we had not previously considered. And I describe how such analysis might make us reconsider our prior commitments concerning what the senses are, and what we should count as being token instances of senses.

2. Background Issues

There are several questions concerning the individuation of the senses that are distinct. Key among these are the following:

(a) How should we define what a sense is, rather than something that is not a sense? For example, what makes a sense different from an aardvark or an armadillo? Or different from a part of the body, like an arm or an artery? Or different from a bodily process, like arousal or arterisis?

(b) How do we determine what is one sense rather than multiple senses? For example, we know that while most people today think of touch as one sense, Plato held that it was three separate senses on phenomenological grounds: pressure, temperature, and pain (Classen 1993); and some modern scientists think likewise on the grounds that there are three distinctive types of receptors in the skin for detecting pressure, temperature, and pain (Craig 1996).

(c) What kind or type is a token sense, on the assumptions that it is indeed a sense, and that it is one and only one sense?[1]

The key question that is addressed in this chapter is (c), although, as we shall see, the answer to (c) may shed light on answers to (a) and (b).

Traditionally, there have been four main competing philosophical approaches to answering question (c). Each approach has identified a different

[1] Types are general sorts of things and tokens are particular concrete instances. (In this respect they are different from occurrences, which need not be concrete.) See Wetzel (2011).

criterion that should be used to determine what kind of sense a token sense is. The four criteria are these:

(i) *Representation*: a sense is the kind it is on account of the objects and properties that the experiences associated with that sense represent (or present) or about which the brain states associated with the sense carry information. For example, vision might be classified as the sense that represents three-dimensional coloured objects at a distance from our body. Touch might be categorised as the sense that represents the heat, texture, and shape of objects that press against our body. Hearing could be considered the sense that represents sounds around and in our body, and so on.

(ii) *Phenomenal character*: a sense is the kind it is in virtue of the nature of "what it is like" (Nagel 1974) for a subject to have the experiences that the sense produces. For example, typically, experiences produced by vision seem to us to have a certain subjective feel that is different from those of touch, which again is different from those of smell, taste, and hearing, and so on.

(iii) *Proximal stimulus*: a sense is the kind it is in virtue of the nature of the physical stimuli that stimulate the sense organ. For example, vision might be characterised as the sense that utilises electromagnetic waves, or those of 390–750 nanometres. Hearing could be regarded as that which makes use of pressure waves in a medium—or pressure waves between 0.0172 and 17.2 metres. Smell might be branded the sense whose proximal stimuli are volatile molecules.

(iv) *Sense organs*: a sense is the kind it is on account of the nature of the sensory organs and physiology, which may include the relevant parts of the brain, associated with it. For example one might claim that vision is the sense that uses eyes and, perhaps, the visual cortex. And hearing is the sense that utilises ears, and, perhaps, the sort of processing that the auditory cortex does, and likewise for the other senses.

I have argued (Macpherson 2011a and 2011b) that the senses cannot be clearly divided up into a discrete and limited number of different kinds and that the differences between the senses amounts more to a difference of degree than of such kinds. I will not go into all the details for thinking that here, but one reason is the large number of actual senses that vary to a greater or lesser degree from the human senses, and the existence of some that seem to share as many features with one type of human sense as another. When we also consider the possible senses that could exist, there seems no good principled way to carve up the senses into a limited number of discrete kinds.

Take vision as an example. There is 20/20 normal human vision. Slightly different are various forms of short- and long-sightedness, and the different forms of colour blindness. More different still are the various forms of visual agnosia. Then there is what we typically classify as vision in other creatures. For example, consider bee 'vision' that detects ultra-violet, as well as the light visible to humans, by means of five eyes of different types. Bee 'vision' is different from our vision in respect of the type of eyes being used. The proximal stimulus overlaps with that of human vision but includes also wavelengths of electromagnetic radiation smaller than those humans can detect. Like human vision, bee 'vision' can represent three-dimensional objects at a distance but, unlike human vision, it is sensitive to differences in surfaces that make them reflect more or less ultra-violet light. Is it similar enough to our vision to warrant the title of vision? And consider senses that are typically not classified as vision but which share features of human vision. For example, some snakes have an infra-red sense that detects electromagnetic waves of longer wavelength than visible light, which, like ordinary human vision, gives them information about three-dimensional objects at a distance from their body. The snakes' infra-red sensitive pits below their eyes are the organs of this sense. The pits are not usually classified as eyes, perhaps because the snakes have a sense that is more similar to our sight than their infra-red sense that utilises organs more like our eyes than their pits. However, it is unclear whether the pits should be thought of as form of eye or not. Moreover, the proximal stimulus of their infra-red sense is not visible light, but electromagnetic waves of longer wavelength. Is that proximal stimulus similar enough to that of human vision to allow us to count this sense as a form of vision? Consider also bat echo-location, which is like vision in that it detects three-dimensional objects at a distance from the body, yet it involves using organs best thought of as ears sensitive to sound waves. This sense is somewhat like human vision, somewhat like human hearing, and somewhat unique. Then there are possible forms of senses that we can imagine, such as the X-ray 'vision' of Superman, the 'vision' of the Terminator that can analyse something's composition, or a sense that combines features of vision, and other senses. Reflection on these cases suggests that there is no principled way to draw divisions between token instances of senses that carves them up into a limited discrete number of kinds.

If we adopt this view—that the senses do not come in a limited number of discrete kinds—we needn't suppose that the four criteria specified earlier are competing criteria. I've previously argued that the thought that they are competing arises only if we think that we need to pigeonhole each sense into a small number of discrete categories, such as the Aristotelian five: sight, hearing, touch, taste, and smell, or those plus a small number more. It is because the four criteria sometimes differ as to how they would categorise a sense, on the assumption that it has to be one of a fairly limited number, that it seems as if they are in competition. If we reject such a coarse-grained classification then we simply need note what each sense is like with respect to each of the four criteria. Reflecting on the properties that it has, we can note how similar or how different a sense is from one of the Aristotelian senses if we like, but we don't need to categorise it as being one or the other, or as something different altogether.

This view has the advantage that we can then use all the criteria in question to individuate the senses—which rightly allows all the criteria to matter to us practically and to matter to us given philosophical or scientific concerns. For example, suppose that I tell you that I am seeing the pot of green tea in front of me. I might do so because the representational criterion is of import—I want to let you know which properties of the teapot I can detect, such as its colour and shape at a distance from my body. Or the proximal stimulus criterion may be at issue: I may wish to indicate that you should leave the lights on if you want me to notice more about it. Or I may be trying to convey the phenomenal character of my experience to you. Finally, I may wish the sense-organ that is being used to be taken into account by you. I may be informing you that if you want me to continue using the sense that I am using to find out about the teapot then you had better not gouge out my eyes. Contra Nudds (2004), I don't believe that only one of these criteria is what makes the senses significant to us.

Finally, note two things. First, I want to be liberal about the criteria required to individuate the senses. If you think that others are needed and do not fall under those specified, then I would (prima facie) be happy to include them. Second, a problem for those who wish to divide the senses into natural kinds is how to use these criteria. For example, suppose we wanted to use the proximal stimulus criterion to determine which of a variety of candidate senses were vision, and which were not. We would have to decide what the proximal stimulus is in the case of vision. Should we say that vision is any sense that has electromagnetic waves as its proximal stimulus? Or is its proximal stimulus some limited portion of the electromagnetic spectrum, such a those waves detectable by human vision? My view allows us to neatly sidestep this worry. When we no longer need to identify the kind to which a sense belongs, we can simply note for each of the criteria the exact values that the sense takes with respect to it. This is an added advantage of my view.

3. The Role of the Criteria for Individuating the Senses

Suppose that one has noted for each sense the properties it has with respect to the four criteria mentioned: representation, phenomenal character, proximal stimulus, and sensory-organ. The senses, it seems, will be more or less different from one another. If we could construct a space of the sensory modalities, then we would be able to map how similar and how different the senses are from each other. We could place the senses in such a space in a way that recorded their nature vis-à-vis each of the criteria and which compared each to the other with regard to these. In my essay 'Individuating the Senses' (2011a), I expressed this as follows:

> We can think of [the] four criteria as defining a multidimensional space within which we can locate each of the . . . senses . . .Thus, human vision, bee 'vision', snake infrared perception and TVSS perception would each be located at a different place in the multidimensional space. . . . Plotting the actual senses in this space will allow one to see the similarities and connections between them yet, at the same time, to individuate the types in a nonsparse, fine-grained manner. . . . [W]e might find that the actual senses are to be found in clusters in this space. For we will find, for example, that human vision and bee 'vision' are closer together in this space than human vision and bat echolocation. Perhaps these clusters would correspond to the Aristotelian senses or the Aristotelian senses plus a few others. I suspect this might be the case. This would show us that the folk were trying to reflect complex facts about the types of senses that we find in the world using an over-simplistic model, but one whose origin is explicable given the facts. (2011a: 37–38)

One might hope to get a space of the kind illustrated in Figure 18.1 (whose dimensions are unspecified as yet).

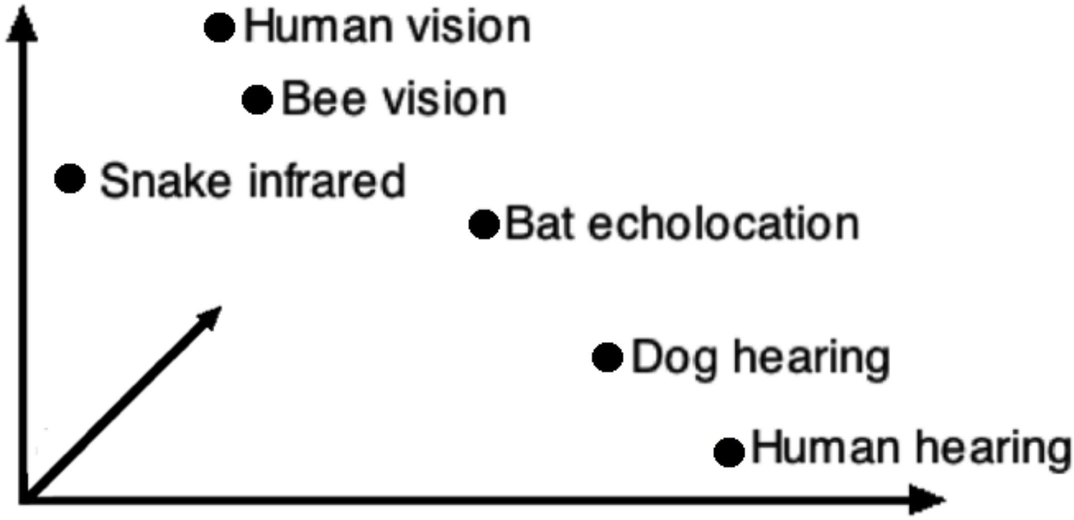

Figure 1: An illustration of an imagined desirable space of the sensory modalities

Such a space may shed light on the nature of the different kinds of senses and may yield insights into our actual commitments concerning what the senses are and what makes for one or more tokens of a sense or, indeed, suggest revisions to our prior commitments, in a way that I outline later.

In such space we could plot types of senses like human vision, bee vision, human smell, and so on. When doing so, we would plot values for normal or ideal instances of each of these types of senses. Or we could do it for token instances of each of the senses, such as my vision, your vision, your granny's vision, and your dog's vision. In this chapter I will suppose that we are doing the former, but for ease of exposition only. One could easily do either and, indeed, there may be advantages of doing the latter; nothing I will go on to say will rule out one's so doing.

This space of sensory modalities was only given a preliminary sketch in my previous work. My earlier quotation says that the four criteria 'define a multidimensional space'. One might, not unreasonably, have taken this to mean that each of the criteria forms a dimension of the space. However, it is very hard to see how one could construct such dimensions, as Richard Gray (2012) has, rightly, complained. More specifically, he claims that if we take the four criteria as dimensions then they could not generate a non-arbitrary ordering (that is, they would not generate an ordering that reflected actual similarity). Thus he concludes that such dimensions can't generate a space, that is to say a multidimensional ordering of the senses.

Compare a geometrical model of physical space, as illustrated in Figure 18.2.

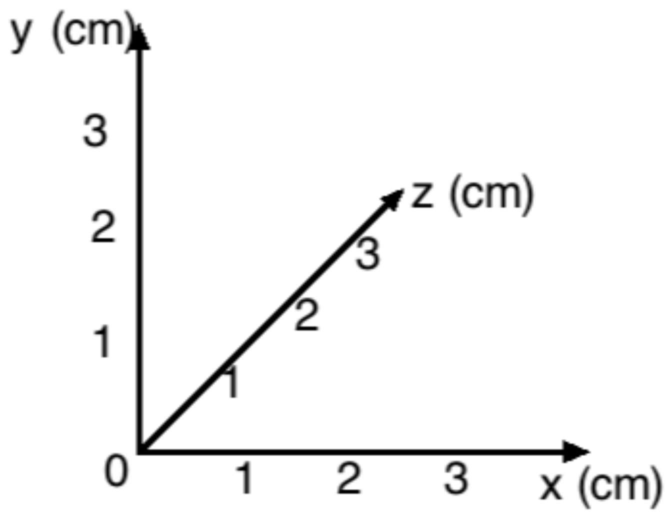

Figure2: Physical space defined by Cartesian Coordinates

The *x*, *y*, and *z* dimensions represent orthogonal directions in space away from the common origin point. The value that defines the *x* dimension is distance in centimetres away from the origin in the *x* direction—and similarly for the *y* and *z* dimensions. These dimensions determine the similarity relations—in this case, distance relations—between all the points. The dimensions do not generate an arbitrary ordering because points that are plotted nearer to each other in the model space correspond to points that are closer together in physical space.

Unlike the dimensions of this Cartesian space, the four criteria for individuating the senses do not constitute suitable dimensions of a space. To illustrate, we might try to specify a proximal stimuli dimension on which we could plot the senses depending on what their proximal stimulus was. But such a dimension would not give us a non-arbitrary ordering. For example, there would be nothing to choose between the two different dimensions illustrated in Figure 18.3.

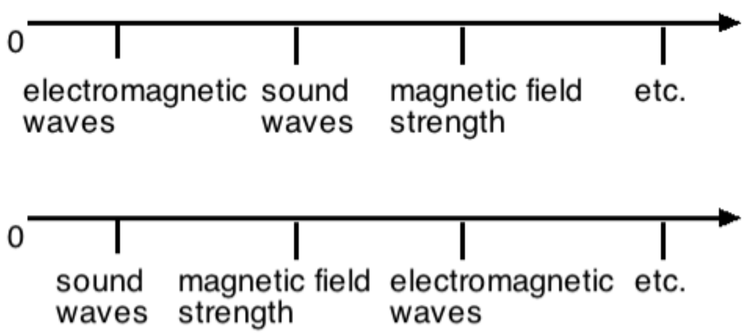

Figure 3: Potential Proximal Stimuli Dimensions

One might think that we can impose an ordering by just choosing one of the listed dimensions, rather than another. However, while we could just stipulate an ordering along a dimension, the position of any sense on the dimension would be arbitrary. The space between any two senses would not represent actual similarities and differences between the senses.

Thus, it is clear that the four criteria can't *constitute* the dimensions of the space of the senses. However, instead, we can let the four criteria *determine* the dimensions that do constitute the space. It is in this sense that we should hold that the four criteria define a multidimensional space within which we can locate each of the senses. In other words, we should take each of the four criteria in turn and define a set of dimensions that captures the properties that the criterion identifies as being those that are important for individuating the senses. If we are careful, I believe that we can find suitable non-arbitrary dimensions of that kind that will together define the space of sensory modalities.

In the next section I illustrate at length how to find suitable non-arbitrary dimensions that correspond to the proximal stimulus criterion. One reason for picking this dimension is that it is a particularly tricky case and one that Gray highlights as being problematic. In looking at this criterion, I will identify features, besides non-arbitrariness, that any dimension used to construct the space of the sensory modalities must have. Thus, the lessons learned by considering the dimensions corresponding to the proximal stimulus criterion should be applied to the dimensions corresponding to the other criteria for individuating the senses. In the section after that, section 5, I also indicate, more briefly, how one might construct suitable dimensions for the other criteria.

4. Defining Suitable Dimensions

How do we come up with a set of dimensions that captures the similarities and differences between the senses with respect to their different proximal stimuli? We can look to physics to specify all the nomologically possible kinds of proximal stimuli—for example, different frequencies of electromagnetic waves, sound waves, magnetic field strength; different temperatures and pressures; different amounts of chemicals in the air, chemicals in solution. With these different proximal stimuli identified we can create dimensions on which we can plot sensitivity to these stimuli. However, one has to be very careful in the way that one does this.

It is tempting but, as we will see, ultimately wrong, to think that we can specify one dimension corresponding to sensitivity to each of the different kinds of stimuli that physics identifies, to yield a space containing the sort of dimensions depicted in Figure 18.4.

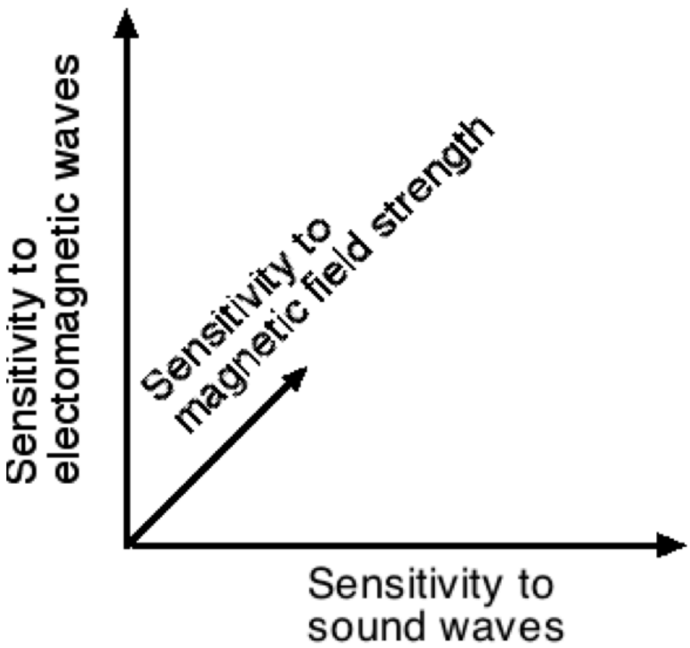

Figure 4: some dimensions one might think could be joined together in space corresponding to the proximal stimulus criterion

Note that an important and, for our purposes, very helpful fact about creating a space with a number of dimensions is that we don't have to order the dimensions that form a multidimensional space. In other words, if we have three dimensions—call them 1, 2 and 3—it doesn't matter if we make 1 the *x* dimension, 2 the *y* dimension, and 3 the *z* dimension, or whether we make 2 the *x* dimension, 1 the *y* dimension, and 3 the *z* dimension. One ends up with the same space either way. The naming of the dimensions as *x*, *y*, and *z* (and so on if more are required to form a higher-dimensional space) is purely for convenience of reference. Such naming makes no difference to the space created. In short, we simply need to identify dimensions and specify that they are joined together to form a space. We don't have to sequence the dimensions. This removes what one might have worried was a potential source of arbitrariness in the endeavour of creating the space of the sensory modalities. If there is no choice to be made about the order of the dimensions there is no arbitrariness that could creep into that choice.

One might have thought that the dimensions in Figure 18.4 could be used to record the sensitivities of each sensory organ to kinds of proximal stimulus. For example, consider the dimension labelled 'sensitivity to electromagnetic waves'. We might think that we could construct a dimension that takes different lengths of electromagnetic wavelengths as its value and then record the sensitivities of the different senses on it. See Figure 18.5 for illustrative purposes, on which I have plotted the following: human vision, which is sensitive to what is often called 'visible light', about 400–700 nanometres; bee vision, which is sensitive to visible light plus some infra-red wavelengths; snake infrared perception, which is sensitive in the infra-red area of the spectrum; and fictional X-ray and microwave senses. The order of values on this dimension are not arbitrary. It is fixed by the specification of the dimension as being one of sensitivity to lengths of electromagnetic waves.

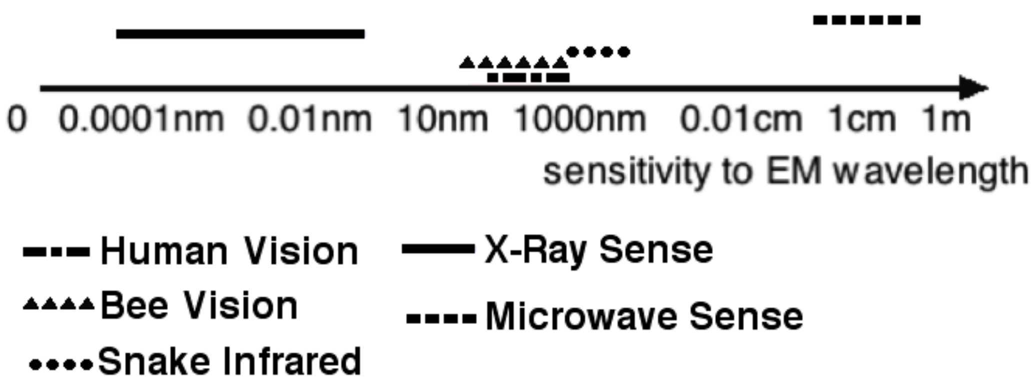

Figure 5: a dimension (not to scale) recording sensitivity to electromagnetic wavelength with different senses mapped on it. The different senses are shown at different heights above the dimension line only so that they can be clearly seen.

We can imagine joining together such dimensions to record further sensitivities of the senses to different proximal stimuli. 18. 6 provides rough illustration (although the values are not accurate). And further, we can imagine joining together yet more similar dimensions. See Figure 18.7, which consists of imagined values for different actual and non-actual senses. Thus, in such spaces we can see that we can plot not only actual senses but also possible senses, as illustrated.

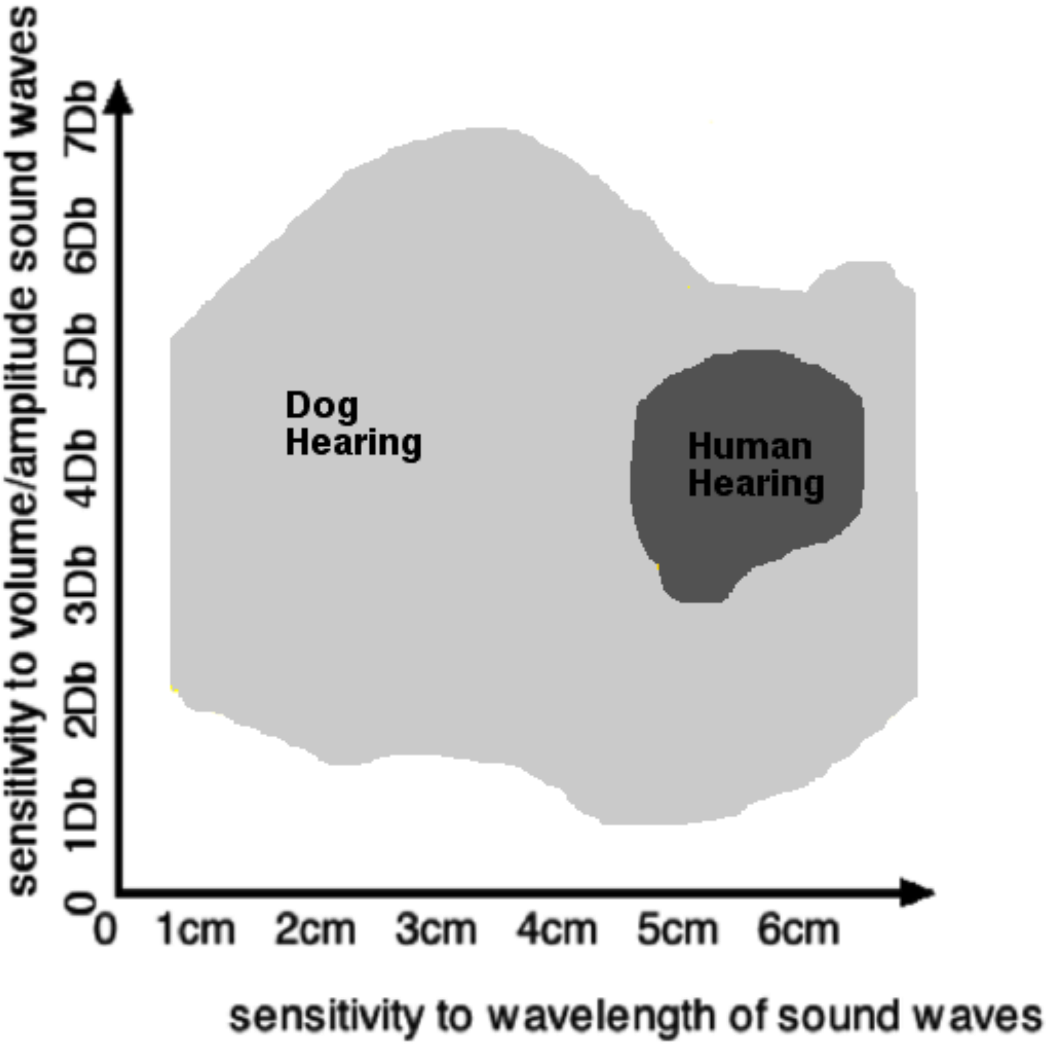

Figure 6: a fictional plotting of two different senses' responsiveness to volume and wavelength of soundwaves in air

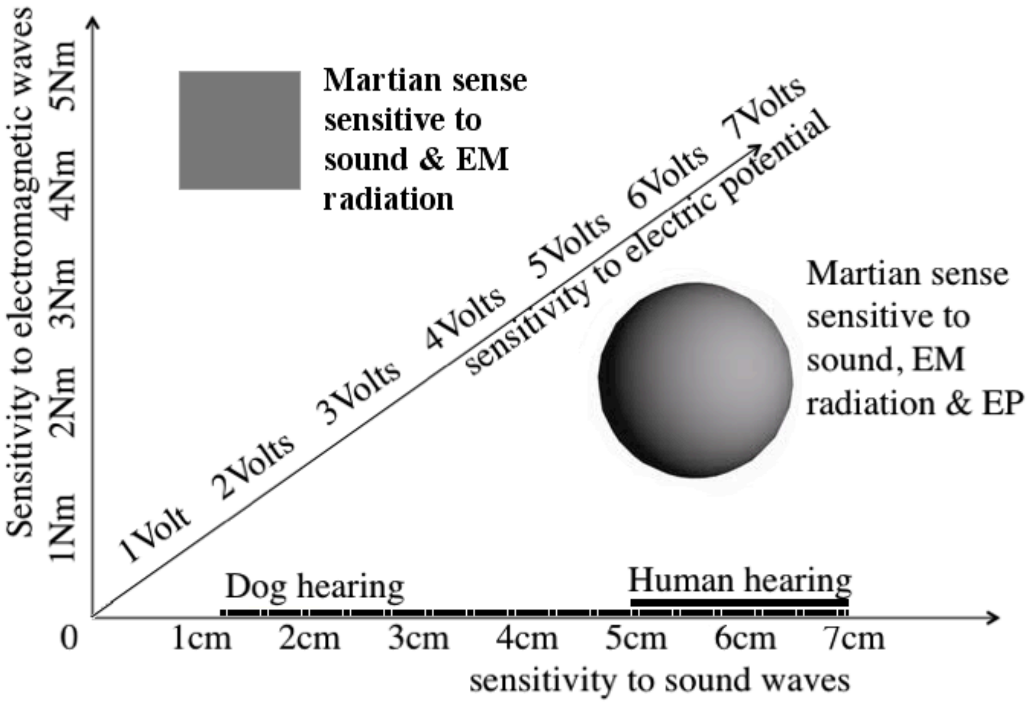

Figure 7: construction of space using dimensions corresponding to the proximal stimulus criterion and fictional values for human and dog hearing plotted together with possible alien senses

However, as indicated earlier, using these particular dimensions—ones corresponding to sensitivities to different kinds of proximal stimuli—will not work. The reason is that there is not a way to plot at least one value for each sense on these dimensions. To see this, consider the 'sensitivity to electromagnetic wavelength' dimension, along the length of which we mark out the lengths of wavelengths, as in Figure 18.8.

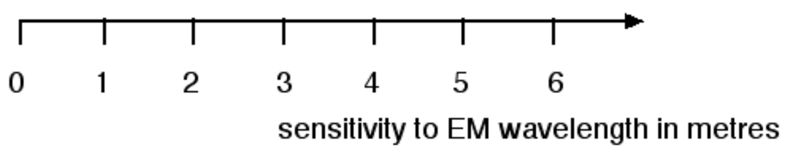

Figure 8: sensitivity to EM wavelength dimension

On such a dimension there is no way to plot a value for a sense that is completely insensitive to the presence or absence of any electromagnetic frequencies, such as smell. How do we register a total lack of sensitivity to electromagnetic wavelength? Not by plotting the sense at '0' for that means 'sensitive to 0 metres of electromagnetic wavelength', which means sensitive to the absence of electromagnetic wavelength. But smell is not sensitive to that. If it were, it would be able to register when there were no wavelengths present, rather than some. But it cannot do so. It is simply insensitive to both the presence and absence of electromagnetic waves.

To plot something in a multidimensional space, that thing must take at least one value on *each* dimension for to put it in the space is to plot its position on each and every one of the dimensions that form the space. Thus, we cannot take sensitivities to general kinds of proximal stimuli as dimensions, as this does not allow us to plot on them senses that are not sensitive at all to that kind of proximal stimulus. So, we must find other dimensions that correspond to the proximal stimulus to form our space of sensory modalities.

As we have just seen, we cannot use dimensions corresponding to sensitivity to *kinds* of proximal stimuli, but we can move towards identifying suitable dimensions by considering more specific proximal stimuli. For example, consider a dimension that records the sensitivity of a sense to a specific wavelength of electromagnetic radiation, such as that illustrated in Figure 18.9.

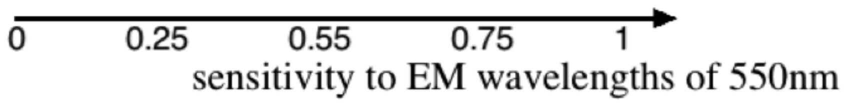

Figure 9: a dimension of sensitivity to 550 nanometre electromagnetic waves expressed as the relative response per light energy, normalized to a peak value of 1

Spectral sensitivity is often expressed as a quantum efficiency—the probability of getting a quantum reaction to a quantum of light. It is also sometimes expressed as a relative response, normalized to a peak value of 1, and quantum efficiency is used to specify the sensitivity at that peak.

On this dimension, the order of the values is not arbitrary. It is fixed by the specification of the dimension. And, importantly, every sense can take a value on the dimension. For example, smell, which is insensitive to 550nm electromagnetic waves, can take a value of zero on this dimension, for zero on this dimension means that a sense is not sensitive to this wavelength (not that it is sensitive to the absence of this wavelength). Clearly every sense can take a value on the dimension as either it is responsive or not to this wavelength: that exhausts the possibilities. And those facts can be recorded on this dimension.

We can have as many dimensions of this sort as we require—each corresponding to a different wavelength of the electromagnetic spectrum. For example, we can have one corresponding to a wavelength of 550nm, one to 551nm, one to 552nm, or one to 550.5 nm, et cetera. Having infinite dimensions may be problematic for forming a space, but we can have as large a finite number of dimensions as we desire and that we think are required to capture differences between the senses. Similarly, we can have dimensions corresponding to sensitivities to different particular lengths of sound waves, particular strengths of electromagnetic fields, particular temperatures, particular pressures, particular chemicals in solution and in the air, and so on. We can now imagine a highly multidimensional space in which all the senses take a value along each of these dimensions.

But there is a problem with this solution. In constructing a space in which there are different dimensions corresponding to sensitivities to different wavelengths of electromagnetic radiation, we have lost a vital piece of similarity information we want to capture. One attractive feature of the otherwise flawed dimension drawn in Figure 18.5, a dimension recording sensitivity to different wavelengths of electromagnetic radiation, was that it captured the fact that senses that are sensitive to similar electromagnetic wavelengths are more similar to each other than those that are sensitive to different wavelengths. And the more overlap, the more similar they are. The less overlap, and the more they respond to very different wavelengths, the more different they are. For example, human vision is more similar to bee vision than to an X-ray sense, with respect to its sensitivity to electromagnetic wavelength. On the present proposal, however, we simply have a large number of dimensions, on each of which can be recorded sensitivities to a very specific proximal stimulus, all joined together to form a space. But, as they are different dimensions joined to form a space, there is no ordering among them (as explained earlier in this section). Thus we can't capture the similarities and differences between the senses just mentioned. This is a problem, for it is precisely these kinds of similarities and differences that we wish to capture in our space.

Can this problem be solved? One might hope that other dimensions would capture these kinds of similarities and differences, but there is no guarantee of this—and the onus clearly seems to be on me to show that they can be captured. So, at the present point in the dialectic, we have before us two undesirable options: dimensions that capture similarity, but on which we can't plot all of the senses, and dimensions on which you can plot all the senses, but not their similarity.

However, there is a way forward. We can construct a space that takes sensitivities to electromagnetic wavelengths as one dimension and electromagnetic wavelengths as another. Figure 18.10 illustrates approximate actual values for the human visual and human smell systems.

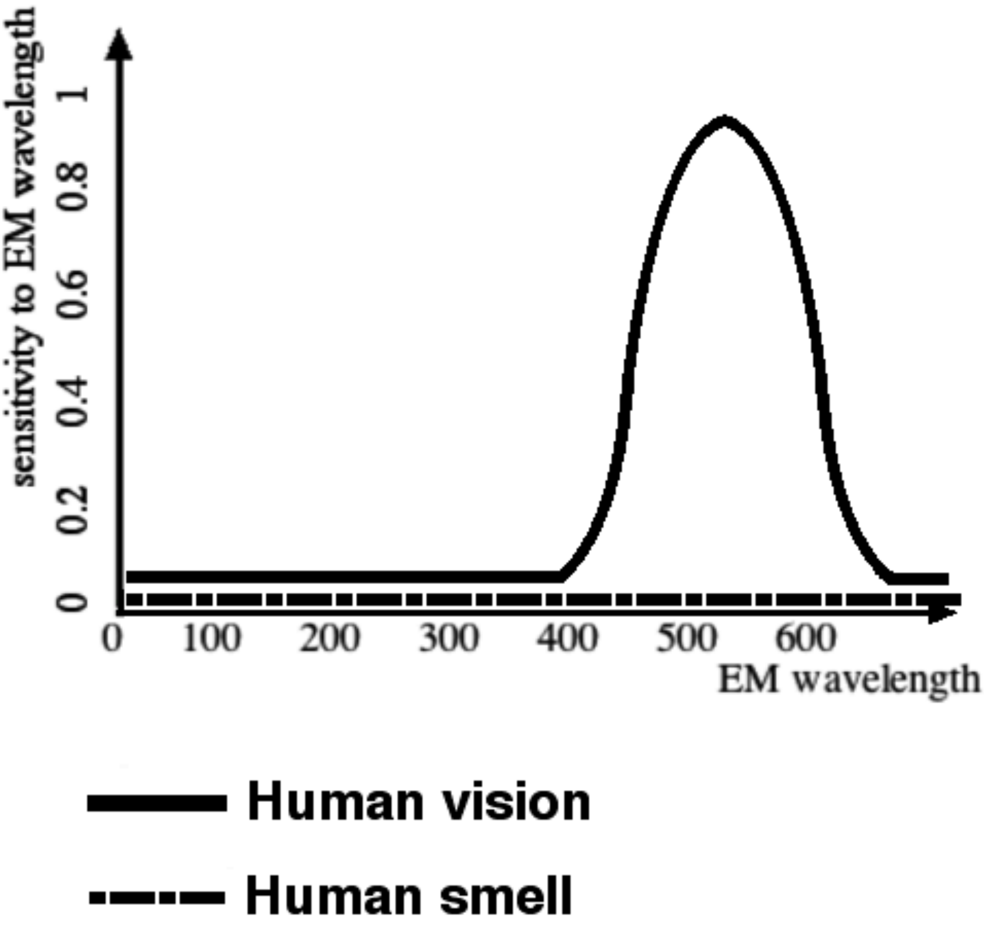

Figure 10: a space created from the dimensions of sensitivity to electromagnetic wavelength and length of electromagnetic waves

In this kind of space, we can register sensitivities to different wavelengths, preserving the order of the wavelengths, thus capturing the sort of similarities and differences between the senses noted previously. At the same time, we can record a sense, like smell, that is sensitive to no electromagnetic wavelength, including its absence. Moreover, we no longer need a vast number of different dimensions each corresponding to different wavelengths of each type of wave to form our space, for we can capture them all on one wavelength dimension for each type of wavelength. All we need are dimensions corresponding to the different kinds of proximal stimuli conjoined with dimensions corresponding to sensitivity to those stimuli.

Choosing these dimensions provides a solution. However, note that these dimensions are such that a sense cannot take a value on those dimensions if each is considered alone. It does not make sense to ask what value vision or smell takes on a wavelength dimension or on a sensitivity to wavelength dimension.

But now we might wonder whether such dimensions are permissible. We might think they are not, as previously I noted that for a dimension to be part of a space in which we can plot the senses, each sense must be able to take at least one value along each dimension. Are we now going back on that?

The answer is 'no'. It is true that each sense must be able to take a value along each dimension if they are to be plotted in a space defined by these dimensions. But it may do so *either* because the dimension is such that just in virtue of its nature alone all senses can take at least one value on it, *or* because the dimension is in a space created by joining it to another particular dimension. Thus, as long as we have both dimensions in our space, each sense can take a value on both dimensions.

One might be worried that it could be problematic to join two such dimensions that depend on each other's presence to other dimensions that we will need to form our space. But it is not. For example, we can join our electromagnetic wavelength dimension and sensitivity to electromagnetic wavelength dimension to unrelated dimensions. To see this, consider the dimension, relevant to the criterion of the nature of the sensory organs, that measures, from 0 to 100, the percentage of the sensory organ made of carbon. One could clearly join this dimension to those under consideration, and plot all the senses in such a space. See Figure 18.11 for how one would do so, supposing that the percentage of carbon in the eyes and the nose was 50%.

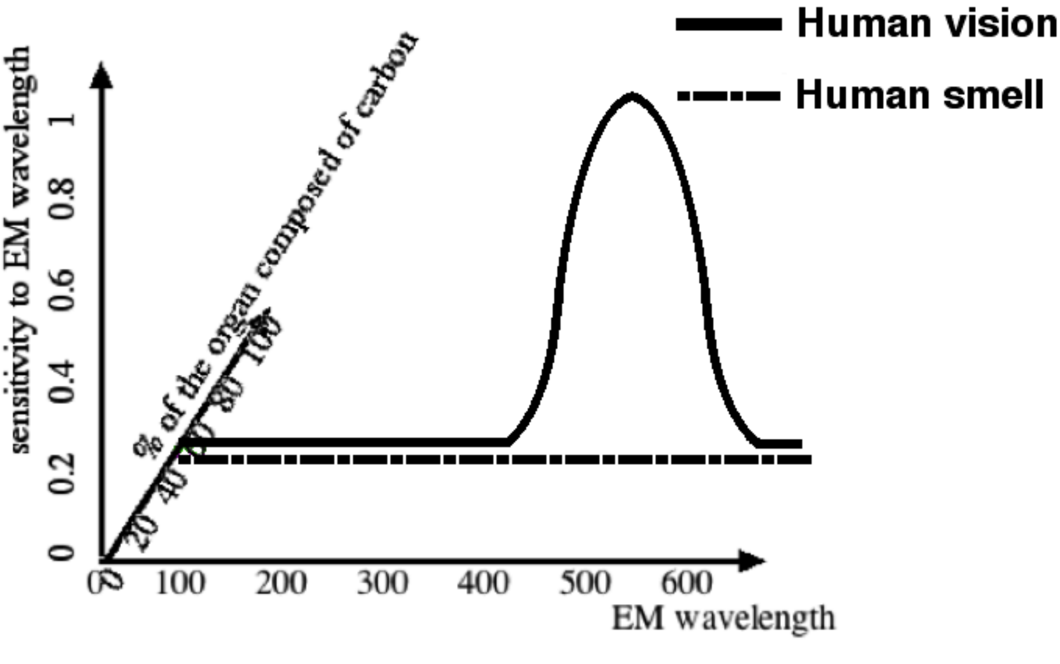

Figure 11: joining different kinds of dimensions together to form a space

And we could join our electromagnetic wavelength dimension and sensitivity to electromagnetic wavelength dimension to other dimensions that depend on one another's presence to form a viable space, such as 'length of sound-wave' and ‘sensitivity to sound-wave’. So there is no problem of making a two-dimensional space created from dimensions that depend on one another’s presence to form a viable space, a subspace of a larger space.

In this way, we can see clearly how to specify dimensions that characterise the proximal stimulus criterion that generate a non-arbitrary ordering and that capture similarities and differences between the senses. We have also found a constraint that all dimensions that form the space of the sensory modalities must conform to: each sense must be able to take at least one value along each dimension.

However, at this point, one might think that another issue concerning arbitrariness emerges. Gray (2012) raises a problem regarding the arbitrariness of using certain dimensions to form the space of the sensory modalities—dimensions such as ‘sensitivities to different frequencies of wavelengths of type X’. He says:

> Given that the dimensions are based on the values of types of energies, one suggestion might be that the spaces can be united into a single space via their origins at zero energy. However, this cannot be the right way to generate an ordering of all the senses because it would result in the paradoxical outcome that a sense that detects electromagnetic radiation of lower frequencies (for example, human or rat [vision] …) would be more similar to human hearing than a sense that detects electromagnetic radiation of higher frequencies (for example, butterfly vision).[2]

Examine Figure 18.12, in which three senses, call them A, B, and C, are plotted. Gray’s point, I think, is this: if we create a space using a dimension of sensitivity to sound wave frequency and sensitivity to electromagnetic wave frequency as axes, then the B sense, sensitive to low frequencies, turns out to be closer, and hence more similar, to the A sense than the C sense, which is sensitive to high electromagnetic frequencies. But this is arbitrary, says Gray, for there is no reason to think that the B sense, rather than the C sense, is more similar to the A sense.

[2] Although Gray says that he thinks a “paradoxical” outcome would result, in fact there would be nothing paradoxical about such an outcome, only something incorrect or puzzling.

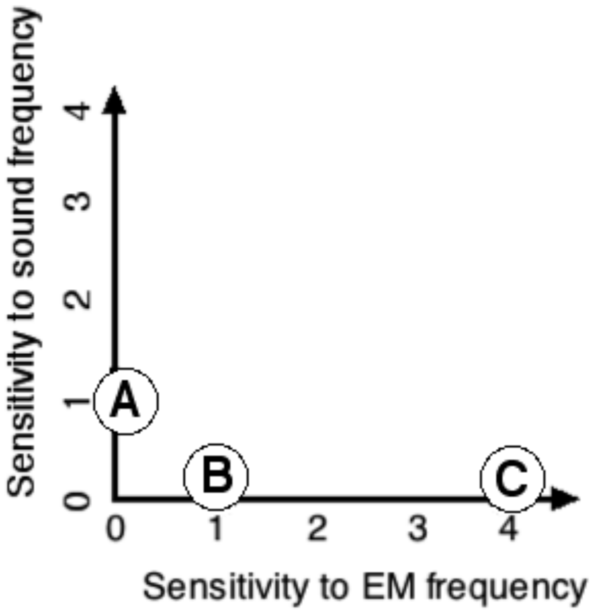

Figure 12: Plotting different senses, A, B, and C on a space created from different dimensions of sensitivities to frequencies of different types of waves

However, if this is Gray's point, then it is spurious. *If* the senses can be plotted in the space, that is, *if* they take a value on both dimensions, *then* their place in this space reflects their similarity with respect to these two values. (We should note that this does not mean that their relation in this space reflects the overall similarities and differences of the senses. That will be determined by their place in the highly multidimensional space formed by all the dimensions that we should consider.) But the similarity of the three senses with respect to the values on the two axes is exactly correct. For consider that the A sense takes a zero value on the *x*-axis so *in this respect* it is more similar to B than C. The B and C senses take a zero value on the *y*-axis, and so *in that respect* they are equally dissimilar to the A. Overall then, B is more similar to A than C *with respect to these values*.

The crucial point of the preceding paragraph is the antecedent of the conditional: 'if the sense can be plotted on the above graph'. I think that Gray is confused in part because of a point that I have dealt with already, namely, that dimensions of the sort 'sensitivity to wavelengths of different lengths of type X' or 'sensitivities to different frequencies of wavelengths of type X' are not good dimensions to have as part of the space of the sensory modalities because not all the senses can be plotted on such a dimension. For example, the A sense in Figure 18.12 is represented in this space as being sensitive to a lack of frequency of electromagnetic wavelengths—not as being insensitive to electromagnetic frequencies, as human hearing is. If we want to represent a sense that is insensitive to the presence or absence of frequencies of electromagnetic wavelengths, then we cannot do so in the space in Figure 18.12. However, we have already solved that problem. It is done by ensuring that all senses can take a value on each dimension. So long as we do this we ensure that we do not face the sort of arbitrariness to which Gray refers.

However, two other worries about arbitrariness now arise. The first comes from considering how we join the relevant dimensions together to form our space of sensory modalities. One could think of joining together all the dimensions at their 'zero' value, but one might wonder whether doing so introduces an arbitrary element, for what counts as a zero value on some dimensions is, surely, merely arbitrary, such as zero degrees on the Celsius temperature scale, or zero degrees on the Fahrenheit scale. But, in truth, no such arbitrariness is created for it does not matter which point along each dimension is chosen as the intersection point, so long as there is one common intersection point. This is because the relative distance and relative ordering between the plotted points will remain the same in all cases, wherever the dimensions are joined. One can see this illustrated in Figure 18.13. Thus there is no source of arbitrariness in the joining of the dimensions.

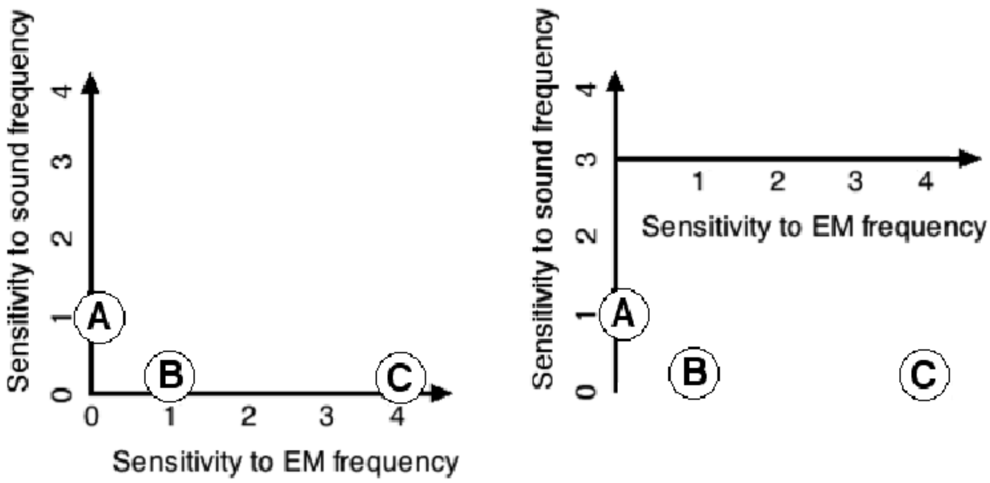

Figure 13: two identical axes joined at different points to make two spaces. Note that the distance between the points plotted is the same in both spaces.

The second additional worry about arbitrariness is, however, more serious. The worry arises because we need to decide on a scale for our dimensions. For example, suppose we have two dimensions each of which measures a different value, *x* and *y*, in units of 0, 1, 2, 3 … We could choose a scale of 1 unit per one centimetre on both the *x* and *y* axes as is illustrated in the left-hand diagram in Figure 18.14. However, we could also have chosen a scale of 1 unit per one centimetre on the *y*-axis and 1 unit per two centimetres on the *x*-axis, as in the right-hand diagram in Figure 18.14. It seems arbitrary which scale we use for the *x*-axis. However, which one we choose will affect the relative similarities and differences between the points plotted in the spaces. For example, the B and A points remain one centimetre apart in both the left-hand and right-hand space. But the distance between the B and the C points is larger by one centimetre in the right-hand diagram, compared to the left-hand one. Thus, in the left-hand space, B is represented to be as similar to A as it is to C, whilst in the right-hand space B is more similar to A than it is to C. Surely, then, the choice of scale introduces an arbitrary factor that incorrectly determines how similar and how dissimilar different senses will be?

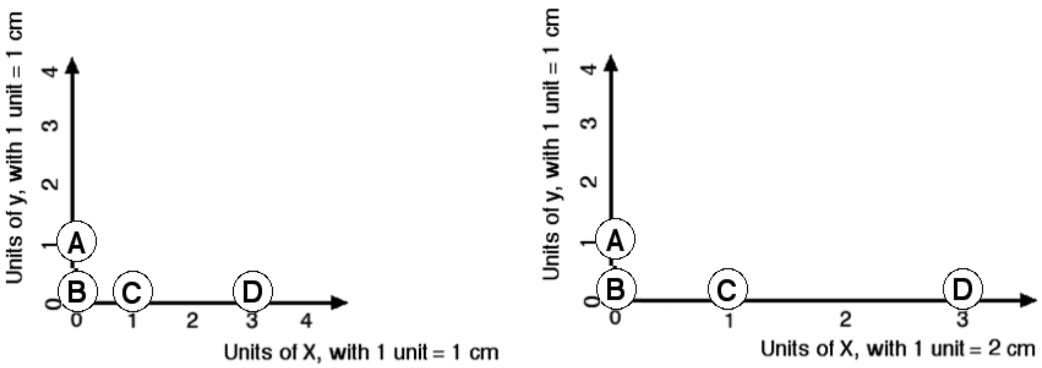

Figure 14: two spaces, with points plotted on them, which differ only in the scale of the axes

There are three points to be noted that alleviate this worry. The first is that the ordering of points along a dimension will not be affected by issues of scale. Thus, if in the left-hand diagram of Figure 18.14 C is between B and D on the *x* dimension, then this remains the case in the right-hand diagram when we change the scale, and will do so for any change of scale.

The second point is that when scientists create a space in which to plot things by specifying dimensions, they are always faced with choices regarding the scales that they should use, particularly when the units of the dimensions are different. That we face such choices does not reflect a deficiency in our method of space building but reflects mandatory features of such an enterprise that is to be expected. When scientists have no intuitions or a priori knowledge about the relationships between things to be recorded on more than one dimension then scaling to unit variance (UV) is the standard practice (Eriksson et al. 2006). If this was not done, a dimension that measured a quantity on a scale of 1–100, with each unit represented by one centimetre, would have more influence than a dimension that measured a quantity on a scale of 1–10, with each unit represented by one centimeter, on how similar or how different things were that were plotted in a space created by conjoining those dimensions. Scaling to UV takes the units along a dimension and places them on a scale between zero and one, so that '"long" variables are shrunk and "short" variables are stretched so that all variables will rest on an equal footing' (Eriksson et al., 2006: 208). The scaling weight employed is one divided by the standard deviation of the variable. Such a scaling 'is the most objective way to scale the data', giving each variable an equal opportunity to influence the data analysis. In particular, 'UV-scaling is useful when variables are of different kinds and not directly comparable numerically.' (Eriksson et al., 2006: 208)[3]

[3] This technique is also the standard technique to be done to data prior to principal component analysis, described in section 6 together with mean centering in which 'the average value of each variable is calculated and then

The third point is that if one has intuitions or knowledge about how different dimensions contribute differently to the relative similarities and differences of the senses, then one can build these into one's model. One way to proceed is to start building the space by selecting one dimension with a particular scale. One can then choose the scale of a second dimension to be added to the space in a way that reflects the sorts of similarities and differences that one believes would hold between the things plotted on them. For example, if it seems to you that two senses that differed by one unit on the *y*-axis would be just as similar as things that differed by 5 units on the *x*-axis, then one should set up one's scales to reflect these facts. One could then add a third dimension by comparing it to either the first or the second dimension or to both.

Now a worry arises with this methodology. It is a worry about circularity. One might think that the point of constructing the space of sensory modalities was to discover the similarities and differences between the senses, not just to map out what one's intuitions are. But if we use our intuitions or knowledge about the relationships between the senses to create it, then we are in danger of not finding out anything that we did not know from our model building.

One way forward would be to not use any of one's knowledge and intuitions and simply to give equal weighting to the dimensions using the UV methodology described three paragraphs earlier. However, one could still gain fresh insights into the nature of the senses in two ways even if one used the intuitions and knowledge that one had. First, one could build a model where some dimensions are given an equal and otherwise neutral weighting as per the UV method, while others are given a more substantial or less substantial weighting relative to that, in accord with one's intuitions. Eriksson et al. say

> if prior information about the importance of variables is available, this insight should be reflected in the scaling. Variables that are known to be 'important', for a given problem can be up-weighted, to at least partially account for this relevance. Analogously, variables that are noisy, or known to be irrelevant for the problem at hand, can be assigned a lower weight to diminish their contribution to the modeling. (2006: 208)

Further, consider that one could have an intuition of how similar and different a sense was with respect to the values on two dimensions. One might also have an intuition of their similarity and difference with respect to the values on one of those dimensions and a third, and an intuition of their similarity and difference with respect to the values on the third and a fourth. However, one may have no intuition of their similarity and difference with respect to the values on the first and the fourth. But by building a space that reflected these intuitions would one be able to discover their similarity and difference with respect to all the dimensions.

---

subtracted from the data'. In effect, this simply places the intersection of the dimensions at the mean point of the data plotted on each dimension. It is said that '[t]his improves the interpretability of the model [after PCA]' (Eriksson et al. 2006: 43). I discuss the question of interpreting the dimensions of the model after PCA in section 6.

The second way to gain knowledge even by inputting one's intuitions or knowledge into building the space is by performing an analysis of the space that may reveal new dimensions that individuate the senses and which may lead us to revise, reconsider, and hone the scales that we have chosen. Further details of this analysis (principal component analysis) and how it can be used in this way are given in section 6.

In the next section, I consider whether we can identify likely dimensions that correspond to the other criteria, besides the proximal stimulus criterion: the sensory-organ, the representation, and the phenomenal character criteria.

5. Dimensions Corresponding to the Other Criteria

Now that I have given indication of how to construct the dimensions corresponding to the proximal stimulus, I consider how to construct dimensions corresponding to the sensory organ criterion. I believe that the dimensions corresponding to the proximal stimulus may be a subset of the dimensions that correspond to the sense organ criterion. This is because which proximal stimulus an organ is sensitive to will be a property of the sense organ. Figure 18.15 illustrates this.

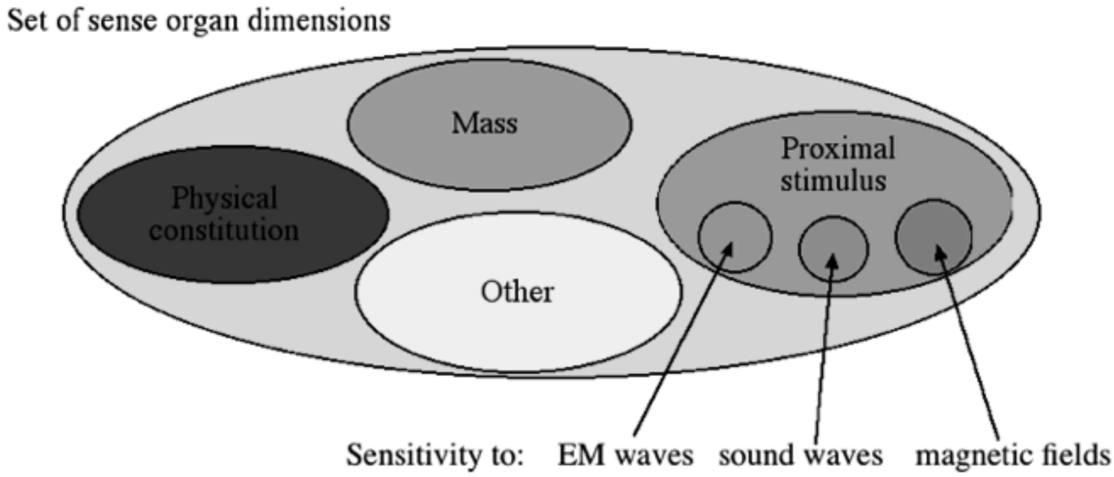

Figure 15: the set of dimensions corresponding to the sense organ criterion

Thus we already have an idea of some of the dimensions that correspond to the sense organ criterion. But what about sense organ dimensions other than those corresponding to the proximal stimulus?

In the previous section, I mentioned the possibility of having a dimension corresponding to the percentage of the sense organ made up from carbon. Such a dimension creates a non-arbitrary ordering on which every sense organ can take a value. We could create similar dimensions for each of the elements, and, indeed, for molecules or any other important physical feature that we deem to be relevant. See Figure 18.16.

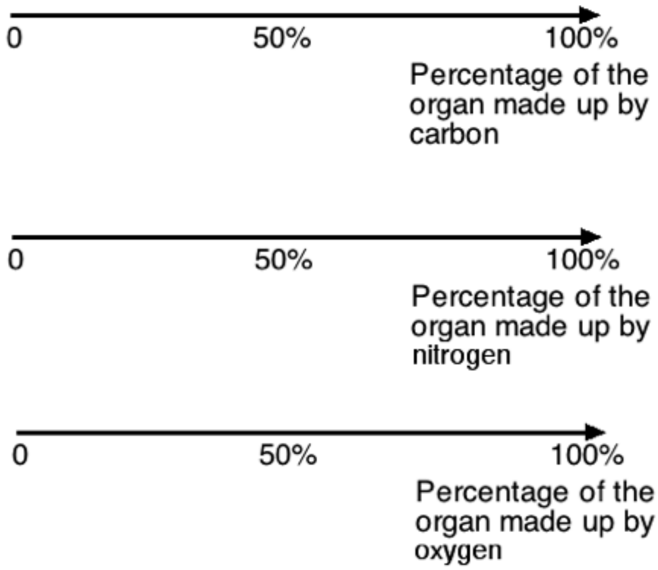

Figure 16: some dimensions of composition of the sense organ

We could also capture the three-dimensional shape of the sense organ. We could create spatial dimensions centred, say, on the centre of mass of the organ. See Figure 18.17. These dimensions establish a non-arbitrary ordering that reflects a real property of the senses. Each sense can be plotted on such dimensions, supposing only that it has some sense organ. In fact, there has been some dispute as to whether we must hold that each sense has its own dedicated sense organ. For example, some people have thought that there is a moral sense that has no sense organ associated with it. However, even if there was such a sense, it is reasonable to think that there are brain areas that are active when such a sense is operative. I think it would be reasonable to record the spatial nature of these brain areas on the spatial dimensions. Although such brain areas would not constitute a sense organ in the traditional sense, as they contain no immediate world-facing sensors, nevertheless, because sense organs are increasingly thought of as being constituted not just by external facing organs such as eyes and ears, but also by the nervous system and areas of brain associated with processing information from these organs, it is reasonable to map the associated brain areas on the spatial dimensions proposed to capture what physical structures in the body are associated with a particular sense.

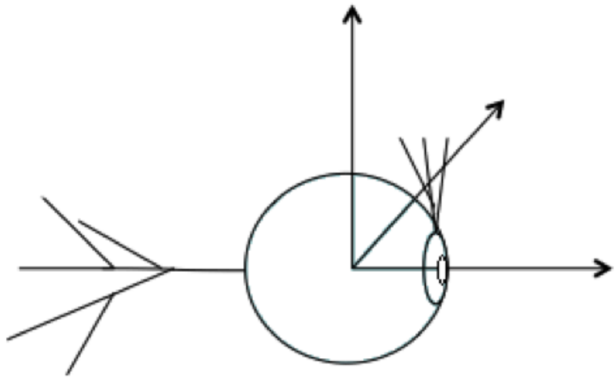

Figure 17: an example of spatial dimensions
centred on the mass of an eye-like sense organ

We could also set up dimensions that capture how much of the world a sense is responsive to. For example, as many birds have eyes that point in directions more disparate than our own, they have a much wider field of view than humans, whose eyes both point forward. In an ego-centric framework, typical human vision is just under 180 degrees from left to right, that is, horizontally, and about 100 degrees from top to bottom, that is, vertically. The vision of some birds is nearly 360 degrees horizontally and around 180 degrees vertically. Human hearing is typically sensitive to sounds 360 degrees all around us. Smell is rather limited in its field of perception—at least on some views of smell according to which we can only detect things at our nostrils. (On this view, to the extent that we can determine that a smell is emanating from the left or the right, it is by inference from a series of 'snapshot' smells of different levels of intensity detected in different places in space.) On such views, smell would be plotted as having very small or no degrees of sensitivity in both horizontal and vertical directions. The plotting of these points in illustrated in Figure 18.18.

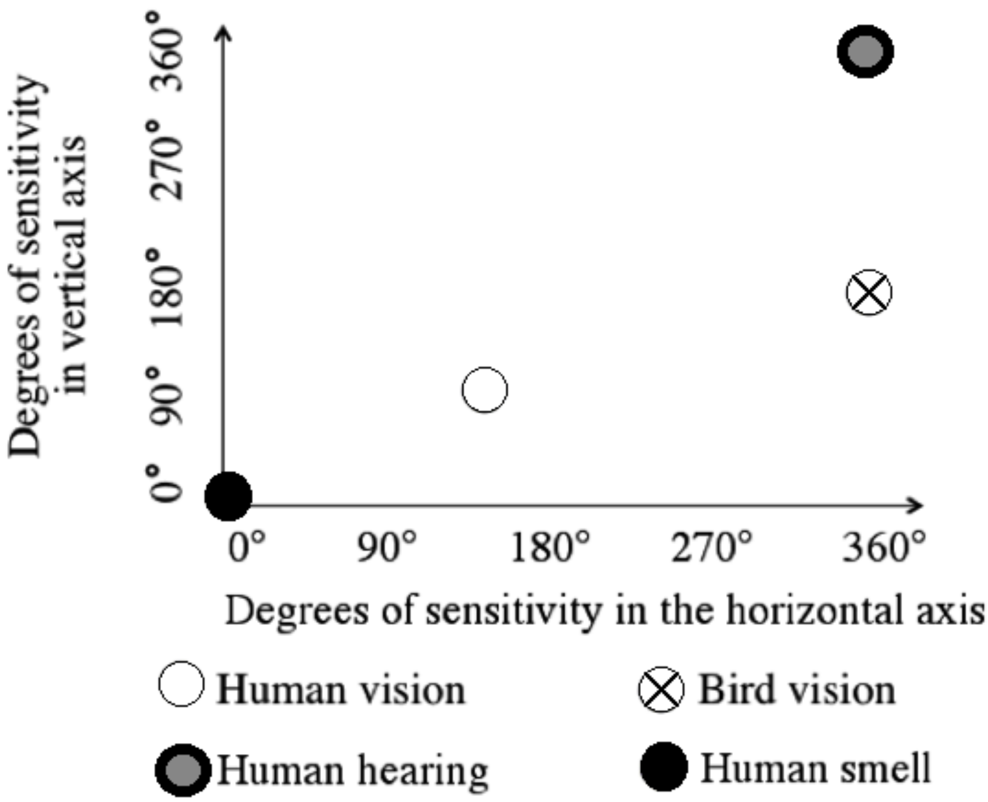

Figure 18: field of view plotted for different senses

One could also use dimensions that captured the following properties of sense organs, which would impose a non-arbitrary ordering, and on which each of the senses could be guaranteed to take a value: mass of the organ, number of neurons composing the organ, number of temperature cells in the organ, number of pressure cells in the organ, and so on. Thus, one can see that it will be relatively straightforward to come up with the dimensions corresponding to the sensory-organ criterion.

When it comes to the representation and phenomenal character criteria, those who suppose that these two properties are really the same will claim that there is only one set of dimensions corresponding to these two criteria. Whether this is true or not, finding dimensions corresponding to these dimensions will be, I predict, possible with suitable ingenuity. Finding suitable dimensions for the proximal stimuli and sense organ criteria should, I believe, give us hope. I end this section with just one example of how one might specify some phenomenal character dimensions.

Consider the colour space created by asking subjects to rank numerous instances of three colour patches *x*, *y*, and *z* in such a way as to indicate whether *y* or *z* is more similar to *x*, and thereby ordering the patches. We find that such a space can be characterised by three dimensions: lightness, hue, and saturation. White and black are antipodal points, with shades of grey running along the straight line connecting them. These are the achromatic shades. All the other shades have some chroma. Particular shades of red and green, and blue and yellow, respectively, are also antipodal points. See Figure 18.19 for illustration.

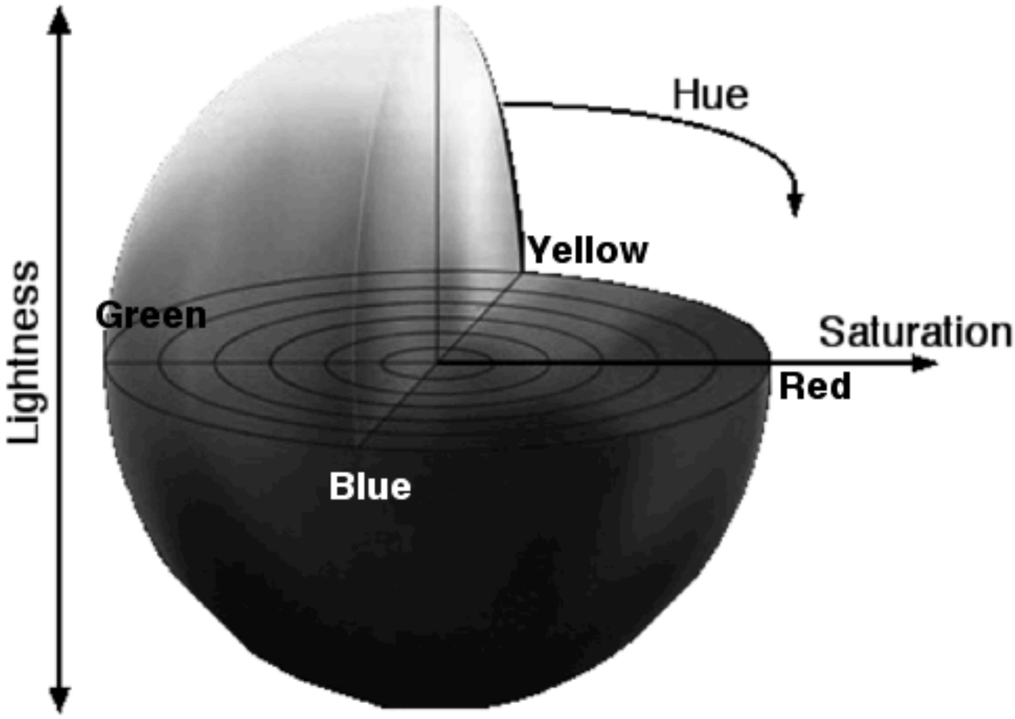

Figure 19: The colour space - a roughly spherical shape illustrated here with a section removed to facilitate visualisation

To create a phenomenal dimension, consider the straight line in Figure 18.19 leading from the centre of the sphere out to a shade of red, the line labelled as the saturation axis. The point at the start of the line is a shade of grey. All the other points will be increasingly saturated shades of a particular hue of red. Now consider a slightly different line. This one is very similar to the previous one except that the point at the start of the line, which in the previous dimension corresponded to a shade of grey, corresponds simply to a lack of one of the varying shades of differently saturated shades of red. We can imagine taking this as a dimension and plotting whether a sense yields an experience with the phenomenal character that one gets when one looks at the various shades of red along the dimension. Normal human vision will occupy all the points along the line. If there is a sense, such as smell, that yields no experience as of this hue of red, then we can plot it at the zero point that indicates a lack of all such phenomenal character. In this way, we can ensure that every sense can be plotted on this dimension. All senses either produce one or more of these phenomenal characters or none of them. We can imagine creating similar phenomenal dimensions for the other colours and for other phenomenal qualities.

To summarise this section, the space of the sensory modalities will be created by joining together all the dimensions corresponding to each of the four criteria at a common origin. Each dimension in the space of sensory modalities must be such that each sense can take at least one value on the dimension, and the nature of the dimensions non-arbitrarily fixes the position of the senses in that space. I have indicated likely dimensions that we might use to construct such a space that corresponds to the proximal stimulus, and a few for the sensory-organ and phenomenal character criteria. I believe that with some thought we could go on to identify all that would be required. The goal of this chapter is not to identify all the dimensions that one would need to construct the space of the modalities: rather, it is to alleviate concerns about whether it is in principle possible to construct the space of sensory modalities. The main lesson of the chapter so far is that if there is a dimension along which one wants to record the nature of senses, but that dimension is such that it looks as if placing values along it will be arbitrary or in some other way unsuitable, there are ways of capturing what one wanted to record on that dimension by using instead a combination of other dimensions that are not problematic. I have also suggested ways of constructing dimensions corresponding to aspects of the senses that one might have thought difficult to build.

In the next section I discuss a very powerful technique that allows one to extract information from the space of sensory modalities that I have been considering building, and at the same time refining it.

6. Refining the Space of Sensory with Principal Component Analysis

In this section, I discuss principal component analysis (PCA) and how we can use it to analyse the space of sensory modalities created in the manner indicated previously. Doing so may lead us to refine and alter the space in which we wish to plot the senses. It may yield interesting and important insights into the factors important for individuating the senses and may make us reconsider what we should count as being senses. It may also make us reconsider what we should count as being token instances of senses. In addition, PCA allows us to reduce the dimensions of a multidimensional space, creating spaces with as many dimensions as we please—including a three-dimensional space in which we could visualise the position of the senses.

The space of sensory modalities that we have been considering creating will have many, many dimensions. The idea behind the space is that plotting the senses with respect to their similarity and difference on all the dimensions we consider important will produce a space that captures the similarities and differences between the senses. We might feel that such a highly multidimensional space is in practise of very little use. It does not help us visualise the relationship between the senses because of the high number of dimensions, and we might be concerned about how much new information we will be able to glean from the space that we did not explicitly put into it.

Fortunately, however, PCA is a powerful technique developed to analyse multivariate spaces and it promises to yield interesting results. It was first developed by Pearson (1901) but has only come to the fore in the last three decades with the arrival of the massive computing power required to handle the computations on the large data sets that it was designed to analyse.

PCA is a mathematical procedure that is used to analyse a high dimensional space in which data are plotted. It can be used to produce a number of different things when given a multidimensional space as input:

(1) a space with as few dimensions as possible that still retains all the variability in the data. While the input space may have been constituted by dimensions of correlated variables, the new space will now contain dimensions of linearly uncorrelated variables, called principal components. The number of the dimensions may remain the same, be reduced slightly, or be substantially reduced. The dimensions are also ordered, in the following sense: the dimension that accounts for the most variability in the data is identified as such, the dimension which accounts for the next most variability is identified as such, and so on for all the dimensions of the space.

(2) a space with a specified number of dimensions that retains as much of the variability of the data as it is possible to do with that limited number of dimensions. The new space will now contain dimensions of linearly uncorrelated variables that are ordered as just described. The amount of variability in the data that is retained can vary but, typically, due to the nature of the technique, a very large substantial proportion is retained while reduction to a few dimensions is achieved. The amount of variability retained can be measured.

(3) a two- or three-dimensional model that provides a visualisation of the space and the data plotted in it. PCA is often used precisely for visualisation purposes.

Thus, PCA is a technique of exploratory data analysis that identifies hidden patterns in the data of a high multidimensional space and classifies the newly produced dimensions according to how much of the information, stored in the data, they account for. This means that the dimensions of the newly created space are typically not the dimensions that constitute the space that is analysed.[4]

To illustrate with a simple example of a PCA, consider Figure 18.20. The points plotted in the space given by the dimensions D1 and D2 can be re-plotted in a space with dimensions D1* and D2*, where D1* and D2* are uncorrelated dimensions, and D1* is the dimension that accounts for most variability between the items plotted, and D2* the next most. Moreover, the space created by D1* and D2* preserves all the similarity and difference information as was in the space created by D1 and D2. Also, the dimensions D1* and D2* represent different values from those represented by D1 and D2.

[4] PCA is described at length in Eriksson et al. (2006).

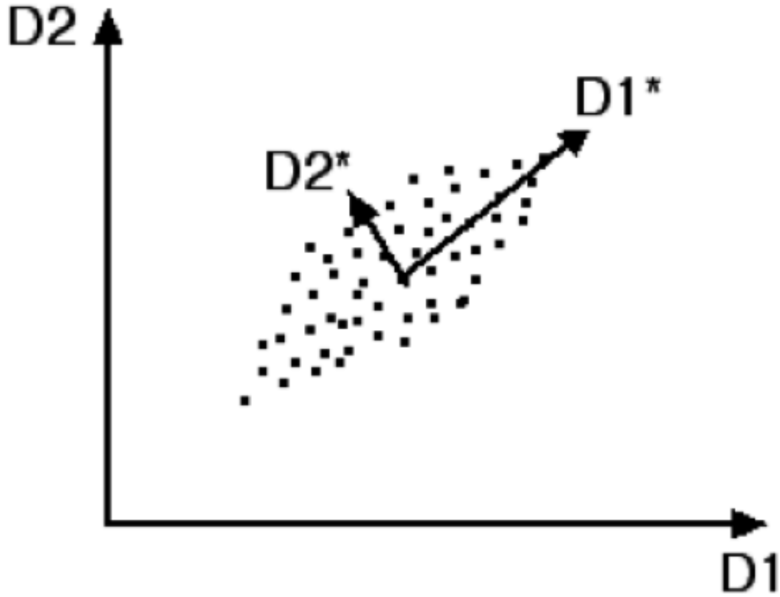

Figure 20: an illustration of the results of a PCA analysis of a space with points plotted in it

In the space of sensory modalities that I have suggested creating, we could plot various categories of senses. For example we could plot the following senses:

- all human senses
- all actual senses
- the senses that folk-psychology recognises
- all nomologically possible senses
- all metaphysically possible senses.

Then, using PCA, we could work out the smallest number of dimensions needed to capture all the similarity and difference information for each of these groups of senses.[5] This technique is used all the time in science for taxonomical purposes.

[5] In fact, you could plot whatever group of senses interests you. For example, someone might think that creating the space of all the sensory modalities is too difficult and of little value and may just want to create the space of visual senses: human vision, octopus vision, bee vision, and so on. As I stated in the first section of the chapter, I've argued elsewhere (Macpherson 2011a and 2011b) that I don't think the senses fall into a small discrete number of natural kinds like vision, audition, and so on. I think there is no non-arbitrary answer to whether snake infra-red, bee 'vision', X-ray 'vision', bat echo-location, or the sense used in tactical visual sensory substitution should count as vision or not. I believe that all the senses are simply more or less different along a number of different

To illustrate with just one of many examples, consider Crescenzi's and Giuiani's (2001) construction of their taxonomy of tumours. They first selected 1,416 dimensions corresponding to different gene expression levels—properties that tumours may have. They then conjoined these dimensions to form a high-dimensional space and plotted 60 tumours in it. They then performed a PCA and discovered that only five dimensions were required to explain the vast majority of the variance among the tumours. These five dimensions were different to each of the dimensions that formed their original highly multidimensional space. Interestingly, and importantly, they could *interpret* the new dimensions and work out what characteristics of tumours they corresponded to—characteristics that they had not considered prior to doing the PCA.

Crescenzi and Giuiani state a fact that is always noted about this technique, 'From the interpretive point of view, the major merit of PCA resides in the possibility to attach a biological meaning to the components' (2001: 117). In other words, when one performs a PCA, one finds the principal components, which can turn out to be new dimensions—dimensions different from those that one puts into the space on which one performed the analysis. These new dimensions can often be given an interpretation in terms we are familiar with, so that they turn out to represent biological, physical, or psychological properties that we had not previously considered.

Importantly, when performing PCA, scientists often tinker with

- which dimensions to use to form the initial space on which the PCA is performed,
- the scale of the dimensions, and
- the examples to plot,

until PCA

- reduces the dimensions greatly while not losing too much similarity and difference information,
- produces dimensions that account for much variability in the data that were not used to construct the original space, and
- produces dimensions that can be identified with biological, physical, psychological, or other natural or scientific properties.

Hence, PCA can yield insights into whatever is being taxonomised.

Note that the issue of which scale to use for the dimensions is included in what one might alter to produce an interesting analysis. As mentioned in section 4, this is one way, in addition to others mentioned there, in which choices about the appropriate scales to use for dimensions can be made.

---

dimensions. Thus, I don't see any value in trying to plot a subset of senses, like those someone might be tempted to call the 'visual' ones. However, if others were to disagree, they could nonetheless plot whichever senses they liked using the techniques I've described in this chapter.

Tinkering with the PCA in the way just described might include toying with which things we take to be senses, and what properties we believe something must have to be a sense, for excluding or including certain cases might produce the best analyses. For example, does excluding or including the interoceptive senses alter the analysis? Does excluding vomeronasal reception produce a more simple or elegant PCA? We might also consider whether to count, for example, touch as one sense or as multiple senses, for it might turn out that we produce the best analysis if we treat it as one, rather than three separate senses of pressure, temperature, and pain. We might also see whether it is best to do the analysis on types of normally functioning senses in each species, or whether to plot for token individual normal and non-normal instances of senses.

The PCA might also yield interesting results about which properties really matter for individuating the senses. For example, it might turn out that a PCA allows one to identify a small number of dimensions that are required to characterise the space of sensory modalities while retaining most of the variability data, and that we can interpret these dimensions as corresponding to certain properties of the senses. For all we know, we might even end up with four dimensions, the best interpretation of which might be a property of each of the four criteria that we started with, namely, the proximal stimulus, the sense organ, the representational, and the phenomenal character criterion. Thus, it is possible that we might end up with a space with four dimensions corresponding to properties of the four criteria that we initially might have thought formed the space of sensory modalities. Or we may end up with an elegant three-dimensional space that allows us to visualise the space of the sensory modalities. We can certainly forcibly reduce the number of dimensions to three to see such a space while noting how much, if any, variance data we lose when we do so.

What exactly we will find remains to be seen. However, let me emphasise most strongly that tinkering with the model on which one conducts the PCA is often done *until* a useful analysis is yielded in which the principal components can be identified with certain biological, physiological, or psychological properties.

This kind of tinkering with data sets with respect to which examples to plot, which dimensions of which scales to use, the number of dimensions to which one will reduce the space, and the tolerance for data variance loss is often the subject of debate within a community of taxonomists spanning many years. See, for example, Noy-Meir and Whittaker (1977) who summarise and contribute to the debate concerning patterns of variation in the composition of species of vegetation at different places.[6]

---

[6] Using different axes, or plotting different points, one might end up with two (or more) spaces that each is reducible to a small number of axes using PCA while not losing much information. One might find that the dimensions in the different spaces are identified with different properties. In such a situation arguments would have to be had about which space was preferable and why—as occurs at present in certain cases of taxonomy such as the one concerning patterns of vegetation. However, both sets of properties identified with the two different

In short, when constructing a space of the sensory modalities, one hopes to produce a space in which the senses are plotted that reveals more about the relationships between the senses than one explicitly put in. The high-dimensional space that we have considered creating will contain such information. However, PCA is an incredibly powerful tool that is likely to allow us to extract more and different information still—an extra deep layer of information hidden within that created when building the high-dimensional space. And results of PCA can feed back to inform the sorts of choices that we should be making when choosing the initial dimensions of the space of sensory modalities and the senses to plot within it.

7. Summary

I believe I have laid out, in a programmatic way, what would be required to create a space of the sensory modalities. And I have shown that completion of such a task, while large and difficult, should be possible.

For each of the four criteria that we should use to individuate the senses—the proximal stimulus, the sense organ, the representational, and the phenomenal character criterion—we need to specify a set of dimensions that captures the nature of each. These may overlap, some may be subsets of the others, et cetera. We then use these dimensions to form a space of the sensory modalities that captures the similarities and differences between them. However, in order to ensure that these dimensions really form an appropriate space we need to choose dimensions that do not generate an arbitrary ordering and ensure that each sense can take at least one value on each dimension. We need to consider the scale that should be used for the dimensions. We can blindly suppose that they all contribute the same to the space and scale them to unit variance. Or we can deviate from this and feed in our intuitions about the relative significance of the dimensions, and feed in our limited intuitions or knowledge about how similar the senses are with respect to a limited number of dimensions.

I have tried to show that it is plausible to think that we can find suitable dimensions in the case of the proximal stimulus criterion and, to a lesser degree, in the case of the sense organ, phenomenal character and representational criteria. These dimensions will then define a space in which the actual senses can be plotted and where at least many (nomologically and metaphysically) possible senses could be plotted too. Such a space should reveal the similarities and differences between the senses.

When we have created such a space, we can then carry out a principal component analysis on it, in the hope, and reasonable expectation, that such an analysis will lead us to refine the space, reduce the dimensions, and identify new properties crucial for individuating the sensory modalities, and give us further insight into the relationships between the senses.

---

spaces could be ones that are philosophically illuminating when considering the nature of the senses.

It is interesting to consider whether construction of a multidimensional space, and a principal component analysis of it, could help to determine the nature, individuation, and taxonomy of other objects, states, events, and properties that philosophers are interested in. To my knowledge, to date, philosophers have not used this methodology. However, I predict that there are cases, in addition to that of the sensory modalities, in which it will be illuminating.

Acknowledgements

I would like to thank Clare Batty, Jon Bird, Adam Rieger, and Barry Smith for discussing the ideas in this paper with me, and Craig French, and the editors of this volume, for helpful comments on a draft of the paper. This work was supported by two grants from the Arts and Humanities Research Council (grant numbers AH/I027509/1 and AH/L007053/1).

References

Classen, C. (1993) *Worlds of Sense: Explaining the Senses in History and across Cultures*. London: Routledge.

Craig, A. D. (1996) 'Pain, Temperature, and the Sense of the Body', in O. Franzén, R. Johansson, and L. Y. Terenius (eds.) *Somesthesis and the Neurobiology of the Somatosensory Cortex*. Basel: Birkhäuser, 27–39.

Crescenzi, M. and Giuiani, A. (2001) 'The main biological determinants of tumor line taxonomy elucidated by a principal component analysis of microarray data', *FEBS Letters*, 507:114-118.

Eriksson, L., Johansson, E., Kettaneh-Wold, N., Trygg, J., Wikström, C., and Wold, S. (2006) *Multi- and Megavariate Data Analysis, Part I,* 2nd revised and enlarged edition. Umeå: Umetrics AB.

Gray, R. (2012) 'Is There a Space of Sensory Modalities?' *Erkenntnis,* published online first, DOI 10.1007/s10670-012-9409-0.

Macpherson, F. (2011a) 'Individuating the Senses', in F. Macpherson (ed.), *The Senses: Classical and Contemporary Readings*. Oxford: Oxford University Press, 3 - 36.

Macpherson, F. (2011b) 'Taxonomising the senses', *Philosophical Studies*, Vol. 153 (1): 123-142.

Nagel, T. (1974) 'What is it Like to Be a Bat?', Philosophical Review, 83: 435–450.

Noy-Meir, I. and Whittaker, R. H. (1977) 'Continuous multivariate methods in community analysis: Some problems and developments', *Vegetatio*, 33(2-3): 79-98.

Nudds, M. (2004) ‘The significance of the senses’, *Proceedings of the Aristotelian Society*, 104(1): 31–51.

Pearson, K. (1901) ‘On Lines and Planes of Closest Fit to Systems of Points in Space’. *Philosophical Magazine,* 2(11): 559–572.

Wetzel, L. (2011) ‘Types and Tokens’, E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy (Spring 2011 edition)* http://plato.stanford.edu/archives/spr2011/entries/types-tokens.